




DEEP SPACE 1


Technology
Validation
Reports




Deep Space 1 Technology Validation Reports

Reference Information

Report	Reference
Results from the Deep Space 1 Technology Validation Mission	<i>Acta Astronautica</i> 47 , p. 475 (2000)
Autonomous Optical Navigation	JPL Publication 00-10, No. 1
Beacon Monitor Operations	JPL Publication 00-10, No. 2
Ion Propulsion System	
NSTAR	JPL Publication 00-10, No. 3A
Diagnostic Sensors	JPL Publication 00-10, No. 3B
Ka-Band Solid State Power Amplifier	JPL Publication 00-10, No. 4
Low Power Electronics	JPL Publication 00-10, No. 5
Multifunctional Structures	JPL Publication 00-10, No. 6
Plasma Experiment for Planetary Exploration	JPL Publication 00-10, No. 7
Power Actuation and Switching Module	JPL Publication 00-10, No. 8
Remote Agent Experiment	JPL Publication 00-10, No. 9
Scarlet Solar Array	JPL Publication 00-10, No. 10
Small Deep Space Transponder	JPL Publication 00-10, No. 11





Results from the Deep Space 1 Technology Validation Mission

Copyright © AIAA 1999 - Printed with permission.

Marc D. Rayman
*Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California 91109*



RESULTS FROM THE DEEP SPACE 1 TECHNOLOGY VALIDATION MISSION

Marc D. Rayman, Philip Varghese, David H. Lehman, and Leslie L. Livesay
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109 USA

Launched on October 24, 1998, Deep Space 1 (DS1) is the first mission of NASA's New Millennium program, chartered to validate in space high-risk, new technologies important for future space and Earth science programs. The advanced technology payload that was tested on DS1 comprises solar electric propulsion, solar concentrator arrays, autonomous on-board navigation and other autonomous systems, several telecommunications and microelectronics devices, and two low-mass integrated science instrument packages. The technologies were rigorously exercised so that subsequent flight projects would not have to incur the cost and risk of being the first users of these new capabilities. The performances of the technologies are described as are the general execution of the mission and plans for future operations, including a possible extended mission that would be devoted to science.

INTRODUCTION

NASA's plans for its space and Earth science programs call for many scientifically compelling, exciting missions. To make such programs affordable, it is anticipated that small spacecraft, launched on low-cost launch vehicles and with highly focused objectives, will be used for many of the missions. To prevent the loss of capability that may be expected in making spacecraft smaller and developing and operating missions less expensively, the introduction of new technologies is essential.

With many spacecraft carrying out its programs of scientific exploration, NASA could accept a higher risk per mission; the loss of any one spacecraft would represent a relatively small loss to the program. Nevertheless, the use of new technologies in space science missions forces the first users to incur higher costs and risks. The concomitant diversion of project resources from the scientific objectives of the missions can be avoided by certification of the technologies in a separate effort.

Overview of New Millennium

The New Millennium program (NMP) is designed to accelerate the realization of ambitious

missions by developing and validating some of the high-risk, high-benefit technologies they need. NMP conducts deep space and Earth orbiting missions focused on the validation of these technologies. The spacecraft flown by NMP are not intended to be fully representative of the spacecraft to be used in future missions, but the advanced technologies they incorporate are. As each NMP mission is undertaken, the risk of using the technologies that form its payload should be substantially reduced because of the knowledge gained in the incorporation of the new capability into the spacecraft, ground segment, and mission design as well as, of course, the quantification of the performance during the flight.

Although the objective of NMP technology missions is to enable future science missions, NMP missions themselves are not driven by science requirements. They are dedicated to technology, with the principal requirements coming from the needs of the advanced technologies they are testing. The science return from NMP missions is in the subsequent science missions that become feasible.

By their very nature, NMP projects are high risk. The key technologies that form the basis for each mission are the ones which require validation to reduce the risk of future missions. Indeed, if an advanced technology does not pose a high risk, testing by NMP is not required. In many cases, these unproven technologies will not have functionally equivalent back-ups on their test flights. Nevertheless, the failure of a new technology on

Copyright 1999 by the American Institute of Aeronautics and Astronautics, Inc. The U.S. Government has a royalty-free license to exercise all rights under the copyright claimed herein for Governmental purposes. All other rights are reserved by the copyright owner.

an NMP mission, even if it leads to the loss of the spacecraft, does not necessarily mean the mission is a failure. If the nature of the problem with the technology can be diagnosed, the goal of preventing future missions from accommodating the risk can be realized. Showing that a technology is not appropriate for use on subsequent science missions would be a very valuable result of an NMP flight. The acquisition of this information would achieve the goal of reducing the cost and risk to candidate future users of the technology. Of course, it is likely that such a determination would lead to modifications of the implementation of the technology, thus restoring its potential value to future space science missions.

Overview of DS1

Deep Space 1 (DS1) is the first project of NMP. Its payload consists of 12 technologies. The criteria for "complete mission success," agreed to by NASA Headquarters and JPL, are:

1) Demonstrate the in-space flight operations and quantify the performance of the following 5 advanced technologies:

- Solar electric propulsion (SEP)
- Solar concentrator arrays
- Autonomous navigation
- Miniature camera and imaging spectrometer
- Small deep space transponder

and any 3 of the following 6 advanced technologies:

- K_a-band solid state power amplifier
- Beacon monitor operations
- Autonomous remote agent
- Low power electronics
- Power actuation and switching module
- Multifunctional structure

2) Acquire the data necessary to quantify the performance of these advanced technologies by September 30, 1999. Analyze these data and disseminate the results to interested organizations/parties by March 1, 2000.

3) Utilize the on-board ion propulsion system (IPS) to propel the DS1 spacecraft on a trajectory that will encounter an asteroid in fiscal year 1999.

4) Assess the interaction of the IPS operations with the spacecraft and its potential impact on charged particle, radio waves and plasma, and other science investigations on future SEP-propelled deep space missions.

The first criterion clearly indicates that the goal of the mission is to determine how well the technologies work. Indeed, the wording reflects the recognition of the high risk of the technologies by allowing for the possibility that some might not be operable.

A twelfth technology, a miniature integrated ion and electron spectrometer, was not included in the success criteria, because it was so late in being delivered that even six weeks before launch it was uncertain whether the device would be ready. (This is another facet of the risk in planning to fly with advanced technologies.) Nevertheless, it was delivered and has performed very well.

All the technologies except autonomous navigation received 100% or more of their required testing by the end of June 1999. An asteroid encounter planned for July 29, 1999 tests 5% of the autonomous navigation system.

In addition to its technical objectives, DS1 was intended to probe the limits of rapid development for deep-space missions. The initial study of DS1 was undertaken only 39 months before launch, an unprecedentedly short time for a NASA deep-space mission in the modern era. At the time the preliminary concept study was initiated, the only definition of the project was that it should validate solar electric propulsion and other unidentified technologies in deep space and that launch should occur sometime in 1998. The level-1 requirements and goals¹ were formulated 26 months prior to launch.

Further background on the project, including the selection of technologies and the mission and spacecraft design, and additional information on NMP are presented elsewhere.^{1,2}

TECHNOLOGY RESULTS

Overviews of DS1's advanced technologies and the results from flight testing follow. The mission in which the technologies were used is discussed in the next section.

Solar electric propulsion

Solar electric propulsion (SEP) offers significant mass savings for future deep-space and Earth-orbiting spacecraft that require substantial velocity changes. The objective of the NSTAR (NASA SEP Technology Application Readiness) program, to validate low-power ion propulsion,

was a good match to NMP's goals. NSTAR involved a collaboration among JPL, NASA's Glenn Research Center, Hughes Electron Dynamics, Spectrum Astro, Moog, and Physical Science, Inc.

The ion propulsion system (IPS) on DS1³ uses a hollow cathode to produce electrons to collisionally ionize xenon. The Xe⁺ is electrostatically accelerated through a potential of up to 1280 V and emitted from the 30-cm thruster through a pair of molybdenum grids. A separate electron beam is emitted to produce a neutral plasma beam. The power processing unit (PPU) of the IPS can accept as much as 2.5 kW, corresponding to a peak thruster operating power of 2.3 kW and a thrust of 92 mN. Throttling is achieved by balancing thruster electrical parameters and Xe feed system parameters at lower power levels; and at the lowest PPU input power, 525 W, the thrust is 19 mN. The specific impulse ranges from 3200 s with about 2 kW delivered to the PPU to 1900 s at the minimum throttle level.

Because the purpose of flying the IPS was to validate it for future space science missions, a comprehensive diagnostic system is also on the spacecraft.⁴ This aided in quantifying the interactions of the IPS with the spacecraft, including advanced-technology science instruments, and validating models of those interactions. The diagnostic instrument suite includes a retarding potential analyzer, two Langmuir probes, search-coil and fluxgate magnetometers, a plasma wave sensor, and two pairs of quartz-crystal micro-balances and calorimeters. One of these pairs has a direct view of the ion thruster exit, while the other is shadowed by spacecraft structure. Measurements included the rate and extent of contamination around the spacecraft from the Xe⁺ plume and the sputtered Mo from the grid, electric and magnetic fields, and the density and energy of electrons and ions in the vicinity of the spacecraft. As a bonus, the sensors will be used to complement science measurements of DS1's ion and electron spectrometer (see below) during the small body encounters.

By June 30, 1999, the IPS had operated for nearly 1800 hours. This included several dedicated tests, but the majority of the time was devoted to placing the spacecraft on a trajectory to reach asteroid 1992 KD (in accordance with the third mission success criterion).

The IPS operated over a broad range of its 112 throttle levels, from input powers of 580 W (throttle level 6) to 2140 W (throttle level 90). The corresponding specific impulses were 1975 s and 3180 s. Measured thrust (determined through radio navigation) was within 2% of the prelaunch prediction throughout the range.

Comparison with the extensive ground-test program showed that operation in space is more benign and contamination is lower. The vast body of data from the diagnostics sensors on the effects of the IPS allows the development of guidelines for future designers on how to make fields and particles measurements on future IPS-propelled spacecraft without interference from the propulsion system.

In the first attempt to thrust with the IPS (on November 10, 1998), it operated for about 4.5 minutes and then switched to a standby mode. It is believed that the unplanned termination of the thrusting was the result of a contaminant causing a short between the two grids. Attempts to restart the thruster on that day were unsuccessful. Thermal cycling during the subsequent two weeks changed the spacing between the grids, thus stressing the contaminant, and when the IPS was commanded on again it operated as desired. Similar phenomena have been observed with other ion thrusters in space.

In the 1799.4 hours of thrusting (for deterministic thrust, trajectory correction maneuvers, and dedicated tests), the total Xe consumption was 11.4 kg, providing 699.6 m/s. After the first day of unsuccessful attempts to resume thrusting, all 34 IPS starts in the mission were successful.

All spacecraft systems operated normally during IPS thrusting. Telecommunications during IPS thrusting, even with the radio signal passing through the plasma, were unaffected. Sensors 0.7 m from the thruster with a direct line of sight to the exit grid recorded about 10 nm of surface contamination. Nearby sensors, without a direct line of sight, accumulated an order of magnitude less.

Solar concentrator array

Because of the IPS, DS1 required a high-power solar array. The Ballistic Missile Defense Organization (BMDO), working with NASA's Glenn Research Center, AEC-Able Engineering,

and Entech, developed the Solar Concentrator Array with Refractive Linear Element Technology (SCARLET II).⁵ BMDO wanted a flight test for SCARLET, and because it could provide the necessary high power, including it on DS1 was mutually beneficial.

SCARLET uses cylindrical silicone Fresnel lenses to concentrate sunlight onto GaInP₂/GaAs/Ge cells arranged in strips. Including the optical efficiency of the lenses, a total effective magnification greater than 7 is achieved. With relatively small panel area actually covered by solar cells, the total cost of cells is lowered, and thicker cover glass becomes practical, thus reducing the susceptibility to radiation. The dual junction cells display significant quantum efficiencies from 400 nm to 850 nm, and achieved an average efficiency in flight of about 22.5%.

The pair of arrays produced 2.5 kW at 1 AU, within 1% of the prelaunch prediction. Each array comprises four panels that were folded for launch, and a single-axis gimbal controls pointing in the more sensitive longitudinal axis. The two wings include a total of 720 lenses, each focusing light onto 5 cells. DS1 is the first spacecraft to rely exclusively on refractive concentrator arrays; it also is among the first to use only multibandgap cells.

The array is one of the three new technologies that had to work correctly immediately after launch in order for the mission to proceed; stored battery energy was sufficient only for a few hours. A substantial part of the validation of the array was the mechanical deployment and subsequent pointing. The deployment was so accurate that, following dedicated tests, no pointing adjustments were deemed necessary, and the array provided stable operation throughout the mission.

Autonomous navigation

Because mission operations is a significant part of its science budget, NASA explicitly included autonomy in its guidelines to NMP. A reduction in requirements for Deep Space Network (DSN) tracking of spacecraft will come from the placement of a complete navigation capability on board the spacecraft. (Other autonomy technology experiments are discussed below.) In addition, autonomous navigation allows a smaller navigation team during flight.

One portion of the core of the autonomous

system validated on DS1, AutoNav⁶, began functioning immediately upon activation of the spacecraft after separation from the launch vehicle, which occurred in Earth's shadow. The attitude control system (ACS) used a commercial star tracker to determine its attitude. Then the real-time part of AutoNav provided ACS with the position of the Sun so that ACS could turn the spacecraft to the attitude needed to illuminate the solar arrays upon exiting the shadow.

Data stored on board for use by AutoNav include a baseline trajectory, generated and optimized on the ground; the ephemerides of the DS1 target bodies, distant "beacon" asteroids, and all planets except Pluto; and a catalog of the positions of 250,000 stars (all contained in the Tycho catalog).

Throughout the mission, about once per week, AutoNav was invoked by the operating sequence to allow it to acquire optical navigation images. It issues commands to ACS and the integrated camera and imaging spectrometer (see below) to acquire visible-channel images, each with one beacon asteroid and known background stars. On-board image processing allows accurate extraction of the apparent position of each asteroid with respect to the stars, thus allowing the spacecraft location to be estimated. The heliocentric orbit is computed with a sequence of these position determinations combined with estimated solar pressure, calculated gravitational perturbations, and on-board knowledge of the thrust history of the IPS and incidental accelerations from unbalanced turns by the hydrazine-based reaction control system (RCS). The trajectory then is propagated to the next encounter target, and course changes are generated by the maneuver design element. In general, those course corrections are implemented through changes in the IPS thrust direction and duration, but in certain cases described below, the maneuvers are accomplished with dedicated IPS or RCS maneuvers.

After AutoNav parameters were tuned in flight, typical autonomous cruise heliocentric orbit determinations differed from radiometric solutions (developed to provide a reference against which to test AutoNav) by < 1000 km and < 0.4 m/s. With simple ground-based removal of some images (based on an algorithm that would be straight-forward to implement in the flight software), accuracies improved to < 400 km and < 0.2 m/s.

For encounters, navigation is target-relative, and 1σ delivery accuracy is ~ 3 km. AutoNav also performs target tracking at encounters to provide accurate pointing information to ACS, and it initiates the encounter sequences based on its estimate of the time to closest approach.

During periods of IPS thrusting, AutoNav controls the IPS. It selects the throttle level based on models of SCARLET power generation and spacecraft power consumption; pressurizes, starts, and stops the IPS; and commands ACS to achieve the attitude needed for thrusting. AutoNav also commands updates to the throttle level and spacecraft attitude every 12 hours.

During periods of ballistic coast, AutoNav is given time windows, in each of which it can execute a trajectory correction maneuver (TCM), which it designs autonomously if it has established that a TCM is necessary. In most cases, the TCMs are conducted with the IPS. To save time during the final 2 days before an encounter (and for the purposes of dedicated AutoNav testing), the hydrazine RCS is used. With either propulsion system, if thrust is required in an attitude that is prohibited by ACS, the TCM is autonomously decomposed into two allowed maneuvers.

Integrated camera and imaging spectrometer

If NASA is to conduct missions with smaller spacecraft, it is essential to have correspondingly smaller science instruments. One of the advanced technologies DS1 tested is the Miniature Integrated Camera Spectrometer (MICAS), conceived and developed by a team from the United States Geological Survey, the University of Arizona, Boston University, Rockwell, SSG, and JPL. In one 12-kg package, this derivative of the original concept for a Pluto Integrated Camera Spectrometer⁷ includes two panchromatic visible imaging channels, an ultraviolet imaging spectrometer, and an infrared imaging spectrometer plus all the thermal and electronic control. All sensors share a single 10-cm-diameter telescope. With a structure and mirror of highly stable SiC, no moving parts are required; the detectors are electronically shuttered. Spacecraft pointing directs individual detectors to the desired targets.

The instrument includes two visible detectors, both operating between $0.5\ \mu\text{m}$ and $1\ \mu\text{m}$: a 1024×1024 CCD with $13\text{-}\mu\text{rad}$ pixels and a

256×256 $18\text{-}\mu\text{rad}/\text{pixel}$ CMOS active pixel sensor, which includes the timing and control electronics on the chip with the detector. The imaging spectrometers operate in push-broom mode. The infrared spectrometer covers the range from $1.2\ \mu\text{m}$ to $2.4\ \mu\text{m}$ with spectral resolution of $12\ \text{nm}$ and $54\text{-}\mu\text{rad}$ pixels.

MICAS serves three functions on DS1. First, as with all the advanced technologies, tests of its performance establish its applicability to future space science missions. Second, AutoNav uses the visible channels for optical navigation. Third, as a bonus, it will collect science data during the primary mission at the asteroid and at other encounters if an extended mission is conducted.

The ultraviolet channel, designed to operate between $80\ \text{nm}$ and $185\ \text{nm}$, did not function properly and never returned interpretable data. Several tests were conducted to diagnose the problem, and indications are that the malfunction is in the signal chain after the detection of the photons.

MICAS images and IR spectra revealed scattered light. Stray light analysis and dedicated tests established the multiple paths responsible. The scattered light is the result of spacecraft surfaces directing off-axis light to reflective components inside MICAS, particularly the multilayer insulation surrounding the IR detector. The problem is easily avoided for future missions with different mounting of the instrument and alteration of the internal baffling. Modifications to AutoNav significantly increased its immunity to the light, and the flux is sufficiently low that it is not expected to interfere with encounter science.

Integrated ion and electron spectrometer

Just as MICAS integrates several different measurement capabilities into one low-mass package, the Plasma Experiment for Planetary Exploration (PEPE)⁸ combines multiple instruments into one compact package. At $5.6\ \text{kg}$ and $9.6\ \text{W}$, PEPE is less than 25% of the mass and consumes less than 50% of the power of a comparably performing (but more expensive) instrument on Cassini. Designed and built by Southwest Research Institute and Los Alamos National Laboratory, PEPE determines the three-dimensional plasma distribution over its $2.8\pi\ \text{sr}$ field of view.

PEPE measures the energy spectrum of

electrons and ions simultaneously from 8 eV to 33 keV per unit charge with at least 5% resolution. Rather than using moving parts, it electrostatically sweeps its field of view. PEPE measures ion mass with a resolution of 5% in the range of 1 to 500 amu per unit charge.

PEPE plays three roles on DS1. It has validated the design for a suite of plasma physics instruments in one small package; it has assisted in determining the effects of the IPS on the local plasma environment, including interactions with the solar wind and photoelectrons⁴; and it makes scientifically interesting measurements during the cruise and the encounters.

PEPE made measurements of the solar wind with the IPS on and off, and a very important result is that the data suggest that SEP can be used on future science missions without interfering with the scientific payload. PEPE data showed Xe⁺ returning to the spacecraft from the 1 ampere exhaust plume of the IPS and allowed limits to be placed on electrical charging of the spacecraft. In January 1999, a favorable alignment of the DS1 and Cassini spacecraft allowed 36 hours of collaborative solar wind measurements, with the two spacecraft separated by nearly 0.5 AU.

Telecommunications technologies

DS1 validated a small deep-space transponder (SDST), built by Motorola, and a K_a-band solid state power amplifier developed by Lockheed Martin.^{9,10} Combining the receiver, command detector, telemetry modulator, excitors, beacon tone generator (see below), and control functions into one 3-kg package, the SDST allows X-band uplink and X-band and K_a-band downlink. To achieve the SDST's functionality without a new technology development would require over twice the mass and 4 or 5 individual subassemblies. The SDST, along with SCARLET and AutoNav, had to function correctly from the beginning of the mission. Based on extensive routine use and dedicated experiments, its performance was in excellent agreement with preflight tests.

The SDST's K_a-band signal is amplified by the 0.7-kg power amplifier to 2.3 W with an overall efficiency of 13%. In addition to characterizing the operation of the hardware device, DS1 provided K_a-band signals for DSN use in verifying systems for acquiring, demodulating, decoding, and processing telemetry

as well as in producing 2-way Doppler and ranging data. The DSN also applied the K_a-band signals to the validation and improvement of system designs in preparation for upgrading to operational use of K_a-band. As the Earth-spacecraft range increased, certain tests were repeated to assure that the transition through threshold in a selected K_a-band data rate would be observed. All communication and radiometric tests proved to be in good agreement with models or with X-band results for the tests that were enhanced by simultaneous X-band operation.

Beacon monitor operations

The SDST generates the tones needed for beacon monitor operations¹¹, conceived to reduce the large demand that would be expected on the DSN if many missions were in flight simultaneously, as envisioned by NASA. In beacon monitor operations, an on-board data summarization system determines the overall spacecraft health. The system then transmits one of four tones to indicate to the operations team the urgency of the spacecraft's need for DSN coverage. Without data modulation, these tones are detected easily with small, low-cost systems, reserving the large, more expensive DSN stations for command radiation and data reception when a beacon indicates that such services are needed. The four tones correspond to *i*) the spacecraft not needing any assistance because all is well; *ii*) informing the ground that the spacecraft has encountered an unusual but not threatening event, so a DSN track should be scheduled when convenient; *iii*) alerting the ground that intervention is needed to prevent the loss of important data or to assist in resolving a threat to the mission, so DSN coverage should occur soon; and *iv*) requiring immediate assistance because the spacecraft has encountered a mission-threatening emergency it was unable to solve. In each case, when tracking is initiated, the data summarization system provides a synopsis of the pertinent spacecraft data.

This artificial intelligence technology uses adaptive alarm limits, which allow tighter monitoring than traditional limits. Furthermore, the spacecraft parameters that are monitored and their limits depend upon the spacecraft activity. The system adaptively filters data so instead of using fixed limits, it can compute variable limits on the fly; it can apply this not only to single data parameters but also to functions of multiple data parameters. These alarm limit functions are

“trained” using a neural network on the ground with actual DS1 engineering data to create functions that can perform more precise anomaly detection and detect important trends sooner than with conventional limits. Although this ground software is quite complex, only the resulting functions are uploaded to the spacecraft.

Experiments conducted during DS1 addressed both the data summarization and the tone generation and detection (in both X-band and K_a-band), which agreed well with preflight models. Beacon monitor operations may be relied upon during an extended mission if it occurs.

Autonomous remote agent

For the third autonomy technology DS1 tested, an artificial intelligence system was placed on board to plan and execute spacecraft activities.¹² The team that developed this system was drawn principally from JPL and NASA’s Ames Research Center. Rather than standard remote control, this technology uses an agent of the ground team on board the spacecraft. This remote agent was tested in a restricted case on DS1, in preparation for more ambitious experiments on subsequent flights. The remote agent includes an on-board mission manager that carries the mission plan, expressed as high-level goals. A planning and scheduling engine uses the goals, comprehensive knowledge of the spacecraft state, and constraints on spacecraft operations to generate a set of time-based or event-based activities, known as tokens, that are delivered to the executive. The executive expands the tokens to a sequence of commands that are issued directly to the appropriate destinations on the spacecraft. The executive monitors the response to these commands and reissues or modifies them if the response is not what was anticipated. A mode identification and reconfiguration engine aids in assessing the spacecraft state and in recovering from faults without requiring help from the ground except in extraordinary cases.

In the experiments on DS1, the remote agent operated selected subsystems based on plans formulated on board. Injection of four (simulated) faults tested remote agent’s ability to resolve or work around different classes of problems, and in each case it devised the correct response. A bug in the executive interrupted the first experiment, and the successful diagnosis of the problem was one important benefit of the testing; it also illustrated the value of trying out a

new technology on a dedicated test mission. The bug proved to be easily correctable for future uses of the technology. Analysis showed that it was safe to continue experiments on DS1 without fixing it, so a second experiment was devised, and it captured the remaining remote agent test objectives.

Microelectronics and structures

Electronics mass, volume, and power consumption are important drivers for overall spacecraft design. DS1 included tests of two microelectronics technologies and a mechanical/electronic experiment intended to contribute to the achievement of NASA’s vision of spacecraft in the future. To reduce the power consumption of electronics, one experiment used devices with very low voltage and low capacitance.¹³ This low-power electronics experiment contained four ring oscillators and some discrete transistors to test 0.9-volt logic and 0.25- μm gate lengths (achieved with 248-nm lithography) based on silicon-on-insulator technology. Provided by the Massachusetts Institute of Technology Lincoln Laboratory, the functioning of the devices in flight was in good agreement with prelaunch tests. DS1 also tested two power actuation and switching modules, the result of a joint development among Lockheed Martin, Boeing, and JPL.¹⁴ Each device contained four power switches, controlled by a mixed-signal ASIC, providing voltage and current sensing and current limiting. High-density packaging technology quadruples the packing density over the current state of the art. Designed to be capable of switching up to 40 V and 3 A, the experiment switched an internal test load on DS1. Regular tests showed that the performances of both PASM were consistent with prelaunch tests.

A multifunctional structure¹⁵ was provided to DS1 by the United States Air Force Phillips Laboratory and Lockheed Martin Astronautics as a test panel that was attached to the spacecraft bus. This new packaging technology integrates electronic housings and thermal control into load-bearing elements, thus offering great reductions in the mass of spacecraft cabling and traditional chassis. The DS1 experiment returned data on the performance of the electronic connection systems for embedded test devices and on the thermal gradients in the panel. The connectors displayed no evidence of degradation, and the thermal gradients were consistent with preflight predictions.

MISSION

Two objectives provided the impetus for a short mission. The principal requirement of DS1 was to return results promptly to the future users of the technologies. Except for tests of lifetime, most technologies could be evaluated on short (but intense) missions as well as long ones. In addition, in general shorter missions are less expensive than longer ones. As a result, it was decided that the primary mission would be about one year. This allowed sufficient time to exercise all the technologies under a wide range of conditions while keeping costs low and not forcing eager potential users to wait unreasonably long before being confident about the new systems. It also allowed sufficient time to accomplish the objective of thrusting with the IPS long enough to place DS1 on a ballistic trajectory to an asteroid (the third criterion for success).

DS1 was planned for launch in July 1998, based on the earliest expected spacecraft readiness in a schedule that was extremely aggressive (particularly given the large number of unproven technologies incorporated into the mission). The mission design, including solar system encounter targets, was based on that plan.¹⁶

In the spring of 1998, it became clear that launching DS1 in its planned launch period presented an unacceptable risk to mission success, so the launch period was shifted to October - November 1998. DS1 was given the slot vacated by the Far Ultraviolet Spectroscopic Explorer when its launch was moved to 1999. Still, an unusually dense schedule of launches and other activities at the Eastern Test Range made scheduling DS1's launch difficult. Once the launch period was selected, a new mission was designed with the requirements that it necessitate changes in neither the spacecraft nor ground systems and still be compatible with the secondary payload (see below). The original mission plan was sufficiently robust that its architecture did not need to be changed, but the encounter targets and thrusting and coasting times did change.

DS1's launch occurred at 12:08:00.502 UTC on October 24, 1998. It was launched on the first Delta II 7326-9.5 (from The Boeing Company), the smallest vehicle in the Delta stable, and was the first launch of NASA's Med-Lite program. This launch vehicle was selected largely on the basis of prompt availability and low cost, but its

capability exceeded what was needed for DS1, with relatively low mass and low injection energy (in part attributable to the high performance of the IPS). Including 81.5 kg of Xe and 31.1 kg of hydrazine, DS1 was 486.3 kg at launch, and the Delta provided a $C_3 = 2.99 \text{ km}^2/\text{s}^2$; the launch vehicle could have delivered approximately 600 kg to DS1's escape trajectory. The excess launch vehicle performance allowed the manifesting of another spacecraft on this launch. SEDSAT-1, built by the Students for the Exploration and Development of Space at the University of Alabama in Huntsville, in collaboration with NASA's Marshall Space Flight Center and Johnson Space Center, was mounted on the second stage, which accomplished insertion into Earth orbit. After the second stage's second burn, to raise the orbit of the third stage and DS1, the stage separated and carried SEDSAT-1 to its intended orbit, where it was separated. The third stage completed DS1's injection to heliocentric orbit.

Following launch, several days were spent conducting an initial evaluation of the spacecraft, verifying its health and preparing it for early mission operations. Dedicated technology experiments began within one week of launch. Of course, some technologies were used as part of regular spacecraft operations, in particular the solar array, transponder, and AutoNav, but those and all other technologies also were subjected to in-depth characterization tests.

Radiometric determination of the actual trajectory was combined with results of the first SCARLET and IPS tests to generate and optimize an updated low-thrust trajectory that was transmitted to the spacecraft. After verification of its functional capability, AutoNav was tuned in flight, particularly to account for discrepancies between the predicted and the actual MICAS images. As the mission progressed, more reliance was placed on AutoNav, with conventional radio navigation used to validate its performance.

Initial IPS thrusting was conducted with the thrust vector along the Earth-spacecraft line to maximize communications rates and the Doppler signature, in order to quantify the actual thrust at selected throttle levels. After 10 days of thrusting, the spacecraft was turned to thrust along the optimal vector (subject to a variety of pointing constraints) for reaching the encounter targets for the primary and extended mission.

To meet the demanding schedule prior to launch, some software development was completed after launch. The launch load did not include all functions needed to conduct tests with the low-power electronics, power actuation and switching modules, multifunctional structure, beacon monitor operations, and remote agent. These technologies were selected for exclusion from the earlier software because they were not needed for the basic operation of the mission. In February 1999, a completely new software load of 4.1 megabytes was installed. This new software enabled the testing of four of the previously excluded technologies (remote agent was not in this load), upgraded AutoNav (principally to accommodate scattered light in the MICAS images), corrected bugs identified after launch, and improved spacecraft operability.

To accommodate the remote agent experiments in May, the flight software was patched; in addition, the remote agent software was updated. In June, following the remote agent experiment, the entire flight software was replaced again. This last load contained new operational enhancements and upgrades to a number of systems, but primarily it included further AutoNav upgrades for enhanced image processing (such as image differencing to gain greater suppression of scattered light and more powerful corrections for MICAS' large geometric distortions) and the functions needed to execute encounters. All three changes to the flight software, which included substantial development and testing, large uplink volumes, and rebooting of the nonredundant main spacecraft computer, were completed without incident.

The mission operations system made extensive use of standard tools and mission services JPL provides to support a wide range of missions. DS1 employed JPL's multimission ground data system to provide the uplink and downlink data transport capabilities as well as much of the telemetry processing and display system. Project-developed applications augmented the system to be consistent with the autonomous capabilities of the spacecraft.

DS1 mission operations were significantly different from that of typical deep space missions at JPL. This was primarily attributable to the technology-validation focus of the DS1 mission. Unlike typical deep space missions, with its very active technology testing campaign, DS1 did not

have a quiet cruise. Because of the experimental nature of the spacecraft and the technologies, early sequence development was confined to implementing and validating command activity blocks that could be modified readily and executed on board by real-time commanding to achieve a desired technology experiment. In the first three months of flight, about 1800 real-time commands were executed by the spacecraft.

The judicious use of multimission tools and services and standards such as CCSDS (Consultative Committee for Space Data Standards) kept the cost of the ground system and mission operations to a minimum. The small operations team averaged about 50 full-time equivalents, including system and subsystem analysts, flight controllers, technology support teams, testbed engineers, and project management and staff.

During the routine IPS thrust periods, one DSN pass each week allowed high-rate commanding and return of spacecraft engineering and technology validation data through the high gain antenna. This weekly track was immediately preceded by AutoNav's collection of optical navigation images, and both activities were conducted with the IPS off. The IPS thrusted for the remaining 90% of the week. One or two shorter passes were scheduled between the longer ones. Conducted only with one of the low gain antennas, to allow communication in the preferred thrust attitude, the shorter passes were used to verify that the IPS was thrusting. On occasion this coverage also was used to conduct IPS or SCARLET experiments.

The strategy for selecting IPS thrust and coast times was based on compromises between optimizing the trajectory and conducting the technology experiments and other mission activities incompatible with the attitudes or other spacecraft states required for thrusting.¹⁶ As illustrated in Figure 1, the deterministic thrusting for the primary mission was accomplished in two major periods. The brief hiatus in the first major thrust arc was inserted to allow several days for activation and initial testing of PEPE in the absence of the IPS plasma, and SDST and K_a-band experiments incompatible with the IPS thrust attitude. When the second thrust segment ended on April 27, 1999 (under direction of AutoNav), the spacecraft was on a ballistic trajectory that would encounter asteroid 1992 KD, thus




Figure 1. DS1 trajectory for the primary mission (through September 18, 1999) and proposed extended mission. The dotted portion is for ballistic coast; the solid portion indicates the IPS thrust is on.

accomplishing the third mission success criterion. The thrust plan was developed to maintain the option for an extended mission (see below).

On July 29, 1999, the spacecraft will encounter (9969) 1992 KD at 15.5 km/s. The size and shape are poorly known, but this asteroid is believed to be elongate with a mean radius of roughly 1 km. During the final 20 days of the spacecraft's approach to the body, AutoNav will require optical navigation images and trajectory correction maneuvers at increasing frequencies to

control the targeting of the final encounter. The maneuvers prior to 2 days before closest approach will be executed with the IPS, and in the final 2 days the RCS will be used to save time.

Because 1992 KD is so faint, it will not be detected by AutoNav (using MICAS images) until about 1 day before closest approach; until the asteroid is detected on board, AutoNav will continue to use 1992 KD's *a priori* ephemeris. A flyby 15 km from the center of the body is planned. With an expected navigational delivery

accuracy of about 3 km (1σ), this assures a safe but very exciting encounter. The last opportunity for a trajectory correction maneuver will be 3 hours before closest approach.

During the final approach, AutoNav's MICAS images will be interspersed with MICAS images and spectra collected for science purposes. The late navigation images will contain information AutoNav needs to provide rapid updates to its estimate of the position of 1992 KD, critical for keeping the asteroid in MICAS' field of view. The 4 command sequences, sequentially governing activities during the final 5 minutes before closest approach, will be activated by AutoNav based on its estimates of the time to closest approach. Because MICAS is body-fixed, the termination of imaging is dictated by when the angular acceleration of the line of sight to the asteroid exceeds ACS' capability to keep 1992 KD in the MICAS field of view. Measurements by PEPE and diagnostics sensors for the IPS will continue through closest approach.

The asteroid encounter will allow an opportunity to gather data on the size, shape, geomorphology, albedo, and the mineralogy and compositional heterogeneity of the surface material. It may be possible to measure, or at least constrain, the asteroid's magnetization, interaction with the solar wind (including sputtering), and outgassing.

The spacecraft will point its high gain antenna at Earth about 1 hour after closest approach to begin returning technology validation and science data. Although the data return will require several days, the IPS may resume thrusting as soon as several hours after closest approach. It turns out that with the antenna Earth-pointed, the IPS is within 30° of the optimal attitude for thrusting for the extended mission. For purposes of the extended mission, it is better to thrust in that attitude than to coast.

The end of the primary mission is on September 18, 1999. No new technology experiments are planned after the asteroid encounter. Following the completion of the return of data, some minor engineering activities will be conducted to prepare for the resumption of long-term thrusting, and then the regular cycle of IPS thrusting, interrupted only for weekly acquisition of optical navigation images and DSN communications, will resume.

EXTENDED MISSION

If the spacecraft remains healthy and the resources for an extended mission are available, the DS1 project could conduct a scientifically exciting mission. With the technology testing complete, the extended mission would be devoted to comet science. With AutoNav controlling the IPS, the spacecraft would travel to Comet 107P/Wilson-Harrington and Comet 19P/Borrelly.

As illustrated in Figure 1, most of the extended mission would be devoted to IPS thrusting. By the end of the extended mission, the spacecraft would have expended essentially all of its Xe, providing a total of about 4.5 km/s.

Because of the reduced mission operations staff and the increasing geocentric range during the extended mission, beacon monitor operations likely would be used to augment the team's ability to monitor the spacecraft's health. Demand-access operations have not been implemented by the DSN however, so that aspect of beacon monitor operations cannot be implemented.

In January 2001, DS1 would reach Comet Wilson-Harrington. This comet was lost after its discovery in 1949. In 1992, asteroid (4015) 1979 VA was recognized to be the same body. It is possible therefore that this comet was seen just as its activity was terminating. It is considered to be a dormant comet or a comet/asteroid transition object, with an estimated radius of 2 km.

With DS1's relative speed of 15.8 km/s, the encounter would be similar to the 1992 KD encounter, but it would occur when the comet is near solar conjunction. Although the operations team would have reduced control authority at that time, AutoNav would control the trajectory and timing of sequence activations. Of course, there would be sufficient time to incorporate the results of the final testing of AutoNav at 1992 KD.

In September 2001, DS1 would encounter Comet Borrelly at 17.0 km/s, within days of the comet's perihelion; this is one of the brightest and most active short-period comets. The nucleus is believed to be a prolate spheroid of about $4 \text{ km} \times 2 \text{ km}$ with an active surface area of 7% - 10%. Science data at the comet that could be collected include the structure and composition of the coma and tail (including gas, plasma, and dust), the

nature of jets and their connection to surface features, the interaction with the solar wind, and the same kind of characterization of the nucleus as at the asteroid.

The extended mission plan, although devoted to science, illustrates the benefits of the advanced technologies. If DS1 had used conventional technologies, including a bipropellant propulsion system (and excluding the fraction of the solar arrays needed for operation of the IPS), a transponder similar to that on the Mars Climate Orbiter and Mars Polar Lander (launched a few months after DS1), and science instruments with similar capability but without all the innovations being tested on DS1, the spacecraft would be significantly more massive. Reaching 1992 KD, Wilson-Harrington, and Borrelly would have required an injected mass of approximately 1300 kg compared to DS1's 486.3 kg. And rather than being able to share a launch on the least expensive Delta II, the requirements of this hypothetical mission would have exceeded the capability even of a dedicated Delta II 7925, the most expensive member of that family. The smallest operational US launch vehicle that would have met the requirements is the Atlas IIA, for which a shared launch would suffice.

CONCLUSION

The successful flight of DS1 provided an extensive body of data characterizing the 12 technologies it tested in space. By operating these advanced technologies under actual spaceflight conditions, the cost and risk to subsequent users should be greatly reduced, thus allowing rapid integration of the important capabilities they offer into future space and Earth science missions. Another significant benefit of the testing of technologies on DS1 was the experience gained by engineering teams. In many cases, the technologists had not worked on flight projects, and their experiences in both development and operations should prove helpful in their work on future versions of their technologies. The incorporation of the technologies into an operational mission yielded valuable insights into implementation issues that would not be expected to arise in typical technology development or conceptual mission studies. In addition, spacecraft, ground system, mission planning, and mission operations teams discovered the implications of integrating these new technologies into their designs and, of course, learned how to take advantage of the

capabilities of the technologies to create new designs. Any informed user, seeking to benefit from the capabilities of these advanced technologies, now will encounter lower risk and cost by building upon the successful results of the DS1 project.

ACKNOWLEDGMENTS

The members of the DS1 mission operations and technology teams are gratefully acknowledged for their impressive work, upon which this overview is based.

Some of the research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.


REFERENCES

1. Rayman, Marc D. and David H. Lehman, "Deep Space One: NASA's First Deep-Space Technology Validation Mission," *Acta Astronautica* **41**, p. 289 (1997).
2. Rayman, Marc D. and David H. Lehman, "NASA's First New Millennium Deep-Space Technology Validation Flight," Second IAA International Conference on Low-Cost Planetary Missions, Laurel, MD, April 16-19, 1996, IAA-L-0502.
3. Polk, J. E., R. Y. Kakuda, J. R. Anderson, J. R. Brophy, V. K. Rawlin, M. J. Patterson, J. Sovey, and J. Hamley, "Validation of the NSTAR Ion Propulsion System on the Deep Space One Mission: Overview and Initial Results," 35th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Los Angeles, CA, June 20-24, 1999, AIAA-29-2274.
4. Wang, J., D. Brinza, R. Goldstein, J. Polk, M. Henry, D. T. Young, J. J. Hanley, J. Nordholt, D. Lawrence, and M. Shappirio, "Deep Space One Investigations of Ion Propulsion Plasma Interactions: Overview and Initial Results," 35th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Los Angeles, CA, June 20-24, 1999, AIAA 99-2971.
5. Stella, Paul M., Donald G. Nieraeth, David M. Murphy, Michael I. Eskenazi, and John Stubstad,

- “Validation of the SCARLET Advanced Array on DS1,” *Proceedings of the 34th Intersociety Energy Conversion Engineering Conference*, Vancouver, British Columbia, August 1-5, 1999, paper 1999-01-2630.
6. Riedel, J. E., S. Bhaskaran, S. P. Synnott, S. D. Desai, W. E. Bollman, P. J. Dumont, C. A. Halsell, D. Han, B. M. Kennedy, G. W. Null, W. M. Owen Jr., R. W. Werner, and B. G. Williams, “Navigation for the New Millennium: Autonomous Navigation for Deep Space 1,” 12th International Symposium on Space Flight Dynamics, Darmstadt, Germany, June 2-6, 1997, paper 88.
7. Beauchamp, P. M., R. H. Brown, C. F. Bruce, G-S. Chen, M. P. Chrisp, G. A. Fraschetti, T. N. Krabach, S. W. Petrick, D. H. Rodgers, J. Rodriguez, S. L. Soll, A. H. Vaughan, L. A. Soderblom, B. R. Sandel, and R. V. Yelle, “Pluto Integrated Camera Spectrometer (PICS) Instrument,” *Acta Astronautica* **35**, Supplement 1995.
8. Bolton, S. J., D. T. Young, J. E. Nordholt, J. L. Burch, H. O. Funsten, N. Eaker, and D. J. McComas, “Plasma Experiment for Planetary Exploration (PEPE)”, in *Proceeding of Space Technology and Applications International Forum*, Albuquerque, NM, 1997, AIP Proceedings 387, ISBN 1-56396-679-4.
9. Herman, M. I., S. Valas, W. Hatch, C. C. Chen, S. H. Zingales, R. P. Scaramastra, L. R. Amaro, and M. D. Rayman, “Deep Space One Telecommunication Development,” 12th Annual AIAA/USU Conference on Small Satellites, September 6-10, 1998, Logan, Utah, paper SSC98-IV-1.
10. Karnacewicz, M. K., W. J. Taft, S. Valenti, T. Renna, M. Hirokawa, S. Conway, and L. Newman, “2.6 Watt K_a-Band Solid State Power Amplifier for Deep Space Communications,” *1999 Digest of Papers, Government Microcircuit Applications Conference*, pp. 247-250.
11. Wyatt, Jay, Rob Sherwood, Miles Sue, and John Szijjarto, “Flight Validation of On-Demand Operations: The Deep Space One Beacon Monitor Operations Experiment,” *Proceedings of the 5th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS-99)*, Noordwijk, The Netherlands, June 1-3, 1999.
12. Nayak, P. Pandurang, Douglas E. Bernard, Gregory Dorais, Edward B. Gamble Jr., Bob Kanefsky, James Kurien, William Millar, Nicola Muscettola, Kanna Rajan, Nicolas Rouquette, Benjamin D. Smith, William Taylor, and Yu-wen Tung, “Validating the DS1 Remote Agent Experiment,” *Proceedings of the 5th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS-99)*, Noordwijk, The Netherlands, June 1-3, 1999.
13. D’Onofrio, R. P., C. L. Keast, A. M. Soares, P. W. Wyatt, “Low Power Flight Experiment: Flight Qualifying Sub-0.25- μ m Fully Depleted SOI CMOS Technology,” *1999 Digest of Papers, Government Microcircuit Applications Conference*, pp. 238-241.
14. Salim, Abbas, Gerhard Franz, Jim Jud, David Hogue, and Gregory Carr, “PASM, The Advanced Power Actuation and Switching Module as the Building Block for Space Micropower Systems,” *1999 Digest of Papers, Government Microcircuit Applications Conference*, pp. 242-246.
15. Barnett, David M. and Suraj P. Rawal, “Multifunctional Structures Technology Experiment on Deep Space 1 Mission,” *IEEE AES Systems Magazine*, January 1999, pp13-18.
16. Rayman, Marc D., Pamela A. Chadbourne, Jeffery S. Culwell, and Steven N. Williams, “Mission Design for Deep Space 1: A Low-Thrust Technology Validation Mission,” Third IAA International Conference on Low-Cost Planetary Missions, Pasadena, CA, April 27 - May 1, 1998, IAA-L98-0502.




Stowed Configuration







Stowed Configuration





Deep Space 1 Trajectory





Autonomous Optical Navigation (AutoNav) DS1 Technology Validation Report

AutoNav Team:

J. E. Riedel, S. Bhaskaran, S. Desai, D. Han, B. Kennedy, G. W. Null,
S. P. Synnott, T. C. Wang, R. A. Werner, E. B. Zamani

Radio Nav Team:

T. McElrath, D. Han, M. Ryne

*Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California 91109*



Table of Contents

<u>Section</u>	<u>Page</u>
Extended Abstract	vi
Fact Sheet	vii
1.0 Introduction	1
2.0 Technology Description	1
2.1 Technology Overview.....	1
2.2 AutoNav Technology-Validation Plans and Objectives.....	2
2.2.1 AutoNav Validation Plan Overview.....	2
2.2.2 AutoNav Key-Point Technology Description and Validation Strategy.....	4
2.3 Expected Performance Envelope.....	7
2.4 Detailed Technology Description.....	8
2.4.1 The AutoNav System.....	8
2.4.2 AutoNav File Descriptions.....	8
2.4.3 Software System.....	10
2.4.4 AutoNav Commanding Strategy.....	10
2.4.5 “Uncommanded” AutoNav Functions.....	12
2.4.6 Core Algorithm Descriptions.....	13
2.5 Technology Interdependencies.....	14
2.5.1 MICAS/AutoNav Interface.....	14
2.5.2 Attitude Control System (ACS).....	18
2.5.3 Ion Propulsion System.....	19
2.5.4 Remote Agent and RAX.....	19
2.5.5 Fault Protection.....	19
3.0 Test Program	19
3.1 Ground Test.....	19
3.1.1 UNIX-Based Simulation.....	20
3.1.2 TestBed Testing.....	20
3.2 Flight Test.....	20
3.2.1 Early AutoNav Flight Operations.....	20
3.2.2 The First Validation of Onboard Orbit Determination.....	22
3.2.3 Results from MICAS Calibration Images.....	23
3.2.4 Late Cruise Timeline.....	23
3.2.5 Final Software Load and Final Validation of Cruise AutoNav.....	24
3.2.6 Asteroid Rehearsal Preparations.....	26
3.2.7 Results from the 7/13/99 Encounter Rehearsal.....	27
3.2.8 Cruise to –5 Day TCM.....	28
3.2.9 Acquisition of Target and Countdown to Encounter.....	30
3.2.10 Post-Encounter Reconstruction and Performance Analysis.....	32
3.2.11 Causes of the Braille-Encounter Failure.....	32
3.2.12 Post-Braille Cruise Operations.....	35
4.0 Technology Validation Summary	35
4.1 Summary Overview.....	35
4.2 Pre-Flight Validation.....	36
4.2.1 Prototype Validation.....	36
4.2.2 Development Bench-Testing.....	36
4.2.3 Software Module Delivery and Version Testing.....	36
4.3 In-Flight Validation.....	36
4.3.1 Early Cruise AutoNav.....	36
4.3.2 Late Cruise AutoNav.....	37
4.3.3 Encounter Phase: Rehearsal.....	37
4.3.4 Encounter Phase: Actual.....	37
5.0 Application of AutoNav to Future Missions	37
5.1 Requirements for Use of AutoNav.....	37

5.2 Types of Missions that can Use AutoNav to Advantage	37
5.3 Adaptations Necessary or Desirable for Future Use	37
5.3.1 Adaptations for Cameras	37
5.3.2 Dynamic Model Upgrades	38
5.3.3 Ephemeris Extensions	38
5.3.4 Image Analysis Extensions and Enhancements	38
5.3.5 Software and Spacecraft System Adaptations	38
5.3.6 Picture Planning Full Automation	38
5.3.7 Multiple Spacecraft Navigation	38
6.0 Concluding Remarks	38
7.0 Acknowledgments	39
8.0 References	39

Figures

<u>Figure</u>	<u>Page</u>
Figure 1. DS1 Mission and AutoNav Operations/Validation Plan and Schedule	3
Figure 2. Mission and AutoNav Operations/Validation Plan and Schedule for Braille Close Approach	4
Figure 3. RSEN Time-of-Flight Performance	7
Figure 4. Subsystem Performance Influence on AutoNav	8
Figure 5. The AutoNav Software System and Interacting System Software	11
Figure 6. Reduced State Encounter Navigation Schematic Functional Overview	13
Figure 7. Multiple Cross Correlation of Asteroid and Stars	14
Figure 8 a,b,c. Spacecraft Integration Equations of Motion and Derivation of AutoNav OD Kalman Filter	15
Figure 9. Adjusting a Low-Thrust Burn Arc	16
Figure 10. MICAS Extended Bright-Image Charge Bleed	17
Figure 11. MICAS “Low Solar Cone Angle” Scattered-Light Picture	17
Figure 12. MICAS “High Solar Cone Angle” Scattered-Light Picture	18
Figure 13. Primary Mission Trajectory Plan	21
Figure 14. Pre-(upper) and Post-(lower) Fit Residuals from 3/22/99 Optical Solution	22
Figure 15. Flight vs. Ground-Orbit Determination April 5, 1999	24
Figure 16. Flight vs. Ground-Orbit Determination 5/31/1999	25
Figure 17. Current B-Plane Target Conditions at the –20–10 Day TCMs: Decision Data from 7/15/99	27
Figure 18. Minus 5 Day TCM Solutions	29
Figure 19. Flight OD vs. Ground OD#37, 7/21/99	29
Figure 20. Pre-Minus—1 Day TCM, “Flight OD” Braille B-Plane	31
Figure 21. Pre-ACA—6 hour TCM B-Plane July 27	31
Figure 22. Diagrammatic View of Received RSEN Signal	32
Figure 23. MICAS APS Channel Non-Linear Signal Response	33
Figure 24. Encounter Results Using Post-Encounter CCD Braille Pictures	33
Figure 25. Post-Encounter APS Image of Braille	34
Figure 26. Reconstructed Nominal vs. Perturbed Braille Field-of-View Flight Path	34
Figure 27. Post-Braille AutoNav DataArc and Residuals	35

Tables

<u>Table</u>	<u>Page</u>
Table 1. AutoNav Files, Sizes, Autonomy Status, Locations, and Update Frequency	9
Table 2. Summary of AutoNav Commands	11
Table 3. Imaging System AutoNav Requirements and Attainment by MICAS	16
Table 4. Navigation Encounter Activities	26

Appendices

Appendix A: Listing of Telemetry Channels and Names	40
Appendix B: Date of Turn-on and Frequency of Data Capture	42
Appendix C: Descriptive Paper of AutoNav High-Level Design Considerations	45
Appendix D: Analytic Paper of AutoNav Orbit Determination	64
Appendix E: Analytic Paper of AutoNav Trajectory Control Algorithms	85
Appendix F: DS1/ AutoNav Technology Validation Agreement	97
Appendix G: AutoNav—Extended Mission	102
Appendix F: Navigation of the Deep Space 1 Spacecraft at Borrelly.....	121

EXTENDED ABSTRACT

The first mission of NASA's New Millennium Program, *Deep Space 1* (DS1), has as one of its principal demonstration technologies the first autonomous optical navigation system to be used in deep space. The concept of DS1—to develop and validate new technologies in the context of a low-cost, deep-space planetary mission—was extremely challenging. In practice, the challenges were even greater. Nevertheless, the complete manifest of technologies was validated, with most of them proving highly successful, including the autonomous navigation system, AutoNav.

The theoretical basis of AutoNav is a process in which images of asteroids (typically main-belt) are taken against the distant stars and, through the measured parallax, geometric information is inferred. This information is used in a dynamic filter to determine spacecraft position and velocity, as well as parameters describing the performance of the ion propulsion system (IPS) and solar pressure. With this information, corrections to the mission design as described in the propulsion profile are made and/or predictions for necessary trajectory correction maneuvers (TCMs) are computed. This system is shown diagrammatically in the Fact Sheet.

The AutoNav system is a set of software elements that interact with the imaging, attitude-control, and ion propulsion systems aboard DS1. The principal elements and functions of AutoNav are:

1. NavRT—Provides critical ephemeris information to other onboard subsystems, such as the Attitude Control System.
2. NavExec—Plans and executes various important Nav-related activities, such as image taking and processing, ion propulsion system thrusting events, and TCMs.
3. ImageProcessor—Performs image processing.
4. OD—Performs orbit-determination computations.
5. ManeuverPlanner—Performs computations relative to IPS events and TCMs.

The validation of the AutoNav system was to be accomplished through its use as the principal navigation system. As such, a comprehensive series of activities were planned to, primarily, accomplish the many navigation tasks for DS1 and, secondarily, to validate AutoNav. These tasks and their completion and/or validation status are shown in the table on the Fact Sheet.

From the very first invocation of the higher functions of AutoNav, soon after launch in October of 1998, there were serious challenges. The imaging system onboard DS1 suffered from serious light-leakage problems. As a result of this and a general lack of camera sensitivity, the availability of adequately bright asteroids to image was very limited. The light-leakage problems also seriously degraded the ability of the image-processor to reduce the data. Additionally, the geometric distortions of the camera field were much worse post-launch than pre-launch lab testing had indicated. All of these factors contributed to initial navigation errors of 10,000 km and 7 m/s in the spacecraft state. Nevertheless this was (and is) adequate quality for cruise operations of an interplanetary mission.

Efforts were immediately undertaken to compensate as much as possible for the camera shortcomings. With a new load of software onboard in February of 1999—and a further update in June—performance gradually improved to the level of 250 km and 0.2 m/s, very nearly the pre-launch (and pre-anomaly) predicted performance and substantially better than the validation requirement. On approach to the first of three encounter targets planned for the mission, AutoNav adjusted the IPS-powered course, and computed and executed TCMs. Three weeks before the Braille encounter, a “full dress” rehearsal of the encounter was performed. AutoNav operated without problems, delivering the spacecraft to within the required 2.5-km control parameter, tracking the target to within 30 s of closest approach, and effectively reducing the field-of-view errors to within the required 0.5 km.

During the actual close approach of Braille, not surprisingly, unexpected conditions were encountered. The actual brightness of the asteroid was a factor of 5 to 10 below expectation and the camera channel used was 4 to 5 times less sensitive than designed and anticipated, resulting in previously set thresholds for discriminating real target signals not being crossed. As a consequence, the close-approach target-tracking system of AutoNav did not “lock-on” to the target. Since the encounter sequence was aggressively “success oriented” and early (distant) images were not preserved onboard (due to a lack of storage RAM), the eagerly anticipated high-resolution images were not acquired. Nevertheless, important information was gathered about the operation of the DS1 suite of technologies that will be applied to the encounters with comets Wilson Harrington and Borely in 2001.

This report details the technology development, implementation strategy, testing methodologies, and testing results as well as the actual inflight success of the operation of the DS1 AutoNav system.

FACT SHEET

AUTONOMOUS OPTICAL NAVIGATION (AutoNav) for NEW MILLENNIUM DSI : Technology Validation Fact Sheet


Contact: Joseph.E.Riedel@jpl.nasa.gov; Jet Propulsion Laboratory, Pasadena, CA; 818-354-8724

CONVENTIONAL NAVIGATION

Encounter Phase:
Ground Based Approach Optical Navigation, Limited in accuracy, Large flyby ranges required, also reduce science.

Maneuvers:
Maneuver Computed on Ground, Parameters Uplinked, requiring ground processing and analysis

Earth-Based radio and optical data Processing




DSI AUTONOMOUS NAVIGATION

Encounter Phase: Optimum return of science with onboard Nav closed-loop target tracking. Onboard Rehearsal tracks to within 250km range, knowledge error reduced to 0.5km

Images processed on-board, to 0.4 px acc'y

Maneuvers:
Autonomous Maneuver Computation onboard. IPS profile updated and TCM's (RCS and IPS) performed inflight, keeping Spacecraft on-target.

Cruise Phase:
Spacecraft Position, Velocity and Forces Estimated Onboard from Optical Data triangulation. Accuracies better than 250km and 0.2km/s Achieved.



Candidate Future Mission Types for AutoNav

- Missions with Very Limited Tracking Budgets or Tracking Limitations
 - Missions with Severe Dynamical Control Challenges, Requiring Fast Loop Closure
 - Mission with long "Unsupervised" Cruise Periods.
- (AutoNav Equipped Mission Requirements: Reasonable Quality Imaging Instrument, and High Level ACS Performance and Semi-Autonomous Functionality)

Diagrammatic and Comparative Description of AutoNav Technology and Validation

AutoNav Technology Validation Key Point Summary

A	B: Technology Validation Item Description	C	D	E	F	G	H	I	J
1	Provision of Ephemeris Services	~10 ⁵	~10 ⁵	~10 ⁵	0	≤0.1 km	Req'd	<<0.1 km	<<0.1 km
2	OpNav PhotoOp Process	~40	47	46	1				
3	Image Data Handling and Downlink	~40	47	47	0				
4	OpNav Data Accumulation, Handling, Downlink	~40	47	44	3				
5	Image Processing (RSS ensemble statistics)	~1200	~1500	~500	0	≤0.25 px	Desir'd	≤0.40 px	1.5 px
6	Orbit Determination (Accuracy within data arc)	~32	34	34	0	≤250 km, 1 m/s	Req'd	≤150 km, 0.2 m/s	10000 km, 7 m/s
7	Generation of Onboard Ephemeris and Downlink	~32	34	34	0	0.1–1 km	Req'd	0.1 km	1 km
8	Trajectory Control and Maneuver Planning	~20	12	12	0				
8a	IPS Mission Burn Updates (convergence criteria)	~12	6	6	0	≤1 km	Desir'd	≤1 km	≤1 km
8b	IPS and RCS Maneuver Computations (convergence criteria)	~8	5	5	0	≤1 km	Desir'd	≤1 km	≤1 km
8c	TCM Execution, and Delivery (final TCM and accuracy – position and velocity)	8(2)	5(1)	5(1)	0	(≤2.5 km, 0.25 m/s)	(Req'd)	(≤1.5 km, 0.18 m/s)	(≤1.5 km, 0.18 m/s)
9	Execution of Mission Burns	~12	7	7	0				
10	Encounter Image and OD Operations (RSEN)	2	2	1	0				
10a	Image Processing, and Data Reduction	~80	~80	~40	1				
10b	Ephemeris Generation and Delivery	~80	~80	~40	0	≤0.5 km	Req'd	≤0.5 km	15 km
11	Encounter: Initiation of Encounter Sequences	8	8	8	0	≤5 s	Desir'd	≤5 s	≤15 s

Legend- A: Item Number (Appendix F), B: Item Description, C: No. Planned In-Flight Executions, D: No. Actual In-Flight Executions, E: No. Successes In-Flight, F: No. Failures In-Flight (due to AutoNav Fault and/or Misuse), G: Quantitative Goal-Value (If Applicable), H: Required/Desired Quantitative Value, I: Best Value Achieved, J: Worst Value Achieved

Autonomous Optical Navigation (AutoNav) Technology Validation Report

*S. Desai, D. Han, S. Bhaskaran, B. Kennedy, T. McElrath, G. W. Null, J. E. Riedel,
M. Ryne, S. P. Synnott, T. C. Wang, R. A. Werner, E. B. Zamani
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California*

1.0 INTRODUCTION

Optical Navigation, as it is currently being applied by the deep-space probes of JPL/NASA, is a technique by which the position of a spacecraft is determined through astrometric observations of targets against a background field of stars. The stars and target positions are known by ground or other observations, independently, or concurrently made, and the position of the spacecraft taking the image is inferred from the “error” in the position of the near-field object against the far-field (i.e. the parallax). In practice, there are many complicating details. These include the numerical integration of the spacecraft trajectory, which requires accounting for adequate non-gravitational perturbation models in the spacecraft. Also to be provided is adequate accuracy in the star catalog, including accounting for proper motion. Adequate calibration of the camera field-of-view distortions must be provided, as well as dynamic filtering of the acquired optical data, including stochastic estimation of pointing and spacecraft dynamic parameters.

Early demonstrations of optical navigation on deep-space probes were performed on some of the later *Mariner* series and on the Mars *Viking* mission. However, the first missions that required optical navigation to accomplish the principal mission objectives were the *Voyager 1* and *2* missions. The key technological developments for interplanetary optical navigation were made then [1][2][3][4]. Following the successful use of optical navigation, variations of this system were used for the *Galileo* approach and flybys of *Ida* and *Gaspra* [5], and during the *Galileo* Jovian tour. Due to a failure of *Galileo*’s high-gain antenna, however, new technologies had to be developed for optical navigation, primarily to increase the information content from any single image. These new technologies include the multiple-cross-correlation technique, used for the *Gaspra* and *Ida* flybys, and an autonomous detection and capture algorithm loaded onboard to search through a navigation frame to find the target body (a Galilean satellite) and stars. Both of these algorithms were subsequently put to use onboard DS1 as part of the AutoNav system.

The concept of providing a completely autonomous onboard optical-navigation system arose from several sources. An era of space exploration comprised of many small semi- or fully-autonomous spacecraft would be

impossible to achieve without a means of reducing the cumbersome and expensive ground-communications link requirements, as made necessary, in part, by ground-based radio navigation. By relying on a visual science-quality instrument onboard the craft, these science ships could determine their own position, independent of an Earth-provided radio beacon. Another development enabling an autonomous optical navigation system is the increasing importance and attention to the orbits of the minor planets, which are the principal observational beacons of such a system. With the increased concern of possible Earth impact with Earth-crossing asteroids or comets, an international network of asteroid observers has evolved to track newly discovered objects, as well as to take data on older ones of interest. Accurate determination of the beacon-asteroid ephemerides is an important first step in building an autonomous optical-navigation system.

Autonomous optical navigation was chosen as one of the prime technologies to demonstrate onboard DS1. Furthermore, it was accepted as the principal means of navigation for both cruise and encounter, operation of the ion propulsion system (IPS), and execution of the encounter events. Since navigation of a deep-space probe using continuous low-thrust propulsion had never been done manually or autonomously, there were substantial challenges presented to the DS1 AutoNav team. Additional challenges were the use of a new-technology imaging system, the Miniature Imaging Camera and Spectrometer (MICAS), and the development of operations techniques for a fully autonomous flight system (AutoNav) within the context of a conventionally commanded and sequenced spacecraft.

2.0 TECHNOLOGY DESCRIPTION

2.1 Technology Overview

DS1 AutoNav is an onboard, autonomous, optical-navigation system. When used onboard a spacecraft with an adequate imaging system, AutoNav is designed to autonomously determine the position of the spacecraft using images of distant asteroids. AutoNav then will compute changes to the spacecraft course using the scheduled IPS thrusting profile (if present) or with discrete trajectory correction maneuvers (TCMs). Finally, AutoNav will direct the terminal tracking activities at the closest approach. These high-level activities are accomplished through the following actions and responsibilities:

- Provide ephemeris information to other spacecraft subsystems.

- Plan and execute image taking sessions by
 - Developing an Image-Taking plan from an initial “suggested” target list.
 - Communicating with the attitude control system (ACS) to get specifications of turns.
 - Executing turns and requesting pictures be taken.
- Process pictures and reduce the image data to astrometric-geometric information.
- Combine pictures into a data arc and perform a batch-sequential least-squares solution of spacecraft position and velocity.
- Compute course correction:
 - Propagate current spacecraft state to target and compute impact-plane error.
 - If in a mission burn, compute changes to the burn direction elements, and burn duration.
 - If there are TCM opportunities, compute the magnitudes and durations of each TCM.
- Execute a mission burn:
 - Communicate with the ACS for spacecraft turn specifications.
 - Turn the spacecraft to the correct attitude.
 - Start the main engine and maintain a mission burn with periodic direction updates.
 - Terminate the burn after the appropriate thrust has been achieved.
- Execute a trajectory correction maneuver:
 - Communicate with the ACS for spacecraft turn specifications.
 - Turn the spacecraft to the correct attitude.
 - Start the main engine, or request that ACS perform a ΔV event.
 - Optionally, turn to a second TCM attitude and execute the second segment.
- Perform terminal tracking and encounter operations:
 - Process close-approach images of the target
 - Reduce and filter the picture data.
 - Estimate a target relative state and communicate information to ACS.
 - Start encounter sequences at the appropriate time.

2.2 *AutoNav Technology-Validation Plans and Objectives*

2.2.1 *AutoNav Validation Plan Overview*—Before detailed operations planning for DS1 took place (indeed, long before even the encounter targets had been selected), AutoNav was undergoing development, testing, and validation. These early tests were performed on platforms far different from the actual spacecraft and, as such, were not considered a formal part of the validation plan. Nevertheless, they were a crucial part of the system validation, and will be discussed in some detail in section 3.1.

As has been stated, in the early design phases of the DS1 mission, it was decided to make AutoNav the primary means of navigation for the mission. As such, the driving assumption for planning purposes was that the system

would be operational and would be used soon after launch. Accordingly, extensive planning was undertaken by the Mission Design, Sequence, and AutoNav Teams to construct an operations plan that took full advantage of the capabilities of AutoNav. Figure 1 shows an early version of this plan (for an October 15, 1998 launch). (This diagram was produced by Pam Chadbourne, of the Mission Design Team, as part of that team’s continuous and very successful efforts to plan and schedule the myriad of interconnecting events and processes that comprised DS1 operations, including the technology validation.) Though the actual launch was 9 days later than shown, changing various aspects of the plan (especially the length of the IPS thrust arcs), the layout of events is very representative of the final plan and gives a good impression of the timing and interaction of the validation plans with each other and particularly with AutoNav.

Immediately upon booting of the spacecraft computer as part of the launch sequence, AutoNav would begin its simplest, but, in a few respects, its most important operation and test; and that test would be to provide ephemeris information to the ACS. Without this service properly completed, the spacecraft could not achieve a normal post-launch state and could, in fact, be endangered. Therefore, the validation of AutoNav would commence in earnest within minutes of launch.

Despite this early “must-work” requirement upon the ephemeris server, it was acknowledged that the higher functions of AutoNav (picture taking and processing, orbit determination (OD), etc.) would not be immediately credible. Furthermore, even if fully operable and immediately invoked, AutoNav was not capable of performing the higher-accuracy near-Earth navigation (from immediately after launch to launch plus 2 days) required to assess injection conditions and keep the very spacecraft-position-dependent near-Earth DSN tracking within specification. Consequently, “conventional” radio navigation would guide DS1 “out of the harbor” and, in fact, would continue for the entire 1992KD cruise, being the only independent means of assessing AutoNav orbit determination (OD) performance. (And, in fact, as the actual mission proceeded, there was much dependence upon the radio-navigation function, as AutoNav was validated, but, more importantly, as various and many problems with other subsystems were resolved or work-arounds attempted.) The development of radio navigation techniques for use with a low-thrust mission was a technology development in and of itself. However, the documentation of this important technology has not yet commenced; even an overview of this extensive body of work is beyond the scope of this document. However, those interested can contact Tim McElrath, Mark Ryne, and Don Han at the Jet Propulsion Laboratory for further information about the outstanding work achieved with DS1 radio navigation.




Figure 1. DS1 Mission and AutoNav Operations/Validation Plan and Schedule

It was anticipated that within two weeks of launch, AutoNav would be performing regular OpNav events three times per week. These events (Photo-Op/OD/ManPlan—see section 2.4) would continue at this high frequency for about six weeks, during which time validation and verification of the system would take place. See Figure 1 for a complete overview of all of the validation events. Following those six weeks would be a more relaxed schedule of once per week; this would be roughly coincident with the beginning of the first IPS mission burn thrust arc and the validation of another component of AutoNav, the autonomous operation of the IPS.

The means of verification of system performance depended upon the particular AutoNav function. As stated above, for the crucial measure of accuracy of the orbit determination, ready comparisons with ground-based radio navigation could be made. For other subsystems and functions, AutoNav performance was either self measuring or required parallel and independent measure on the ground using elements of the ground optical-navigation system. This will be discussed further in the next section (2.2.2).

Throughout the IPS burn segments, OpNav operations were to continue (with the main engine shutting down for purposes of picture taking and subsequent telecom), along with

adjustments to the spacecraft course through maneuver planning (Man_Plan) events.

Validation of the encounter operation onboard was planned to be through the validation of those operations in common with cruise (e.g., Photo_Ops) and with a dedicated rehearsal of the encounter a month or so before the actual encounter (see Figure 1). This rehearsal had been planned to be 2 days of operations mimicking those of the real encounter operations. An essential part of the validation was the ability of AutoNav to simulate, onboard, incoming optical data. This provided the capability to “spoof” the entire ensemble of spacecraft elements into thinking an actual encounter was taking place. Success of the terminal approach and tracking system (discussed at length below) was self assessing, in that AutoNav either “locked on” and tracked or did not; in other words, the validation criteria was “binary”, as opposed to quantitative.

Figure 2 shows the intense schedule of planned navigation validation events for the two days approaching encounter. Of particular note are TCMs and the Reduced State Encounter Navigation (RSEN) initialization events.

2.2.2 AutoNav Key-Point Technology Description and Validation Strategy—The AutoNav Technology Validation Key Point Summary table on the Fact Sheet refers to a number of key elements of the validation plan that are broken out as individual items for which flight-validation observables were expected and agreed to (see Appendix F, the Technology Validation Agreement). Additionally, some of these items have quantifiable metrics: requirements in the Technology Validation Agreement, internal requirements of normal spacecraft function, or strong “desirements” of the AutoNav team. Following is a description of the meaning, content, and validation strategy of each of these elements.

Navigation and Related Validation Events

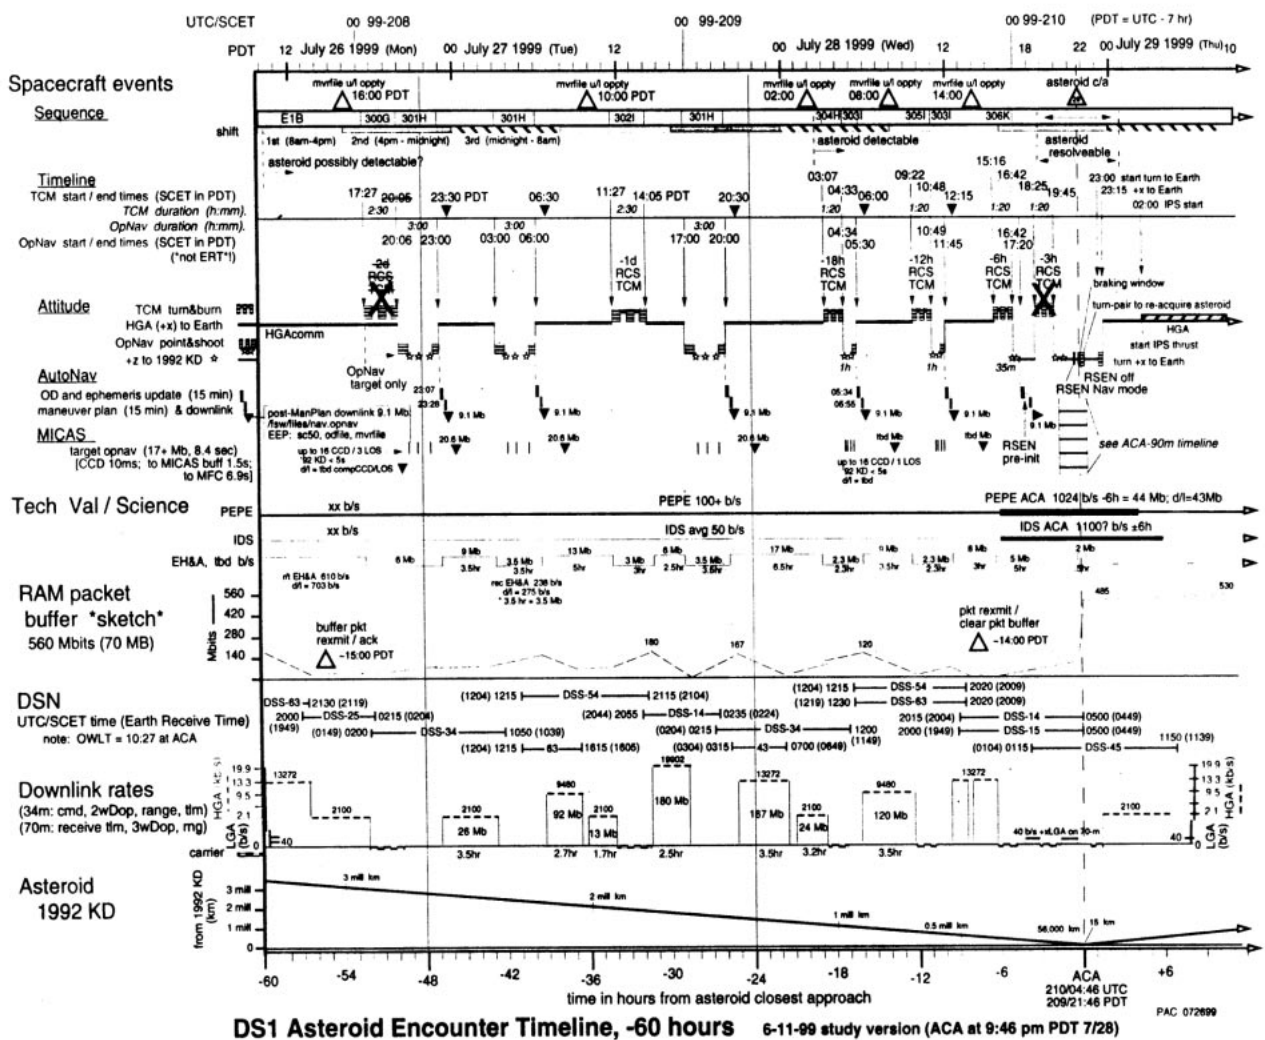


Figure 2. Mission and AutoNav Operations/Validation Plan and Schedule for Braille Close Approach

2.2.2.1 Provision of Ephemeris Services—This is the required function to provide various onboard systems (chiefly ACS) information about the location of the spacecraft and any solar system object of importance to the mission, such as Earth (for telecommunications purposes), and other solar-system bodies for camera targeting. The quantitative measure of the validity of this system is effectively the interpolation error of the Chebyshev polynomial ephemeris files provided from the ground or generated onboard. In effect, no error beyond computational error is desired, but the absolute highest degree of accuracy is in the encounter time-frame, where a maximum of 100 m of error would be tolerable. Validation was by testbed proof and by spot checks onboard.

2.2.2.2 OpNav PhotoOp Process—This is the overall “Photo-Op Machine” subsystem of AutoNav. It entails the coordination and execution of the sub-tasks described in sections 2.2.2.2a through 2.2.2.2c. Validation of this process was by inspection: i.e., evaluation of the EH&A record state, noting the completion of the requested tasks and lack of any tripping of explicit or implicit error states in its own or external sub-systems.

2.2.2.2a Picture Planning—This function retrieves the appropriate “suggested” selection of asteroid beacons from the Picplan file and determines those that are appropriate for imaging given current mandated restrictions in the allowed viewing space of the sky. Validation is by inspection.

2.2.2.2b ACS/APE Interaction & Turn Planning—This function is the extensive network of interactions between AutoNav and ACS and its planning subsystem, Attitude Planning Expert (APE). ACS is queried for current states of the ACS; these results are used to construct the AutoNav sequences. APE is queried for turn specifications for the turns to the desired targets. Validation is by inspection and careful review of the EVR messages from the navigator, wherein the details of the interactions are downlinked.

2.2.2.2c Mini-Sequence Picture/Turn/Fault Execution—This function is the implementation phase of the Photo-Op. At the highest level, this function ensures that all operations are completed in the allotted time. For picture taking and turning, mini-sequences are built with the desired commands and launched into the sequencing engines (one of eight). Additionally, the progress of the Photo-Op is monitored and excessive back-logs of unprocessed pictures is prevented. Finally, this function provides for contingencies in the event of one of a subset of failures of the Photo-Op and recovery or abort action (short of calling the Fault Protection (FP) system). Validation is by inspection and careful evaluation of

downlinked EVRs, which document, in complete detail, these events. Note: In M6, this function ceased being done by mini-sequence and was thenceforth mediated by direct message calls.

2.2.2.3 Image Data Handling and Downlink—This function accomplishes the MICAS picture data handling for AutoNav. This handling involves the compression, deletion, and downlink of pictures as desired, with various levels of combinations of data quantity provided. Validation of this function is by inspection and by successful retrieval of downlinked and compressed pictures.

2.2.2.4 OpNav Data Accumulation, Handling, Down link—This function is the somewhat esoteric but critical process of filtering and compacting the data from the processed pictures, which resides on the OpNav file, onto the OD file. The filtering process attempts to delete bad data through ensemble statistical analysis. Another critical part of this function is to trim two important data files to be of appropriate length: namely, the NonGrav History File and the OD file. Validation is by inspection, through EVRs, and by ground processing of the OpNav and OD files.

2.2.2.5 Image Processing—As its name implies, this function is responsible for extracting useful navigation data from the onboard taken pictures. There are three stages to this process: (1) an initial course registration, wherein the a-priori prediction of the location of objects in the field, good to 10 to 20 pixels, is refined to 1 or 2 pixels; (2) then, precision astrometry takes place, where the locations of objects are determined to 0.1 to 0.25 pixel; (3) finally, using only the star images as reference, the inertial attitude of the camera when the image was taken is computed and that information, plus the location of the target, is written to the OpNav and, subsequently, the OD files. Validation is accomplished in several ways. Raw pictures downlinked can be reprocessed on the ground using related or independent software and the results compared to those of the flight system. Evaluation of EVRs is also very useful for analysis of the image processor.

2.2.2.6 Orbit Determination—This is the purely computational function of reducing the suite of optical observations on the OD file to an estimated state of the spacecraft. Sub-elements of this function include numerical integration of the spacecraft position and velocity as well as partial derivatives of the spacecraft state with respect to dynamic parameters. Of course, estimation and filtering itself is a key function. Validation of this function is in two phases: confirmation of correct action onboard by repeating the onboard computations in the context of ground versions of the flight software and comparisons of the actual computed states with those of radio ground system. Pre-launch analysis indicated that, given nominal camera performance, it would be possible to achieve OD accuracies during the cruise phase of 250 km and 1 m/s in position and velocity respectively; these were the agreed-to

standards in the Technology Validation Agreement (Appendix F). A complete analysis of the expected performance of the OD subsystem is given in Appendix D.

2.2.2.7 Generation of Onboard Ephemeris and Downlink—This function takes the freshly computed solution from the OD function and integrates a new spacecraft ephemeris, produces a file (Spacecraft Ephemeris) of same, and makes this file available to Ephemeris Services. This function is also performed after a maneuver plan. Validation is by inspection, EVR analysis, and evaluation of the downlinked files. The Chebyshev polynomial fitting process has precision requirements. Over a one-month integration, the desire was 1-km precision. For encounter, the requirement was much tighter: only 100 m in a 1-day integration was tolerable.

2.2.2.8a and b Trajectory Control and Maneuver Planning—This is the purely computational function of computing a course correction using a mission burn or a TCM. Computational elements involved in this function include iterative trajectory integration to compute a-priori mistargeting and numerical partial derivatives for the estimation of correction parameters. These parameters can be the elements of a discrete RCS or IPS TCM or the directional and duration parameters for an IPS mission burn. Additionally, the Maneuver Planner must determine, through interaction with APE, whether a proposed TCM is “legal” in the context of spacecraft orientation constraints. If there is a violation, further interactions with APE will decompose the TCM into two allowed legs, via a process called “vectorization.” Given correct nominal computational behavior and the input of a suitably accurate OD, the maneuver calculation is self-assessing, by either converging to a suitable solution or not. The criterion for success is, nominally, a 1-km error in the targeting plane. Assessment of success is by inspection, EVR, and ground evaluation of the downlinked Maneuver file.

2.2.2.8c TCM Execution and Delivery—This is the executive function of a TCM. Similar ACS, APE and mini-sequence interactions and operations as were described above (2.2.2.2b, c) take place here. This function must ensure that all operations are complete within the allotted time, including turns to burn attitudes, executions of the burns themselves (either IPS or RCS), and a turn to the desired “home” attitude. For the final approach TCM, assumed to be 3 hours from closest approach, with a closing velocity of about 15 km/s, performance specifications for execution (really a measure of combined OD, ManPlan, and TCM execution) were set at 2.5 km and 0.25 m/s for the targeting plane position and velocity. Validation is via inspection and

EVRs; however, final delivery accuracy requires indepth post-encounter reconstruction and evaluation (in simulation mode, the success criteria is available by inspection).

2.2.2.9 Execution of Mission Burns—This function is that which accomplishes the operation of the IPS during the mission burns. There are several subfunctions, including ACS and APE interaction (much as was described for the Photo_Op and TCM functions), interactions with IPS (e.g., starting, stopping, pressurising, setting throttle levels, and safing the engine). Lastly, the mission burn function contains the overall management function of coordination of activities of the mission burn. This management includes evaluation of the navigation files to determine the proper direction and duration of the burning and the starting and termination of the burns. Validation is by inspection and EVR evaluation.

2.2.2.10 Encounter Image and OD Operations (RSEN)—This function is the overall control and coordination function of the AutoNav close-approach Nav function, Reduced State Encounter Navigation (RSEN), and includes initiation and termination of RSEN mode, receipt and delivery of pictures to the RSEN picture processing module, and ultimate dispatch of the pictures following image processing. Validation is by inspection and EVR evaluation.

2.2.2.10a RSEN Image Processing and Data Reduction—This function is responsible for the reduction of APS pictures during the encounter. To an extent, this function is self-evaluating by reporting—through EVRs—the success of the reduction of the pictures. The precise numerical validation of the result must be determined through thorough evaluation of ground-analysis tools, in particular ground versions of the flight software. In test mode, however, the quantification of the validation happens “automatically” in the sense that the OD solutions derived from each individual picture should match the input state deviation. This deviation is the difference between the spacecraft’s “best guess” of its current position and the “truth” as known by the simulation software.

2.2.2.10b Computation (and Delivery) of Target Relative State—Given the successfully generated results of the image-processing function described above, this function performs the reduced-state orbit-determination operation and transmission of the data to ACS for tracking of the target. As with the previously discussed functional element, to some extent this function’s success is self-checking and reporting. However, again, precise numerical consistency is validated with ground repetition of the flight processing; also, as above, when in self-simulation mode, the OD answers should be driven (within statistical deviation due to digitization and spatial quantization of the picture field) to the “truth” held by the self-simulation system. Figure 3 shows the expected accuracy of the RSEN system in downtrack (i.e., time-of-flight) on approach to Braille given successful picture delivery and processing at each of the indicated data. Note that two

different a-priori errors were assumed, 10 and 20 s, representing 150 and 300 km of downtrack error respectively. In fact, the actual error was probably closer to 300 km based on the ephemeris errors observed in the cross-track directions during the actual Braille approach. Figure 3 shows a complicated and continuous representation of the expected RSEN performance, which was distilled down to the specific quantities in item 11 of the Fact Sheet table and mentioned as a system-validation requirement in Appendix F.




Figure 3. RSEN Time-of-Flight Performance

2.2.2.11 Initiation of Encounter Sequences—The final step in the encounter process is to start encounter sequences at a time appropriate for encounter science-data gathering. During a close flyby of the target, the acquisition of navigation knowledge about the relative downtrack position of the spacecraft happens only very late. Consequently, parts of the close-approach science activity must be broken up into segments, generally getting shorter as they approach close-approach, and each of these segments is started at an increasingly accurate determined time relative to close approach. The function that starts the encounter sequences is completely dependent upon the computational processes outlined in the previous two sections (immediately above) for the determination of expected time-of-flight. Given this information, this function, when asked to start an encounter sequence, immediately determines the time remaining to encounter and starts a mini-sequence to “launch” the desired sequence at the appropriate time. Validation is by inspection and EVRs; however, for the numerical accuracy of the starting times, validation is accomplished through the validation of the two previously discussed functions.

2.3 Expected Performance Envelope

The expected performance ranges of AutoNav, and how this system can be applied to other missions, is a complex issue. This issue will be addressed somewhat in Section 5, from the standpoint of modifications to the system for extended use. However, some of the quantitative issues will be addressed here. The most important thing to note is the complete dependence of an autonomous optical-navigation system such as AutoNav upon the camera system and other systems. In Table 3 are noted the operable ranges for the camera parameters for AutoNav use; the ranges are quite wide. Varying these parameters can have positive or negative influence on AutoNav performance; there is no “ideal” combination of settings, but only a continuous trade space that is mission dependent. Other subsystems have similar influence on other parts of AutoNav.

Figure 4 is a flowchart depicting the dependence and correlation of performance between AutoNav subsystems and external providers of data or services. Also shown are the dependencies on a very small sampling of AutoNav control parameters; where a positive correlation factor in one component is shown, it enhances the performance of the subsequent component, and *vice versa*.

With the exception of the basic correlations shown in Figure 4, it is nearly impossible to represent the full space of parametric influences on navigation performance. However, a few basic high-level statements can be made on the overall, but variable, capabilities of the system. First, the system is capable of maintaining an adequate navigation state in the cruise phase of most interplanetary missions, given an adequate camera (again, see Table 3) and given “reasonable” ACS performance. Second, flyby delivery to “a few kilometers” is reasonable under a wide range of conditions. Tighter delivery performance requires tougher camera requirements and/or modeling requirements on the target body. ACS performance improvement, particularly inertial attitude determination from the SRU or IMUs can boost delivery accuracy. Third, rendezvous missions present no more additional challenge to DS1 AutoNav than a flyby; in fact, a rendezvous is in many ways easier. All the events that occur during a flyby occur in a rendezvous, but vastly slower; the added time is a huge advantage. There are no different attributes of the targeting problem for navigation and trajectory control (even though the mission design issues are *very* different) between flyby and rendezvous. Fourth, for large body (planetary) approaches, for most of the planets, the AutoNav system of using small “asteroid-like” navigation beacons is applicable, using the small satellites. For Mercury, Venus and Earth, additional software would be necessary to accurately determine the positions of very large, textured, and possibly “hazy” planets. It should be pointed out that the original mission plan of DS1 included a flyby of Mars, where Phobos and Deimos were to be used as targets.




Figure 4. Subsystem Performance Influence on AutoNav

2.4 Detailed Technology Description

2.4.1 *The AutoNav System*—AutoNav is a file-based computational system. Conditions necessary to operate AutoNav—for example, operational parameters, planetary ephemerides, star catalog, etc.—are provided by the ground operators. This information provides AutoNav with sufficient information to start gathering its own data by scheduling and taking pictures. AutoNav updates these data as necessary as a means of storing computed information and communicating between the AutoNav links. A table of the AutoNav files and their update frequency (by AutoNav and the ground) is given in Table 1.

2.4.2 AutoNav File Descriptions—

2.4.2.1 *Star Catalog (Starcat)*—The Starcat is a file that contains the positions and brightnesses of the stars necessary for navigation. For DS1, this file contained 220,000 stars in an annulus of ± 30 degrees of the ecliptic and as deep as stellar magnitude 10.5. This catalog was extracted from a hybrid catalog comprised of the Astrophraphic-Tycho Catalogue combined with Hipparcos data.

2.4.2.2 *Planetary Ephemeris*—The planetary ephemeris contains the positions of nine planets and the Moon represented as Chebyshev polynomials. This file extends for the duration of the primary and extended missions and is based on the JPL DE-403 planetary ephemeris.

2.4.2.3 *TCM Params*—This file contains parameters that moderate the function of the TCM activities. These parameters include the minimum wait times between turns and actual burns of the RCS and IPS engines and parameters such as timing and control.

2.4.2.4 *Encounter (RSEN) Params*—This file contains parameters that regulate the activity of the close approach navigation system (RSEN).

2.4.2.5 *Encounter Star Catalog*—This file contains a small star catalog that is used only for the far-encounter navigation-image processing. A separate catalog is necessary to process the encounter pictures because of the geometry of the approach (e.g., outside the main catalog annulus) or because of the depth of stars necessary to include.

Table 1. AutoNav Files, Sizes, Autonomy Status, Locations, and Update Frequency

File Description	File Size (KB)	File Update Frequency		Location
		From Ground	Auto-Onboard	
Star Catalog	2200	1/mission	Never	EEPROM
Planetary Ephemeris	92	1/mission	Never	EEPROM
TCM Params	5	4/year	Never	EEPROM
Encounter (RSEN) Params	0.3	2/encounter	Never	EEPROM
Encounter Star Catalog	0.1	2/encounter	Never	EEPROM
FrankenKenny Params	0.7	2/encounter	Never	EEPROM
CCD Camera Params	0.6	2/year	Never	EEPROM
APS Camera Params	3	1/encounter	Never	EEPROM
Beacon Ephemeris File	2	2/year	Never	EEPROM
Mass Profile	56	4/year	Never	EEPROM
Picture plan	20	4/year	Never	EEPROM
Control Params	20	4/year	Never	EEPROM
Photo-Op Params	4	2/year	Never	EEPROM
IPSBurn Params	0.4	2/year	Never	EEPROM
Nongrav Params	0.2	2/year	Never	EEPROM
Imageproc Params	0.3	2/year	Never	EEPROM
File of Filenames	1.5	4/year	1/month	EEPROM
Maneuver	33	4/year	Weekly	EEPROM
OD	10	2/year	Weekly	EEPROM
Spacecraft Ephemeris	12	1/year	Weekly	EEPROM
OpNav	1000	Never	Weekly	RAM
Non-grav History	40	Never	Several/day	EEPROM

2.4.2.6 *FrankenKenny Params*—FrankenKenny is the onboard self-simulation subsystem of AutoNav. It creates images based (optionally) on an independent model of the spacecraft position and feeds these images to AutoNav, providing closed-loop simulation. This file contains parameters to control the simulation.

2.4.2.7 *CCD Camera Params*—This file contains parametric descriptions of the MICAS CCD camera, including focal-length and distortion models.

2.4.2.8 *APS Camera Params*—This file is as above, but for the MICAS Active Pixel Sensor (APS) visual channel of the MICAS camera.

2.4.2.9 *Beacon Ephemeris*—This file contains the Chebyshev polynomial description of several dozen asteroids used for navigation.

2.4.2.10 *Mass Profile*—This file contains a table of propellant consumption values; in essence, the predicted mass of the spacecraft at discrete times.

2.4.2.11 *Picture Plan*—The Picture Plan is a file that contains recommended asteroid targets, selected for maximum navigational strength and to minimize the amount of turn time required to move from target to target.

2.4.2.12 *Control Params*—This file contains dynamic modeling parameters for the spacecraft position

integration and targeting parameters (such as the desired flyby conditions). This file also contains parameters used by the orbit-determination routines.

2.4.2.13 *Photo_Op Params*—This file contains the parameters to control the “Photo-Op” operation, the Nav-controlled events that cause navigation images to be taken and processed. These parameters are primarily timing parameters (e.g., delays after turns).

2.4.2.14 *IPSBurn Params*—This file contains the parameters to control the operation of the Nav-directed mission burns, which are long periods of IPS thrusting. These parameters are primarily timing parameters (e.g., delays after turns).

2.4.2.15 *Non-grav Params*—This file contains parameters to direct the writing of the Non-grav History file that has a continuous record of intentional “non-gravitational” events onboard accomplished by the ACS or IPS. These parameters largely regulate the precision in time with which this record is kept.

2.4.2.16 *Imageproc Params*—This file regulates the operation of the image-processing operation, with controls such as thresholds for brightness and filtering gains.

2.4.2.17 *File of Filenames*—This file is the navigation directory, containing the full path-names of all of the navigation files, thereby indicating their locations in the file system. This file is automatically updated when files are updated using the Nav_Data_Update mechanism.

2.4.2.18 Maneuver—This file contains the descriptions of thrusting events, such as TCMs and mission burns. It also divides up “time” into segments for purposes of OD processing. The Maneuver file is autonomously updated by the Nav_ManPlan maneuver-planning function.

2.4.2.19 OD—The OD file contains the current best estimate of the spacecraft position at several junctures in time through the data arc (typically a month). This file is autonomously updated during the Nav_Do_OD orbit-determination function.

2.4.2.20 Spacecraft Ephemeris—This file is a Chebyshev polynomial representation of the spacecraft position and velocity. This file is automatically updated after the Nav_Do_OD and Nav_ManPlan functions.

2.4.2.21 OpNav—This file contains the results of image processing in the Nav_Do_PhotoOp function: edited picture elements, and determined line/pixel positions of objects.

2.4.2.22 Nongrav History—This file contains the continuous record of intentional “non-gravitational” (i.e. thrusting) events onboard accomplished by the ACS or IPS.

2.4.3 Software System—The AutoNav software architecture is shown in Figure 5. The AutoNav system is comprised of three principal parts: the Nav Executive, Nav Main, and Nav Real-Time (NavRT). These parts communicate with each other and with other subsystems through the underlying system-messaging facility. Much of the commanding by AutoNav is through the sequencing subsystem, as will be discussed below.

2.4.3.1 Nav Executive (NavExec)—NavExec is AutoNav’s director of spacecraft activities. It receives messages from other spacecraft subsystems and sends command directives, either through the onboard sequence machine or through direct messages, to other subsystems. When using the sequence subsystem (sequence engine), NavExec will build small sequences and “launch” them. When NavExec needs an activity to occur immediately (for example, to turn the spacecraft to a desired burn attitude), it will build a relative time sequence that the sequence engine initiates at once. Alternatively, when NavExec needs to ensure that an event begins exactly at a certain time, it will build and initiate an absolute timed sequence (for example, to cause the main engine to ignite for a TCM). NavExec contains three main state machines: for Photo-Ops, for TCMs, and for mission burns. These machines are mutually exclusive, the activities involved being clearly incompatible.

2.4.3.2 Nav Real-Time (NavRT)—NavRT is the subsystem of AutoNav that provides critical onboard ephemeris information to other onboard subsystems, but principally to ACS. NavRT operates at a much higher priority level in the flight software than the other AutoNav components due to the need to respond to sometimes frequent and time-critical ACS requests. NavRT also accomplishes file updates, involving ephemeris-related files, by ensuring that changes in files are completed in a way as to not jeopardize ACS ephemeris queries.

2.4.3.3 Nav Main—Nav Main, or just plain “Nav,” is the central computing element of AutoNav. Requests for activity that involve large amounts of computing are either directed to Nav by NavExec or go to Nav directly through the command subsystem. These functions include picture processing requests from NavExec, Do-OD, and ManPlan commands from ground commands. There are several important sub-functions of Nav: trajectory integration, which includes dynamic modeling of gravitational and non-gravitational forces acting on the spacecraft; data filtering, including a U-D factorized batch-sequential filter, and trajectory update computation, which is based on an iterative linear minimum-norm solution for changes to the IPS thrust profile to reduce projected targeting errors.

2.4.4 AutoNav Commanding Strategy—DS1 AutoNav is fully autonomous only by the invitation of ground controllers. Most importantly, AutoNav will cause physical spacecraft activity or intense computational action only when invited to do so by the ground, allowing controllers to be fully aware beforehand when such activities will occur; however, the particulars of each of these events will likely not be completely predictable. For the three autonomous events that involve onboard-engineered sequences of turns, thrusting, or picture taking, the ground limits AutoNav to predetermined periods of time, allowing careful budgeting of onboard time, instrument, and computational resources. Table 2 is a summary of the AutoNav commands. Following is a brief description of each of the AutoNav Commands and its action.

2.4.4.1 Nav_Do_OD—This command causes Nav to: (1) trim the OD file data arc to the predetermined length, (2) trim the history file to a corresponding length, (3) compute data residuals and partials for all data points in the data arc, (4) estimate position, velocity, and non-grav parameters for the spacecraft state for each segment of the arc, (5) repeat steps 3 and 4 iteratively until converged, (6) write these solutions on the OD file, (7) integrate the current best estimated spacecraft state forward to a pre-specified time (usually about a month into the future), and (8) write this to the spacecraft ephemeris file.

2.4.4.2 *Nav_Do_TCM*—This command causes Nav to perform a TCM by (1) obtaining the pre-computed specifications for the next TCM from the Maneuver file, (2) checking that there is a TCM scheduled within a specified time (e.g., 1 hour), (3) querying ACS for the specifications of the turn to the attitude of the burn, (4) commanding ACS to perform the turn, (5) if the TCM is an IPS TCM, commanding IPS to thrust for the specified time, at the specified thrust or, if the TCM uses the RCS, commanding ACS to perform the specified impulsive ΔV , (6) if there is a second (e.g., vectorized) element to the TCM, performing steps 1 through 6 on this leg, and (7) commanding ACS to turn the spacecraft to the terminal attitude.

2.4.4.3 *Nav_IPS_Off_Mes*—The ground uses this command to inform AutoNav that IPS thrust has been forced off. This will terminate the Mission Burn State Machine, if active.

2.4.4.4 *Nav_Man_Plan*—This command causes AutoNav to compute the propulsive plan for the next control opportunity on the Maneuver file, if any. This may be an RCS or IPS TCM or an IPS mission burn.

For a mission burn, ManPlan will cause AutoNav to (1) propagate the last spacecraft state entry on the OD file to the B-plane, obtaining the current miss vector, (2) starting with a fixed number of mission burn segments, compute the partial




Figure 5. The AutoNav Software System and Interacting System Software

Table 2. Summary of AutoNav Commands

Command Name	Description	Arguments	Usage	Time required
Nav_Do_OD	Perform Orbit Determination	none	1/week	10–100 min
Nav_Do_TCM	Execute a TCM	duration	1/week	1.5–24 hr
Nav_IPS_Off_Mes*	Notify Nav of a forced “engine off”	none	1/week*	1 s
Nav_Man_Plan	Perform Maneuver Planning	none	1/week	10–200 min
Nav_Photo_Op	Perform a nav picture taking and processing session, edit and store data.	duration	1/week	1.5–8 hr
Nav_Reset*	Stop all Navexec state machines	none	Seldom*	1 s
Nav_Set_IPS	Start a Mission Burn	none	1/week	5 min
Nav_Start_Enentr	Start an encounter sequence	seq. ID	4/encounter	1 min
Nav_Update_IPS	Update the thrust vector during a mission burn	none	2/day	1 min
Nav_Change_Mode	Change an AutoNav operating mode	Data vectors	2/month	5 s
Nav_Data_Downlnk	Downlink a Nav file	file ID	2/month	20 s
Nav_Data_Update	Update a Navigation file	file ID	2/month	20 s
Nav_IPS_Press	Pressurize the main engine	none	1/week	1–30 min
Nav_ACM_Infoturn	Optional desired pointing of the spacecraft after a nav event	“turnspec”	1/week	5 s
Nav_BBC_Deadband	Optional desired deadband of the spacecraft after a nav event	deadband	1/week	5 s

*Contingency or emergency back-up command

derivatives of B-plane impact position and time with respect to burn angles of each segment and the duration of the final burn, (3) estimate the changes in the burn angle and last-segment-duration, (4) check the estimated angle changes for violations of pointing constraint (if a violation occurs, then that angle is reset to the constraint limit), (5) iterates, using steps 1 through 4, (6) if after a fixed limit of iterations, step 5 has not converged (i.e., targeting is not “close-enough”), adds mission burn segments to the set being updated, and repeats steps 1 through 6, and (7) if the solution converges, then overwrites the Maneuver file with the updated plan; otherwise, if there is no convergence, leaves the Maneuver file unchanged.

For a TCM, ManPlan will cause AutoNav to (1) propagate the last spacecraft state entry on the OD file to the epoch of the next maneuver, (2) compute from that epoch to the next encounter, the state, and state partial derivatives, (3) compute the required ΔV at the maneuver time, (4) repeat steps 2 and 3 iteratively until converged, (5) determine, via interaction with ACS whether the desired burn direction violates spacecraft constraints, (6) if so, ask ACS to “vectorize” this TCM (i.e., decompose the desired—but constrained— ΔV direction into two allowed directions), and (7) via steps 2, 3, and 4 compute the ΔV associated with each vectorized leg. In both of these cases, a new spacecraft trajectory is computed and written to the Spacecraft Ephemeris file.

2.4.4.5 Nav_Photo_Op—This command causes AutoNav to (1) cycle through its list of candidate “beacon” asteroids, taking each in turn, (2) for each asteroid, query ACS for the turn specifications to take the MICAS boresight to that attitude, (3) before turning, determine that there is sufficient time to turn to target, take the required pictures, and turn back to the desired terminal attitude, (4) if there is sufficient time, turn the spacecraft, (5) begin taking a sequence of pictures, sending each when complete to the AutoNav picture processing element, (6) as each picture is processed, write its reduced data (asteroid pixel, line, pointing values) to the OPNAV file, as well as edited picture elements, (7) cycle to the next asteroid target, via steps 2–5, (8) when the list of candidates is exhausted, or the available time (as communicated in the command argument list) is exhausted, command the spacecraft to turn to the terminal attitude, and (9) filter the contents of the OPNAV file for bad data and place the results in the OD file, where the OPNAV file is optionally scheduled for downlink and deletion.

2.4.4.6 Nav_Reset—This command causes any of the three AutoNav state machines—PhotoOP, MissionBurn, or TCM—to reset to the off state, if they are active.

2.4.4.7 Nav_Set_IPS—This command causes the initiation of a mission burn by (1) reading the Maneuver file and determining that a mission burn begins within a specified time, (2) querying ACS for the specifications of the turn to the burn attitude, and (3) building and starting a sequence to start at the mandated burn-start time (or immediately, if the “Set” command has occurred within a burn segment) that turns the spacecraft and commands IPS to go to a thrusting state, at the appropriate throttle level and for the specified duration.

2.4.4.8 Nav_Start_Encntr—This command causes AutoNav to build and start a sequence that in turn starts the specified sequence at the requested encounter relative time (see RSEN description below). This command is only operable while RSEN is active.

2.4.4.9 Nav_Update_IPS—During a Mission Burn (i.e., after a Set_IPS command) this command will cause Nav to update the current burn direction according to the Maneuver file.

2.4.4.10 Nav_Change_Mode—This command updates various control-mode flags and constant settings in AutoNav. The flags and variables so set are those that need to be changed frequently. The flags and variables may also be set due to changes in spacecraft state or mission phase. Other, more stable, parameters are kept in the parameter files.

2.4.4.11 Nav_Data_Downlnk—This command causes AutoNav to downlink a specified AutoNav data file (see section 2.4.2, AutoNav File Descriptions).

2.4.4.12 Nav_Data_Update—This command causes AutoNav to accept a specified AutoNav data file as replacement for an existing file. The AutoNav file of filenames is updated in this process (see section 2.4.2, AutoNav File Descriptions).

2.4.4.13 Nav_IPS_Press—This command causes AutoNav to command the IPS to pressurize the plena in preparation for thrusting at the throttle level determined from the Maneuver file.

2.4.4.14 Nav_ACM_Infoturn—This command allows the ground to inform AutoNav what the desired ACS turn specification is for the desired terminal attitude after a PhotoOp or TCM.

2.4.4.15 Nav_BBC_Deadband—This command allows the ground to inform AutoNav what the desired deadband is after a PhotoOp or TCM.

2.4.5 “Uncommanded” AutoNav Functions—

2.4.5.1 Reduced State Encounter Navigation (RSEN), and Encounter Sequence Activation—This AutoNav subsystem runs the encounter navigation activity. A Nav_Change_Mode command enables RSEN, whereupon the most recent estimated spacecraft state and covariance are mapped to the

current time. When an APS picture is received, RSEN is then activated, the state and covariance are mapped to the picture time by a simple linear motion propagation, the centroid of the target is located in the frame, differenced with a predict to obtain a residual, and a Kalman-filtered estimate of spacecraft position is made. Then, the cartesian spacecraft state is converted into “B-plane” coordinates, including linearized time of flight to closest-approach; the time-of-flight information is made available to other AutoNav subsystems. This process continues with subsequent pictures, with RSEN “boot-strapping” states from picture time to picture time (see Figure 6). When AutoNav receives a Nav_Start_Encntr command (wherein Nav is asked to start an encounter sequence at a specific time), the time of closest approach previously computed by RSEN is compared with the current time, and an absolutely timed sequence is built to start the desired sequence at the appropriate time.

2.4.5.2 Non-Grav History Accumulation—AutoNav must keep a continuous record of propulsive events by RCS and IPS onboard the spacecraft for purposes of accurately integrating the flightpath of the spacecraft. In this effort AutoNav is aided by the ACS and IPS software subsystems, which report periodically accumulated ΔV (in the case of ACS) or impulse (in the case of IPS). The periodicity of reporting varies for ACS, because this system buffers the accumulation, and only reports when a certain threshold is crossed (typically 10 mm/s). For IPS, the reporting is every minute. AutoNav further buffers this data under parametric control, writing “permanent”

records in EEPROM when accumulated ACS ΔV or IPS vector impulse cross internal AutoNav thresholds.

2.4.5.3 Ephemeris Services—Ephemeris Service is the highest priority AutoNav task and is required to give ephemeris information to ACS as often as on one-second intervals under some rare circumstances; however, ephemeris information nominally is queried every few minutes. The ephemeris server reads the ephemeris files of the spacecraft, the beacon asteroids, and the major planets. All of these files have Chebyshev polynomial representations of the orbital states, with velocities computed. All states are in Earth-Mean-Equator-2000 coordinates, as are the directions on the Star Catalog. Ephemeris Services also provide ephemeris data to the internal AutoNav functions.

2.4.6 Core Algorithm Descriptions—

2.4.6.1 Multiple Cross Correlation—Figure 7 shows a diagrammatic representation of the algorithm that forms the basis of the cruise-image processing in AutoNav. The underlying assumption of the algorithm is that long exposures will be necessary to image dim objects; therefore, because of ambient motions of the spacecraft due to attitude maintenance by ACS, the images of stars and targets will be smeared, often in complicated patterns. These patterns, called “glyphs”, will be nearly identical to one another, since the effects of “twisting” deadband motion in the field is small (the attitude maintenance is roughly equivalent in all directions, but maps to a much smaller effect in the field than the two cross line-of-sight pointing directions). Based on initial knowledge of pointing of the spacecraft (as provided by ACS) and

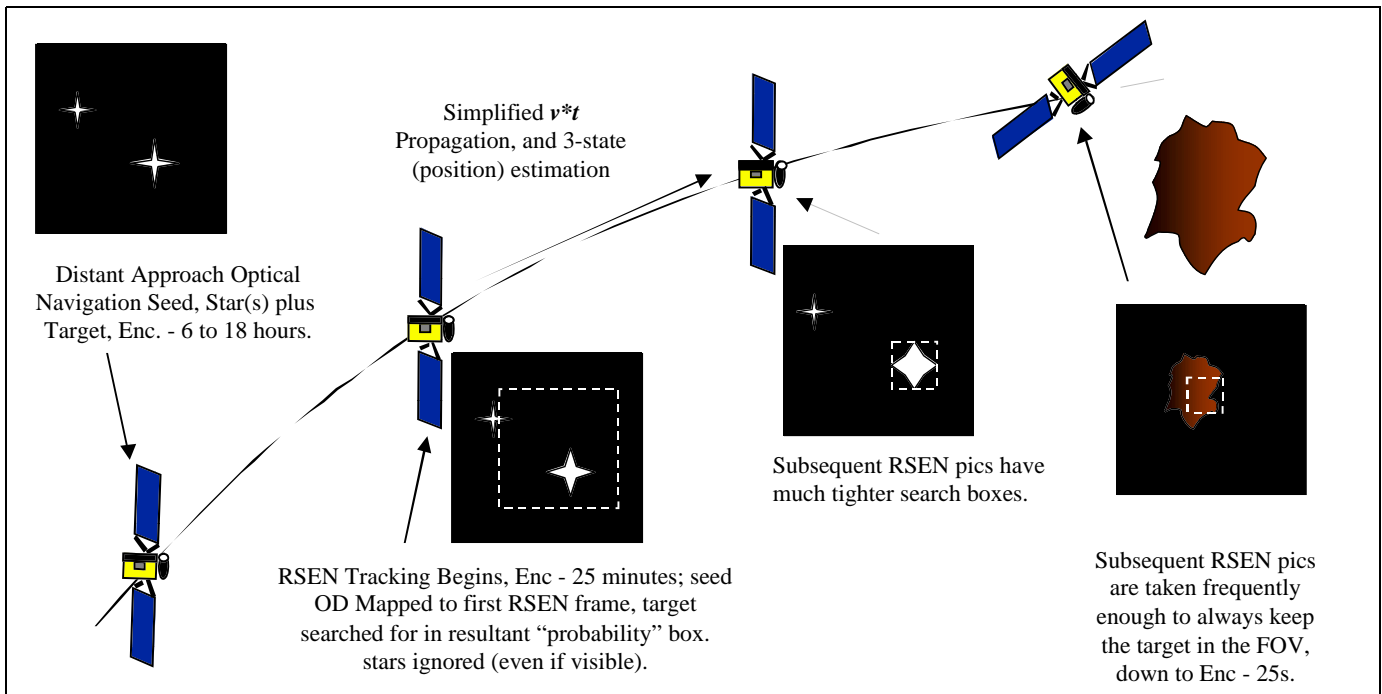


Figure 6. Reduced State Encounter Navigation Schematic Functional Overview




Figure 7. Multiple Cross Correlation of Asteroid and Stars

predictions of the relative locations of the objects in the positions of objects in the field of view (based on the target ephemeris and the star catalog), segments of the pictures are extracted, normalized, and become templates or “filters.” Filters for each object are used to locate each of the other objects. The “location” is accomplished through the convolutional inner-product of filter with data. Once all of the objects are located relative to one another (and these data are filtered for bad or weak signal), a least squares estimate is made of the relative offset of the objects relative to one another. A complete description of this algorithm is given in [5], as it was used for Galileo’s Gaspra encounter.

2.4.6.2 Orbit Determination—Figures 8a, b, c give an outline of OD and related algorithms as used by AutoNav. There are several crucial elements to the OD function: (1) the numerical integration of the spacecraft trajectory (Figure 8a), (2) the dynamic models of the gravitational and non-gravitational perturbations that drive that integration (Figure 8a), (3) the generation of σ and the mapping of the covariance in time with the state transition matrix (Figure 8b), and (4) the formation of the data filter itself (Figure 8c). Appendix D gives a complete development of the filter and related algorithms. As noted earlier, the OD filter used is a Kalman batch-sequential least-squares filter. A typical data arc is about a month long, with four 1-week batches that correspond to the typical one Photo-Op event per week. The estimated parameters for a given solution include the position and velocity at the beginning of the data arc, a constant acceleration 3-vector that applies for the duration of the

arc, and IPS thrust-scale factors that are stochastic parameters for each week. The latter parameters are in force only while there is an IPS Mission Burn in progress during that portion of the arc.

2.4.6.3 IPS Mission Burn Targeting—The process for retargeting the spacecraft trajectory during a mission burn is shown in Figure 9. This is an iterative application of a linear estimation of corrections to the direction of burn of an individual element of the multi-element mission burn and the duration of the final element. Since iterative, the overall algorithm is non-linear. The algorithm will automatically decide how many segments to include in the solution, starting with a minimum acceptable number and increasing the number as necessary to gain sufficient control authority to achieve convergence (i.e., putting the spacecraft on target).

It is important to note that the spacecraft is initially given a “converged” trajectory. This trajectory has been “discovered” and reasonably converged initially with an algorithm known as “differential inclusion” [6] and uplinked to the spacecraft. Then, within well-regulated limits, the maneuver planner is allowed to adjust this trajectory to keep the spacecraft targeted.

2.5 Technology Interdependencies

2.5.1 MICAS/AutoNav Interface—The principal AutoNav dependency on other technologies is with the imaging system. For DS1, MICAS is another “new technology,” with two visual channels: a somewhat conventional Charge Coupled Device (CCD) detector and a much smaller Active Pixel Sensor (APS). The ability to take high-quality astrometric

Dynamical equations of motion

- Includes central body acceleration, 3rd body perturbations from other planets, solar radiation pressure, thrust from the ion engines, and miscellaneous accelerations
- 2nd order differential equation modeled as two 1st order differential equations

$$\dot{\mathbf{r}} = \mathbf{v}$$

$$\dot{\mathbf{v}} = -\frac{\mu_s}{r^3} \mathbf{r} + \sum_{i=1}^{n_p} \mu_i \left[\frac{\mathbf{r}_{ri}}{r_{ri}^3} - \frac{\mathbf{r}_{pi}}{r_{pi}^3} \right] + \frac{AG}{mr^3} \mathbf{r} + \frac{k}{m} \mathbf{T} + \mathbf{a}$$

where

\mathbf{r} = the heliocentric cartesian position vector of the spacecraft
 \mathbf{v} = the heliocentric cartesian velocity vector of the spacecraft
 \mathbf{r}_{pi} = the heliocentric cartesian position vector of the i th perturbing planetary body
 \mathbf{r}_{ri} = the position of the spacecraft relative to the i th perturbing body
 μ_s = the gravitational constant of the sun
 μ_i = the gravitational constant of the i th perturbing planet
 n_p = the number of perturbing planets
 A = the cross-sectional area of the spacecraft
 G = the solar flux constant
 T = the thrust vector from the ion engine
 k = the thrust scale factor
 m = the spacecraft mass
 \mathbf{a} = miscellaneous accelerations acting on the spacecraft

(a)

Given \mathbf{q}^* , the nominal trajectory parameters, as

$$\mathbf{q}^* = [\mathbf{r} \quad \mathbf{v} \quad k \quad \mathbf{a}]$$

Filter estimates corrections, \mathbf{q} , to nominal trajectory parameters

$$\mathbf{q}(t) = [\Delta x \quad \Delta y \quad \Delta z \quad \Delta \dot{x} \quad \Delta \dot{y} \quad \Delta \dot{z} \quad \Delta k \quad \Delta a_x \quad \Delta a_y \quad \Delta a_z]$$

The correction at time t is a linear mapping of the correction from time t_0

$$\mathbf{q}(t) = \Phi \mathbf{q}(t_0)$$

where Φ , the state transition matrix, is defined as

$$\Phi(t) = \frac{\partial \mathbf{q}^*(t)}{\partial \mathbf{q}^*(t_0)}$$

(b)

The partial derivatives of the observed pixel and line locations, p , l , with respect to the state, at time t is

$$\mathbf{H}(t) = \begin{bmatrix} \partial p / \partial \mathbf{r} & 0_{1 \times 7} \\ \partial l / \partial \mathbf{r} & 0_{1 \times 7} \end{bmatrix}$$

This can be mapped back to the epoch, t_0 , via the state transition matrix

$$\tilde{\mathbf{H}}(t_0) = \mathbf{H}(t) \Phi$$

The minimum variance least squares solution to the epoch state corrections is

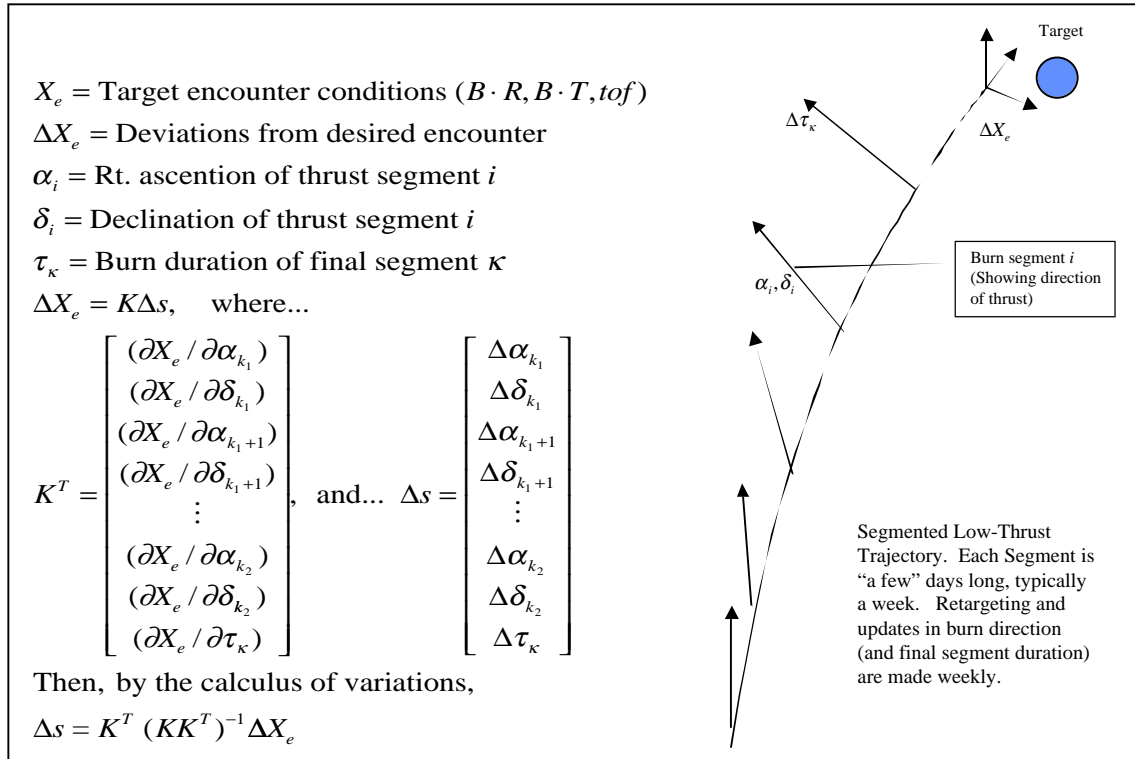
$$\hat{\mathbf{q}} = [\mathbf{P}_0 + \tilde{\mathbf{H}}^T \mathbf{W} \tilde{\mathbf{H}}]^{-1} \tilde{\mathbf{H}}^T \mathbf{W} \mathbf{Y}$$

where

\mathbf{P}_0 = the a-priori covariance of the state parameters
 \mathbf{W} = the weighting values of the pixel and line observables
 \mathbf{Y} = the residual vector between the observed pixel/line locations and their predicted values

(c)

Figure 8 a,b,c. Spacecraft Integration Equations of Motion and Derivation of AutoNav OD Kalman Filter


Figure 9. Adjusting a Low-Thrust Burn Arc
Table 3. Imaging System AutoNav Requirements and Attainment by MICAS

Requirement Description	Value Required	MICAS value	Attained
1. Digitization level	≥ 10	12	yes
2. Field of View	0.6 to 2.0	0.7/0.25 (APS)	yes/no
3. Array Size	≥ 512	1024/256 (APS)	yes/no
4. Geometric Distortion/Errors	$\geq 2 \mu\text{rad}$	$7 \mu\text{rad}$	no
5. Device fullwell and noise	80,000 e ⁻ /50 e ⁻	35,000 e ⁻ /40 e ⁻	no/yes
6. Dimmest obtainable image	magnitude 12	magnitude 9.5	no
7. Long-Exposure Capability	200 s	≤ 100 s	no
8. Encounter Imaging	Target and magnitude 9	Target and magnitude 7	no

images of small asteroids and image a bright, inner-solar-system target against a field of stars presents stringent requirements on a visual detector. The requirements listed in Table 3 were levied on MICAS; the table also indicates the level of success achieved in meeting these.

2.5.1.1 Overview of Camera Requirements and Attainment—Requirement 1 from Table 3 describes the gray levels obtainable in the instrument. 12-bit digitization, providing 4096 levels of gray, was implemented in both the CCD and APS channels, surpassing the requirement. Requirement 2, detector field of view, is met by the CCD, but not nearly by the APS. As will be discussed below, electronics faults in the CCD

channel required AutoNav to use the APS at the Braille encounter. Additionally (also to be discussed below), light leakage and scattered light internal and external to the camera caused the effective field of view to be reduced (severely at times) in the CCD. Requirement 3 was met by the CCD, but not by the APS. Requirement 4 is a complicated statement of the astrometric quality of the instrument. Factors that can effect this ability are the geometric distortion in the camera’s optics, their modelability, and their temporal and/or thermal stability. Observed post-launch distortions in the MICAS optics are well over $70 \mu\text{rad}$ in extent; due to the limiting dim magnitude of the camera, calibrations—so far—have been unable to improve this to better than 10%, or $7 \mu\text{rad}$. Requirement 5 is a statement about the dynamic range of the

instrument and the background noise. Because of the shutterless, fast-cycling readout design, the necessary range of useful signal was reduced in practice by about a factor of two from forecast, even though good noise characteristics were achieved. Requirement 6 was not achieved due to a combination of the reduced dynamic range, response-curve non-linearity, and scattered light (all discussed later). Requirement 7, the need to take long exposures to detect distant “beacon” asteroids, or the approach target, could not be achieved because of the magnitude of the scattered light problems. Requirement 8, the requirement to image the approach target with a navigation star, was not met for the same reasons, substantially limiting the approach-navigation strategies.

2.5.1.2 Other Camera Complications—Eight months before the launch of DS1, it was discovered that the CCD channel had a severe limitation when imaging bright objects (objects as bright as the first two expected targets). When the object of a typical asteroid brightness subtended more than 100 pixels (± 50), severe charge bleed appeared in the picture due to the inability of the CCD read-out to cope with the continuing photon flux during the read-out. Because of this limitation, it was believed that the CCD channel would be unusable during the last few minutes of approach. Figure 10 shows an example of the phenomena, taken during the instrument check-out, pre-launch. As a result of this problem, the less-capable APS channel was used by AutoNav on approach. In partial compensation, the read-out time required for the APS was much shorter than for the CCD, 2 vs. 20 s. At the first use of MICAS, it was apparent that there were substantial light-scattering problems around and in the camera [7]. Depending upon the sun-relative geometry, the CCD would saturate (achieve maximum measurable charge) in as little as 5 s of exposure. In view of the fact that the original feasibility analysis of AutoNav called for exposures as long as 200 s, this clearly represented a reduction in capability by limiting usable geometries and targets.




Figure 10. MICAS Extended Bright-Image Charge Bleed

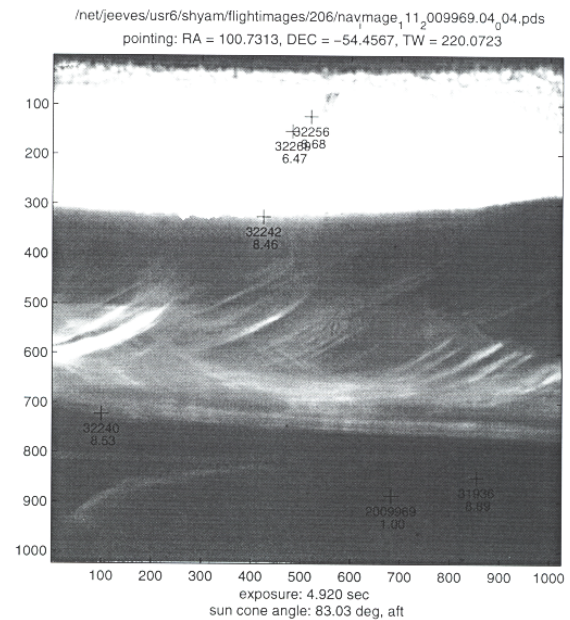


Figure 11. MICAS “Low Solar Cone Angle” Scattered-Light Picture

Figure 11 and Figure 12 show two examples of the scattered-light effect in roughly normal-to-Sun and anti-Sun geometries. A third difficulty with the camera is a highly non-linear response curve (see Figure 23 and the discussion of the encounter results in Section 3). The net effect of this electronics fault is for low flux signals to be non-linearly attenuated. This effect is much more severe in the APS, and largely accounted for abnormally low throughput at the Braille encounter. Another substantial difficulty for AutoNav arose due to light-attenuating scratches in the optics chain over a substantial portion of the CCD center-of-field-of-view. These can be seen as dark scars in the center of Figure 12.




Figure 12. MICAS “High Solar Cone Angle” Scattered-Light Picture

2.5.1.3 MICAS Software Interactions—In addition to the MICAS hardware, AutoNav interacts with the MICAS software subsystem. It is this software set that actually accepts and processes requests for pictures and provides them with important header information packaged in the picture file. Following is an example of such a header:

```

NJPL1100PDS          = XV_COMPATIBILITY
/* FILE FORMAT AND LENGTH
RECORD_TYPE          = FIXED_LENGTH
RECORD_BYTES         = 512
FILE_RECORDS         = 261
LABEL_RECORDS        = 5
/* POINTERS TO STARTING RECORDS OF MAJOR
OBJECTS IN FILE
^IMAGE               = 6
/* ANCILLARY INFORMATION
IMAGE_NUMBER         = 279
EXPOSURE_DURATION    = 0.013700
TARGET_NUMBER        = 5
ONBOARD_FILENAME     =
"/micas/images/ltc300_CCD_2.pds"
IMAGE_TIME           = 58028726.921814
SC_SUN_POSITION_VECTOR = {109905396.260058,-
129004901.095362,-56328752.753662}
SC_SUN_VELOCITY_VECTOR = { 19.890484,
17.517464, 7.523768}
SC_ATTITUDE_QUATERNION = { 0.325205,
0.512832, 0.767046, 0.207087}
DETECTOR             = "VISCCD"
IMAGE_USE            = "SCI"
READOUT_CLOCK        = "300KHZ"
MIN_COMPRESSION_RATIO = 1.00
UV_VOLTAGE_LEVEL     = 13
OBA1_TEMP            = -123.66
OBA2_TEMP            = -126.63
OBA3_TEMP            = -124.74
M1_MIRROR_TEMP       = -124.04
IR_RADIATOR_TEMP     = -165.26
OBA_CUBE_SUPPORT_TEMP = -124.20
IR_DETECTOR_TEMP     = -160.21
UV_DETECTOR_TEMP     = -5.90
ELECTRONICS_CHASSIS_TEMP = 29.52
    
```

```

COVER_ACTUATOR_TEMP    = -10.85
SUBIMAGE_X             = 132
SUBIMAGE_Y             = 640
CLIENT_DATA            =
0x0000000000000000000000000000000000000000000000000000000000000000
00000000000000000000
0
/* DESCRIPTION OF THE OBJECTS CONTAINED IN FILE
OBJECT                 = IMAGE
LINES                  = 256
LINE_SAMPLES           = 256
    
```

In addition to taking and providing the images, the MICAS software set also compresses images with varying ratios of “loss” from 1.0 (no loss) to small fractions. The software will also edit a picture to extract specified regions.

2.5.2 Attitude Control System (ACS)—AutoNav has mission-critical interfaces with ACS. Basic spacecraft health is dependent upon Nav providing ACS with the locations of the spacecraft and requested target bodies. Without this information, the spacecraft will be forced (under certain circumstances) into safing. In order to accomplish its autonomous activities, Nav communicates with ACS in several ways. Though not explicitly called out as a technology demonstration of DS1, the design and implementation of the DS1 ACS system contain a number of important technological advances. These include the operation of the IPS, attitude maintenance and turns with highly constrained attitudes, and autonomous turn planning for AutoNav. Categorized summaries follow.

2.5.2.1 Turn Planning and Execution—ACS’s Attitude Planning Expert (APE) is the service AutoNav uses to plan turns. When NavExec desires to change the attitude of the spacecraft, it queries APE for the particulars of the turn between the assumed beginning attitude and the desired attitude. APE will inform NavExec (1) whether the turn is possible at all, (2) whether it violates (or nearly violates) any pointing constraints, and (3) how long the turn will take. Armed with this information, NavExec decides whether to proceed. When a turn is commanded, it is accomplished with a turn specification (turn-spec) provided by APE. Additional attitude information is conveyed to ACS via updates to the IPS thrust vector (“TVC-pre-aim” vector), which causes ACS to effect small turns using the engine gimbals that point the throat of the ion engine.

2.5.2.2 Mode, Turn Mode, and Deadband Changes—During the course of its autonomous work, AutoNav has the occasional need to alter the operational state of ACS. These changes include changing from normal reaction control system (RCS) mode to thrust vector control (TVC) mode when operating the IPS is required. The mode that controls the pairs of thrusters used to turn the spacecraft must be set to allow for “slow” deadband maintenance during picture-taking is also altered. For most of the spacecraft actions AutoNav commands, the attitude-control deadband itself must be

changed to suit the activity. In addition, the ground generated sequence must set the family of constraints that proscribe areas on the spacecraft from Sun-illumination before certain AutoNav events.

2.5.2.3 Queries for Current State, and ΔV Estimator—As stated earlier, ACS periodically queries NavRT for ephemeris information. These queries always include a request for the spacecraft position and a request for the position of the body (if any) toward which the spacecraft is currently pointing. ACS also records all propulsive activity from the RCS and computes a net translational change in velocity (ΔV). When the value of this ΔV is greater than a predetermined value, a message containing the accumulation is sent to AutoNav and, after further buffering, these quantities are recorded on the AutoNav NonGrav History file.

2.5.2.4 Vectorization and ΔV Requests—Because of the Sun-illumination constraints (and geometric constraints involving keeping the solar panels focused on the Sun), it is impossible to point the spacecraft in certain directions. If it is necessary to accomplish a TCM in one of these directions, it is necessary to break the vector up into two components that are allowed. APE provides a service wherein AutoNav requests a ΔV direction and APE responds with one or two allowed directions for burning the engines. Upon receipt of this information, AutoNav recomputes the magnitudes of the burn elements if it has been vectorized. When the final values of the TCM have been computed, Nav turns the spacecraft (through interaction with ACS) and asks for an RCS ΔV or causes the IPS to burn for a specified time.

2.5.3 Ion Propulsion System—AutoNav has responsibility to perform basic operation of the IPS during mission burns and TCMs that use IPS. Additionally, IPS is responsible to report to Nav the progress of any IPS thrusting. Nav commands IPS through directives to pressurize at a given thrust level, ignite the engine, and stop and safe the engine. IPS, in turn, gives reports of the accumulated impulse over a one-minute period, and reports when the specified duration of the burn has been achieved. When this last message is received, Nav commands the engine to shut down. Accumulated IPS impulse is recorded on the NonGrav History file.

2.5.4 Remote Agent and RAX—Early in the development of the DS1 flight software there existed a high-level autonomous control system called Remote Agent (RA). A year and a half before launch, RA was de-manifested and many of the autonomous functions that were chartered to the RA were taken on by AutoNav. These duties include planning picture-taking sequences, managing the operation of IPS, and accomplishing TCMs, as well as

accomplishing the execution of encounter sequences. A greatly descope version of RA called RA eXperiment (RAX) was flown as a very short (a few hours) run during the prime mission. For the AutoNav-RAX interface, two simple data calls were created that provided RAX with the appropriate asteroids to target at a given time and the directions and thrust levels for a particular mission burn. These interfaces were implemented by simple reads of the AutoNav data files.

2.5.5 Fault Protection (FP)—One of the fundamental guidelines in the design of the AutoNav system was to minimize the possible amount of trouble that the system could cause other systems or the spacecraft overall. AutoNav to a very large degree attempts to trap all of its possible errors internally and exit the faulty function in a manner that to the external system looks “normal.” As a result, there were no explicit connections to the FP system. It was additionally felt that none of the types of internal Nav failures mentioned above warranted notice by FP, even in a monitoring sense. Furthermore, the general use of the sequencing system for most commanding that involved actual spacecraft actions meant that AutoNav requests for action were covered by the usual FP provided by any sequence. There is one indirect method by which FP can detect an AutoNav failure. During certain fault recovery modes when ACS does not receive ephemeris data from AutoNav, it complains to FP, which will variously, depending upon circumstances, merely note the complaint or take the spacecraft to a higher level of fault state. As part of a safing event, FP will run scripts that set the AutoNav Modes into “stand-by” states wherein no attempts will be made to alter EEPROM files, including the Non-Grav History file.

3.0 TEST PROGRAM

3.1 Ground Test

The Ground Testing of AutoNav proceeded on several fronts and on several platforms. The original algorithms and code prototypes were built in a UNIX operating system using the MATLAB® environment. As a feasibility demonstration of the AutoNav concepts, an entire simulation of a flight to an asteroid was created; the prototype version of AutoNav was used to simulate and process pictures, perform OD, and compute course corrections on the way to an asteroid target. A number of the elements of the simulation were adopted from previous flight-support software, including the multiple-cross correlation algorithm used for the Galileo asteroid encounters (see Appendix C). Subsequent developments in image processing and in the orbit determination algorithms also continued to be done in MATLAB®, even after the initial code deliveries, to research and prove approaches. This was especially important as the encounter software was not deemed critical to launch and was, therefore, not completed at the time of the final software load in September 1998 for the late October 1998 liftoff.

3.1.1 UNIX-Based Simulation—As the C-code elements of the AutoNav software were produced, they were tested individually in stand-alone calls, and then assembled into three extended simulations of sub-sets of the AutoNav software. One simulation was specifically for the image processing elements of the flight-software and was comprised of drivers capable of independently testing all of the picture data handling routines of AutoNav, as well as simulating pictures for purposes of testing. Another simulation focused on the robustness and performance of the OD filtering. This simulation took a given set of observations (reduced pictures) with certain noise characteristics and estimated the spacecraft state under varying data conditions (e.g., frequency, quality, and outages). The results of this extensive set of simulations are detailed in Appendix D. The net result in cruise was a capability of achieving 200-km and better than 1-m/s OD accuracy. A third UNIX-based simulation was built to test efficacy and robustness of the maneuver computation algorithms for correcting the IPS mission burn profiles. A number of different strategies were tried; the operational parameters for using the updating algorithm were refined in this simulation. The results of this analysis are given in detail in Appendix E. The net result was the demonstrated ability of the retargeting algorithm to compensate for the expected error sources and, within the expected limiting bounds, keep the spacecraft course on target.

3.1.2 TestBed Testing—Several testbed platforms were available for testing AutoNav software. With the exception of timing, throughput, and overall CPU performance issues, the testbeds were *not* used to assess numerical performance of AutoNav. Once numerical stability and compatibility was established between the UNIX and testbed platforms, computational validity was assumed. Therefore, all testbed tests were used to check overall AutoNav software validity in the FSW environment, including the VxWorks operating system. The testcases were periodically re-checked against UNIX tests when numerical questions arose.

The simplest testbed was dubbed “Babybed,” several of which were available. These had a Power PC–based simulation of the RAD 6000-based operating system. An overall “build” of the entire FSW did not exist, but limited key elements were available, such as timing services and the underlying messaging system (IPC). Nav built background “stubs” for the subsystems that required external interaction, including ACS, MICAS, and IPS. With these, somewhat “stand-alone” testing of the AutoNav modules was possible. Necessarily, these test cases were limited to specific predetermined test cases: without the rest of the onboard software, no closed-loop interaction was possible with other elements. Limited throughput and performance tests could be accomplished

to assess the viability of algorithms under “clean” (i.e., not competing with other FSW elements) conditions.

The next higher fidelity of testbed was called “Papabed” and was comprised of a flight-engineering-model version of the DS1 Rad6K computer and 1553 bus. No flight hardware, spares, or engineering models were attached to Papabed. However, the entire FSW system existed onboard, and tests that invoked the interaction with other subsystems were performed. Also, flight-like commanding and telemetry was available, allowing the test of both uplink and downlink telemetry interactions. It was on Papabed that the first PhotoOps, TCMs, and mission burns were successfully accomplished in a realistic fashion, with AutoNav planning turns through APE and executing those turns with the ACS constraint monitor moderating. All of the AutoNav commands were tested by the Nav team on Papabed under a variety of conditions. For purposes of testing on the higher level testbeds, an AutoNav “self-sim” capability called FrankenKenny (FK) was created. FK is a dynamic simulation which, based on nominal or independently generated spacecraft ephemerides, creates pictures or “paints” images on existing pictures and makes those available to AutoNav. With this feature, it was possible to perform very realistic closed-loop tests of AutoNav functions.

The highest level of testbed fidelity are Hotbench and DS1-Testbed. These testbeds offer the greatest level of hardware integration, including engineering models of IPS and MICAS subsystems. During the final pre-launch software validation and verification, all functions of Nav were systematically tested and the results logged. With each update of the software, regression tests were performed to verify the integrity of the new version. Additionally, post-launch, operational tests of pending sequences on the spacecraft were run on the testbeds. Two months before the Braille encounter, a series of tests were done on the six hours of onboard autonomous operations that comprised the encounter. This test required configuring DS1-Testbed in as realistic a state as possible to the conditions (both physically and logistically) to those expected at Braille. When started, the “Testbed spacecraft” began the AutoNav operations and proceeded to guide itself, in its simulated universe, to the target. During these tests, it was discovered that the full closed-loop capability of the FK sim—including a dynamic modeling of an executed TCM—was not operating correctly (the FK integrated trajectory was, in fact, temporarily neglecting the TCM). Therefore, when this feature was invoked, small or pre-determined TCMs were used to attenuate the problem. For other tests, FK was configured to produce “perfect images” based on AutoNav’s current understanding of the spacecraft position. For all of these tests, when other anomalies were excluded, the performance of AutoNav was consistent with the expectations of the pre-launch analyses referenced above.

3.2 Flight Test

3.2.1 *Early AutoNav Flight Operations*—Figure 13 shows the overall mission plan of DS1. With a launch in late October 1998 and the need to validate onboard systems sufficiently to begin a major mission burn in November of that year, the intense nature of the early mission operations is clear. Following is a timeline of important navigation, navigation validation, and related DS1 events in the early mission.

- 10/24/98 12:08 UTC: DS1 Launch. As soon as the spacecraft computer boots, NavRT begins to successfully provide ephemeris data to ACS.
- 11/06/98: First Picture Taken with MICAS. This shows serious anomalous behavior, later identified as significant scattered light leakage into the instrument.
- 11/10/98: First attempt to light the IPS “main engine.” The engine runs for 4.5 minutes, autonomously shuts down, and does not restart.
- 11/18/98: First AutoNav Photo-Op session. DS1 enters “safe mode” due to ACS/Sun-sensor software error as AutoNav turns spacecraft X-axis more than 140° from the Sun.
- 11/24/98: IPS engine started at low throttle level, with spacecraft HGA (X-axis) on Earth.
- 11/30/98: IPS throttled up to nominal power for achieving mission objectives.
- 12/03/98: 200 hours of IPS thrusting achieved.
- 12/04/98: Spacecraft turned to nominal thrust-vector direction, optimum for achieving mission objectives.
- 12/12/98: Start of IPS burn, spacecraft safes due to battery state-of-charge fault.
- 12/18/98: First operation of AutoNav mission burn, AutoNav turns spacecraft to desired attitude, and starts engine. Thrust vector updated throughout week.
- 12/21/98: Second Photo-Op attempt. All Photo-Op operations worked logistically, but none of the pictures processed due to MICAS scattered light.
- 12/22/98: Second mission burn started. AutoNav operates IPS on the designed mission trajectory over the 1998 holiday season.
- 01/06/99: Nav file load. Parameters in the image-processing software altered in attempt to work around scattered-light problems.
- 01/07/99: Third Photo-Op. No pictures successfully processed.
- 01/07/99: Nav Team begins major overhaul of image-processing algorithms in effort to cope with severe scattered-light infiltration into MICAS.
- 01/18, 01/20, 01/26, 02/01/99 Photo-Ops: Only the very brightest asteroids and stars (brighter than 8.5M) are processable *on the ground*, with the M3 (launch) AutoNav software and extensive parameter manipulation, so heavily damaged are the pictures by scattered light. Downlinked pictures are used to define and test alternative image-processing software.

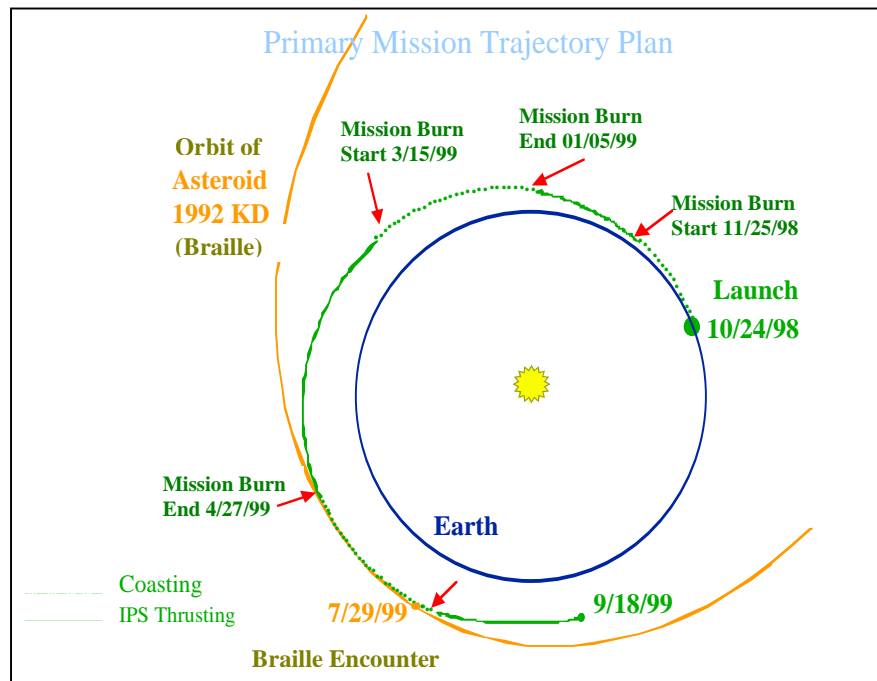


Figure 13. Primary Mission Trajectory Plan

- 02/08/99: M4 Software update onboard, including substantially upgraded AutoNav image-processing software.
- 02/18/99 First PhotoOp on M4 software: Only one picture of 30 processes successfully due to erroneous uplinked parameter-value settings.
- 02/19/99: Nav File Load of new parameters and data-files, including ground-processed picture data in OD file. 36 data from PhotoOps from Jan 7, 20, 26, and Feb 1 are given to AutoNav to “seed” the 2/22 PhotoOp and OD run.
- 02/22/99 PhotoOp /OD/ManPlan run: Of 32 pictures on four lines of sight, six succeeded, three each on two lines. These five added to 36 uplinked data produced the first viable onboard autonomous OD, which is in error from the ground-determined state by about 4000 km and 2 m/s. This solution is intentionally not saved onboard. The ManPlan operation (correctly) declines to perform any computations, as there is no TCM or mission burn pending in the near future (as per plan).
- 02/27/99: Update on AutoNav Control Modes to preserve the OD results (by replacing the onboard ephemeris), effectively putting the spacecraft under AutoNav control after the next OD operation.
- 03/01/99 PhotoOp/OD/Manplan: 13 of 30 pictures taken successfully processed, OD arc spans Jan 5 to Mar 1. OD results are within 5000 km and 2 m/s of

radio-nav determined spacecraft position. This solution is saved onboard in the form of a 60-day spacecraft ephemeris. ManPlan again (correctly) declines performing any maneuver planning.

3.2.2 The First Validation of Onboard Orbit Determination—With DS1 now autonomously computing its course, March activities began a period of 10 weeks of “normal” operations, which included weekly PhotoOp/OD/ManPlan sequences and periods of mission burns. This period of regular data and fairly high-rate downlink capability offered a good opportunity to further analyze and debug AutoNav operations. One of the first items investigated was the geometric stability of the camera. With the initial forays into onboard processing, it was immediately clear that the optical data residuals were larger than expected. Figure 14 shows pre- and post-fit residuals for a solution performed onboard in this investigation period. RMS residuals larger than one pixel, with biases (in some cases) of several pixels, were much higher than expected. Calibration of the camera pre-launch indicated that measurements good to about one pixel should be obtainable without re-calibration. Furthermore, AutoNav’s ability to acquire and locate the dim (on the order of magnitude 10 to 11) asteroids expected (and required) seemed badly disabled; in fact, inconsistent measurements of stellar photometry lead to speculation of strong non-linearity in the CCD channel at low-flux levels. Necessarily, a thorough




Figure 14. Pre-(upper) and Post-(lower) Fit Residuals from 3/22/99 Optical Solution

calibration of MICAS was called for; this was scheduled for March 5. Two star clusters were chosen: one with a dense distribution of moderate to dim stars, another with a few bright stars to aid in both geometric and photometric calibration. Additionally, the MICAS team scheduled a set of calibration frames on March 11.

3.2.3 Results from MICAS Calibration Images—The MICAS and Nav Teams undertook an extensive calibration campaign in early March to attempt to characterize the scattered-light and light-leakage problems. The spacecraft imaged a pair of star clusters for purposes of calibrating the geometric “flatness” of the camera field; these pictures revealed that there were severe distortions, up to 5 pixels in size and of unusual character. Pre-launch calibrations had indicated less than 1 pixel of relatively benign (i.e., readily calibratable) distortion in the field. With the images taken to characterize the scattered light, a quantitative analysis was made of the resulting increased noise in images, which was substantial and damaging to the navigation algorithms.

In order to cope with the geometric distortions, work began on a new distortion model for the flight software, incorporating a sixth-order Legendre Polynomial model. To cope with the high levels of scattered light, algorithms for taking and differencing a background picture are devised, and implementation begun. As part of the calibration suite, Mars pictures indicated that the approach target (1992KD) would be very bright. From these frames, there was observed a nonlinearity in the CCD response, which attenuated weak signals. This nonlinearity had been suspected from the earlier AutoNav frames. The result of this analysis indicated that only the brightest asteroids and stars would be processable by AutoNav. This fact required a change in strategy for picture planning. The original plan was to look at any time at a particular “good” asteroid and, with the expected performance of the camera, acquire in general two to four magnitude 10 stars—more than sufficient for a navigation frame. However, now the suite of “good” asteroids was diminished by 75% and the useable stars were those of magnitude 9 or brighter. Consequently, far fewer asteroid or stellar targets were now available and the picture-planning file had to be carefully “primed” to allow AutoNav an opportunity to image these.

3.2.4 Late Cruise Timeline—The following timeline outlines AutoNav operation and validation activities from 3/1/99 to 6/1/99, the beginning of intensive encounter preparations. This period of time encompasses additional proving of the onboard OD (which continues to be fully engaged onboard) and the first closed-loop operation of the mission burn Maneuver Planner (ManPlan). Analysis

of the picture processing continues and plans are made for further enhancements to the image processing algorithms.

- 3/8/99 PhotoOp: Six 4-lines-of-sight (LOS) pictures. Only the bright asteroid Vesta successfully processes, with five of six Vesta pictures entering the solution. OD error, relative to ground track, climbs to over 6000 km.
- 3/15/99 PhotoOp: 2 lines-of-sight, 12 pictures each. All pictures process normally. OD dispersions grow to near 10,000 km. In this time frame, it is realized that the RCS non-gravitational modeling onboard is severely compromised due to large drops in hydrazine pressure since launch. This factor of 2 drop would result in an approximately equal drop in specific impulse of the attitude thrusters and, thus, in the modeled values of accumulated ΔV sent to AutoNav. Nevertheless, use or non-use of this part of the model makes no appreciable change in the OD performance.
- 3/16/99 Mission Burn: The second of the mission burns to 1992KD begins with Nav mediated thrusting.
- 3/22/99 PhotoOp/OD: 27 of 36 pictures process normally; OD quality still marginal (but adequate for cruise operations). Mission burns continue.
- 3/29/99 PhotoOp/OD: 22 of 36 pictures process normally; however, despite a good distribution of asteroid geometries, the OD quality continues to deteriorate, to 13,000 km. However, the velocity measurements are good to about 1.5 m/s. This quality of velocity determination was inconsistent with the poor position determination, indicating that systematic biases were being observed in the astrometry. It was determined at this time that the largest share of this bias was due to an inconsistency in a model describing the a-priori pointing biases of the camera. These parameters were changed onboard in a subsequent file load.
- 3/29/99 ManPlan: First onboard execution of ManPlan in the presence of a control opportunity. ManPlan correctly assesses that the current OD uncertainties (the OD filter formal errors) mapped to 1992KD encounter are too large to warrant a thrust-plan change. Thrusting on the nominal plan continues.
- 4/05/99 PhotoOp/OD/ManPlan: 29 of 32 pictures process normally; however, due to a dearth of bright asteroids available, the geometry is no longer strong, weakening the OD performance. Nevertheless, with the correction of the pointing a-priori model (see 3/29), the OD performance begins to trend strongly toward improvement (see Figure 15). A file load is accomplished on this day to change parameters such that the mission burn profile will be updated regardless of the formal uncertainties of the OD solution when ManPlan is run on 4/12.
- 4/12/99 PhotoOp/OD/ManPlan: 31 of 36 pictures process normally. OD solution quality is about 6000-km position and a consistent 4-m/s velocity. The ManPlan updates to the thrust profile are considered adequate to use and left in place for the beginning of the mission burn.




Figure 15. Flight vs. Ground-Orbit Determination April 5, 1999

- 4/19/99 PhotoOp/OD/ManPlan: 32 of 36 pictures process normally. OD solution quality improves to about 4000-km position and holds at about 4 m/s. The ManPlan run for the associated mission burn was unsuccessful, due to the combination of relatively poor OD quality, the shortness of remaining burn arc, and the fact that ManPlan was forced to compute statistically insignificant changes. As a result, the nominal plan was reverted to onboard.
- 4/26/99 PhotoOP/OD: 13 of 16 asteroid images process normally, OD quality improves to 2000 km, as the amount of corrupted data from the pointing angle a-priori is systematically trimmed from the OD file. Velocity errors rise slightly to 4.7 m/s. No ManPlan is attempted.
- 5/1–5/5/99 M5 Upload and Reboot. M5 FSW is loaded to enable the inflight RAX test; M5 is identical to M4 for AutoNav.
- 5/6/99 PhotoOp/OD: 27 of 32 pictures process normally, OD quality maintains at about 2000 km and 4.7 m/s. Substantial improvements are seen with ground processing using Legendre polynomial corrections to the asteroid observations and using pre-processed pictures. The pre-processing entails taking a “background” picture with each LOS and differencing this picture from all

pictures on this LOS. The background picture is offset slightly (e.g., 200 pixels) from the Nav pictures to prevent damage to the Nav targets. These two algorithmic changes are factored into the M6 FSW load now building.

- 5/10–5/23/99 RAX Experiment: No Nav operations occur in this timeframe.
- 5/24, 26, 29, and 31/99 PhotoOp/OD Operations: Image processing is more than 75% successful overall. With tuned image-processing parameters (more discrimination of image strength), the use of only strong asteroids and stars, good geometry of asteroids, and a dense late data set (and despite nearly a month hiatus in Nav data acquisition due to RAX preparations and testing), OD improved to 1700 km and 2 m/s (see Figure 16).

3.2.5 Final Software Load and Final Validation of Cruise AutoNav—From 6/1 to 6/9/99, the M6 software set was uploaded to the spacecraft. This included final adaptations to the MICAS problems for cruise, including the Legendre polynomial model of geometric distortions and picture differencing to further reduce problems associated with scattered light. Over the next two months, these new elements were validated in cruise AutoNav operations. AutoNav and ACS software for the execution of TCMs

would be exercised for the first time. Additionally, the first flight use of the now complete encounter software was made during a rehearsal less than two weeks before closest approach. Following is a summary of AutoNav validation and related events down to two days before closest approach.

- 6/1–6/10/99 M6 Software: Loaded and booted on DS1.
- 6/10/99 PhotoOp/OD/ManPlan: The first PhotoOp performed with the M6 software was unsuccessful, due to the presence of an un-updated parameter file, which caused the image processing to work in “M3” fashion. Nevertheless, the ManPlan operated correctly and successfully planned an IPS TCM scheduled for 6/14. The decision criterion used was that it was necessary for AutoNav to reduce the distance remaining to the target at least by half in order to not be overwritten. In this case, the criteria was satisfied. This was computed to be a 1.5 m/s IPS TCM, vectorized along two legs, to correct 830 km in the 1992KD B-plane, and 58 s time-of-flight (or 870 km).
- 6/14/99 First IPS TCM: AutoNav executes the IPS TCM. No problems are encountered.
- 6/16–6/20/99 Photo-Op/OD: 19 of 36 and 20 of 36 pictures process normally, although one of the 4-LOS was at an attitude near the asteroid approach attitude. Because of scattered light effects, none of those pictures were processable though they were very useful for calibration and characterization purposes. The OD quality of these solutions degraded alarmingly to about 3500 and 2130 km and 1.7 and 0.9 m/s, respectively.
- 6/20/99 Anomaly Resolution: It was discovered that *ground* processing of the new Legendre polynomial distortion model had been in error. Consequently, uploaded calibrated data older than 6/10 was erroneous. A new OD file was prepared for uplink, with corrected calibrations, and would be used for OD onboard subsequent to the 6/23 OD (for which the uplink would not be in time).
- 6/23/99 PhotoOp/OD: Only two asteroids were available; 16 pictures were taken of each, with 14 and 11 processed successfully. Still affected by the bad calibrations, the OD was still degraded to 1000 km and 0.5 m/s; however, the effect was diluted by the preponderance of late and correctly calibrated data. The file load was completed after this time.




Figure 16. Flight vs. Ground-Orbit Determination 5/31/1999

- 6/29/99 PhotoOp/OD: Of two available asteroids, only one processed successfully, with 12 of 16 pictures. However, with the calibrations corrected onboard, the OD performance improved dramatically to 662 km and 0.58 m/s.
- 7/2/99 PhotoOp/OD/ManPlan: 28 of 36 pictures processed successfully, OD quality was 904 km and 0.3 m/s.
- 7/4 and 7/6/99 PhotoOp/OD/ManPlan: 22 of 36 and 27 of 32 pictures processed normally, the OD quality was 928 and 1022 km and 0.39 and 0.31 m/s, respectively. The 7/6 ManPlan was used to plan onboard the ACA – 20 day IPS TCM. Figure 9 shows a vast assortment of OD solutions from AutoNav onboard, AutoNav mirrored operations on the ground and from Radio Nav. Within this complex, it can be discerned that the AutoNav solution of 7/6 created a TCM solution (when measured against the radio solution) that would not meet the acceptance criteria for an autonomous TCM (namely, reducing the B-Plane error by 1/2). This would have been the case had a small change in non-grav modeling procedure not changed for the previous maneuver file upload (namely, the lack of forecasting of ΔV s

associated with PhotoOps). This change caused a 400-km discrepancy in the solution (well within the formal uncertainties, as shown), enough to violate the criterion. Since several upcoming TCM opportunities existed, it was decided to cancel the ACA – 20 day TCM.

3.2.6 Asteroid Rehearsal Preparations—Preparations for encounter and for the encounter rehearsal began early in 1999, but focused on the last 90 minutes of operations only until March, when the activities of the last 6 hours before closest approach were planned. By early July, the details of the last two days had been planned. Table 4 summarizes the Nav and related activities and durations of the last two days.

The encounter rehearsal, originally scheduled for 6/25, involved an extensive series of practice runs on Testbed and set-up activity on the spacecraft. In order to accomplish these, rehearsal files had to be created, including spacecraft ephemeris, simulated body ephemeris, a target star catalog, and tailored parameter files. These data create a “simulated universe” in which the spacecraft finds itself upon initialization of the rehearsal. Within this universe, the spacecraft “sees,” through FK modified images, the

Table 4. Navigation Encounter Activities

Encounter Relative Event Time	Duration	Activity	Sequence No.
-2 days 3 hr	180 min	RCS TCM (“Minus 2 Day”)	AN300
-2 days 0 hr	210 min	PhotoOp/OD/ManPlan	AN301
-1 day 21 hr	240 min	High Gain on Earth Telecom Track	
-1 day 17 hr	210 min	PhotoOp/OD/ManPlan	AN301
-1 day 14 hr	240 min	High Gain on Earth Telecom Track	
-1 day 10 hr	210 min	PhotoOp/OD/ManPlan (OD and Maneuver Planning for -1d TCM)	AN301
-1 day 3 hr	180 min	RCS TCM (“Minus 1 Day”)	AN302
-1 day 0 hr	90 min	PhotoOp/OD/ManPlan (OD and Maneuver Planning for -18hr TCM)	AN303
-23.0 hr	210 min	High Gain on Earth Telecom Track	
-19.5 hr	90 min	RCS TCM (“Minus -18hr Hour”)	AN304
-18.0 hr	90 min	PhotoOp/OD/ManPlan (OD and Maneuver Planning for -12hr TCM)	AN303
-17.0 hr	210 min	High Gain on Earth Telecom Track	
-13.5 hr	90 min	RCS TCM (“Minus -12hr Hour”)	AN305
-12 hr	90 min	PhotoOp/OD/ManPlan (OD and Maneuver Planning for -6hr TCM)	AN303
-11 hr	270 min	High Gain on Earth Telecom Track (Last Ground Intervention Opportunity)	
-6.5 hr	90 min	RCS TCM (“Minus -6hr Hour”)	AN306
-5.0 hr	75 min	PhotoOp/OD/RSEN Init	AN307
-5.0 hr	Continung	Low Gain Track, S/C on Target	
-3.5 hr	90 min	RCS TCM (“Minus -3hr Hour”)	AN308
-2.0 hr	30 min	PhotoOp/OD (10m P.O., 20m OD)	AN309
-1 hr 30 min	90 min	Encounter Sequence	SEQ50
-1 hr 30 min	10 min	PhotoOp	Do.
-1 hr 15 min	10 min	PhotoOp	Do.
-55 min	25 min	OD	Do.
-27 min	27 min	RSEN	Do.
-5 min	2.5 min	1 st Close Approach Sequence	SEQ51
-2.5 min	1.5 min	2 nd Close Approach Sequence	SEQ52
-90 s	65 s	3 rd Close Approach Sequence	SEQ53
-25 s	25 s	4 th Close Approach Sequence	SEQ54

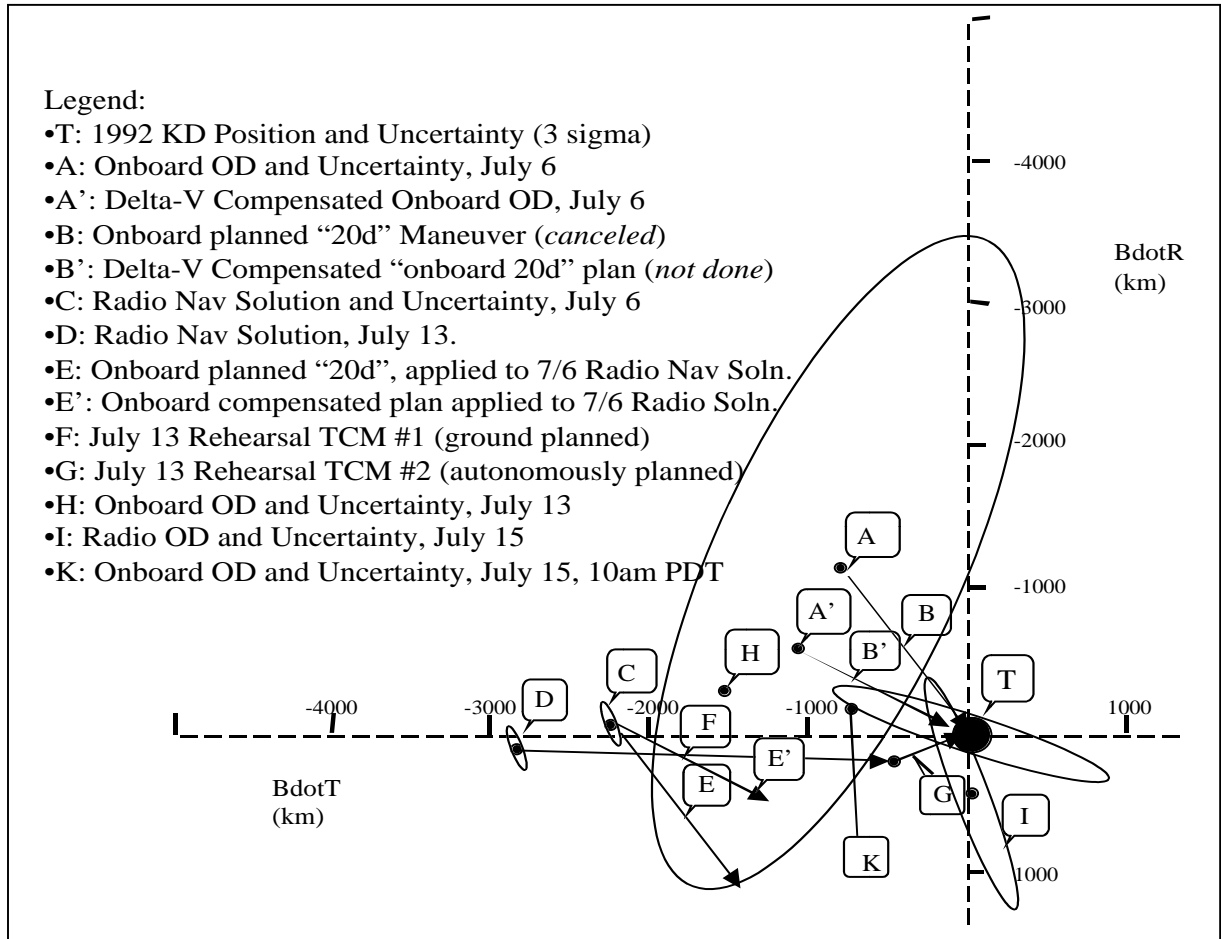


Figure 17. Current B-Plane Target Conditions at the – 20 – 10 Day TCMs: Decision Data from 7/15/99

phantom approach target (dubbed "Spoof") and computes its position relative to Spoof, adjusting course correspondingly. It was desired (and necessary) to use the rehearsal as the first execution of an RCS TCM. It was further desired to use this correction purposefully; in other words, to use the approach TCM to Spoof to correct the actual approach asymptote to 1992KD. The rehearsal maneuver file was tailored to make the first of the rehearsal TCMs that was, for the rehearsal only, deterministic. This TCM was a ground-designed event that would remove much of the then existing residual in the B-plane. At the same time, sufficient residual needed to be left for the second of the two rehearsal TCMs to be a substantive test, and not endanger the 1992KD encounter if it misfired in any way (see Figure 17). The files for the rehearsal were uploaded to the spacecraft on 6/23, while ground tests in the Testbed continued. The results from these tests were good from an AutoNav standpoint, with Nav tracking the target to within 30 seconds of closest approach. However, there was substantial uncertainty about other subsystems; therefore, the onboard rehearsal on 6/25 was cancelled and rescheduled for 7/13. Aside from the requirement that all of the encounter rehearsal-specific files

be regenerated, any opportunity to update the flight software if problems during the rehearsal were encountered was lost.

3.2.7 Results from the 7/13/99 Encounter Rehearsal—The rehearsal was overall very successful. All Nav operations succeeded:

- Execution of Rehearsal RCS TCM-1, a 2400-km B-Plane deflection, or 1.7 m/s, was normal, with performance (determined afterward from radio data) to be within 1.5%.
- FK simulation of images, PhotoOp operations, including image processing, OD, and maneuver planning for RCS TCM-2 occurred normally.
- Execution of Rehearsal RCS TCM-2, a 500-km 0.3-m/s burn, was normal.
- Entry into RSEN mode was normal. RSEN improves position knowledge to better than 0.5 km in the field, and 5 s downtrack.
- Late image processing allowed RSEN to track Spoof to within 30secs of encounter; the approach late-encounter sequences were initiated within their expected uncertainties.

3.2.8 *Cruise to – 5 Day TCM*—A PhotoOp immediately after the rehearsal was cancelled, due to uncertainty in the state of the spacecraft and the near exhaustion of the flight team. There were, however, five more PhotoOps leading up to the ACA – 5 day TCM, with the final one of these designing the TCM itself. Following is that timeline:

- 7/16/99 PhotoOp/OD: 28 of 36 pictures successfully processed along two lines of sight. Accuracy is 658 km and 0.34 m/s.
- 7/18/99 PhotoOp/OD: 28 of 36 pictures of two asteroids, plus 13 of 16 pictures of Mars, were incorporated into the solution. Mars invoked a heretofore unused mode of processing images, wherein extended bodies (Mars being about 5 pixels across) are “brightness-centroided” and then that position is corrected for phase. In OD, these pictures were highly de-weighted (5 pixels, as opposed to 2 for asteroids). As a result, the solution quality onboard remained relatively stable, at 669 km and 0.32 m/s. Post processing on the ground revealed that even with stronger weighting, Mars did not substantially improve the match between the ground radio solutions and flight. This left a concern of the reason for the outstanding observed biases of several hundred kilometers. It was (and is currently) believed that these biases are due to a combination of residual geometric calibration defects and possibly ephemeris errors. Pre-launch, it was expected that the geometric calibration could be made to 0.1 pixel; however, the insensitivity of the camera (inability to acquire dim stars) precluded this. The ephemeris errors, expected to be in the neighborhood of 100 to 200 km were running somewhat larger, perhaps 400-km as would be observed at Braille (1992KD).
- 7/19/99 PhotoOp/OD; Mars-only PhotoOp: 11 of 16 Mars images successfully processed, with the following Radio/Flight agreement: 572 km and 0.25m/s. This Mars observation (as with 7/20) offered unique viewing of Mars against a very bright star. Nevertheless, the substantial challenge in processing the Mars images prevented pushing the quality of the OD past the limiting effects discussed above.
- 7/19/99: The final best-ground-determined Braille ephemeris is loaded onboard the spacecraft, representing the observing efforts of about a dozen astronomers over 18 months, and incorporating observations less than two weeks old. It is believed that this ephemeris is good to about 150 km (1 sigma).
- 7/20/99 PhotoOp/OD: Mars-only PhotoOp; 13 of 16 Mars images successfully processed, with the following Radio/Flight agreement: 710 km and 0.22 m/s.
- 7/21/99 PhotoOp/OD: 12 of 16 Mars images and 20 of 24 asteroid images successfully processed, with the following Radio/Flight agreement: 776 km and 0.11 m/s. Interestingly (and serendipitously), the Braille B-Plane

Radio/Flight agreement was nearly perfect (see Figure 18).

- 7/21/99 Ground Seed Onboard: In order to help compensate for camera deficiencies (believed largely associated with the geometric calibration), an OD file with spacecraft-acquired optical data was put onboard on this day. This data had been “scrubbed” to remove observations that were only marginally good. With the limited data set available to the ground planners it was impossible to set low-pass residual thresholds to a discriminating enough level to accomplish this editing onboard. These scrubbed data sets were regularly achieving Radio/Flight OD agreements of better than 300 km and 0.25 m/s (see Figure 19). Also, in preparation for the ACA – 5 day TCM, a maneuver file was placed onboard with a TCM design based on the radio data (see Figure 18). If after the 7/22 PhotoOp, it was decided that the onboard-planned TCM design was inadequate (recall the decision criteria was to reduce the net deflection from target by one-half); the radio-data-based file would be made the primary maneuver file.
- 7/22/99 PhotoOP/OD/ManPlan: A similar sequence of pictures was scheduled for 7/22 as was scheduled for 7/21. However, a problem occurred (the source of which has not been identified) that caused one or more of the Mars pictures to be off-pointed. This in turn tripped a latent AutoNav software bug, which caused the erroneous writing of large blocks of data into the OPNAV file. This effectively filled the fsw/files file system. The OPNAV file was unreadable by AutoNav; consequently the OD function failed, reverting to the unaltered OD file, which was the “seeded” file uploaded on 7/21. This solution was within 250 km of the radio solution “at epoch” (e.g., on 7/21) and mapped to a maneuver of 400 km in the Braille B-Plane (see Figure 18). This solution did meet the acceptance criteria for the onboard TCM design, but only barely. Because there was an associated anomaly with the PhotoOp and OD, it was decided to revert to the ground design. This was accomplished with a simple Nav_Data_Update command to point AutoNav to the already onboard file. This anomaly had the beneficial effect of alerting the AutoNav team to this bug, which posed a threat to the close-approach sequences. The Picplan file was changed at the next opportunity to ensure that extended-image picture processing would not be used in any of the subsequent PhotoOps, as was then planned for those within 5 hours. With this picture-taking mode disabled, it was believed that AutoNav would receive insufficient improvement in position from the early approach pictures to warrant the ACA – 3 hour TCM. Consequently, the sequence for this TCM was altered and the Nav_Do_TCM call was replaced with a simple turn to Braille.




Figure 18. Minus 5 Day TCM Solutions




Figure 19. Flight OD vs. Ground OD#37, 7/21/99

- 07/23/99 14:30 to 07/24/99 04:00 UTC: ACA – 5 day IPS TCM. This TCM executed normally. Figure 18 shows the effect of the TCM: approximately 500 km in the “B dot R” direction.

3.2.9 Acquisition of Target and Countdown to Encounter—Perhaps the most challenging aspect of the encounter to AutoNav was the lateness of expected acquisition of the target in the images. Had the approach exposures not been limited to 5 s or less due to the scattered light and light leakage into and within MICAS, Braille would likely have been imaged in time for the ACA – 5 day TCM and, possibly, the ACA – 10 day (50- to 100-s exposures would have been taken). As it was, the target was not detected until ACA – 3 days, and then only with extreme post-processing on the ground. The AutoNav system only detected a strong enough signal to “lock on” at ACA – 17 hours, again due to the dual limitation of short exposures and scattered light. Following is a timeline of the Nav activities following the ACA – 5 days TCM:

- 7/24/99 PhotoOp/OD: Following the TCM, there was one “conventional” PhotoOp that took pictures of “beacon asteroids” plus the first attempts to image Braille. Of the former, 14 of 24 were successful, but Braille was not seen. The quality of this OD was 811 km and 0.59 m/s.
- 7/25/99 PhotoOp/OD: Only images of Braille were taken, which were not seen. There were, thus, no changes in the OD quality, since there were no data.
- 7/26/99 05:00 UTC PhotoOp/OD: Onboard, AutoNav makes no detection of Braille; however, with intensive image-processing on the ground, including picture addition, an extremely faint “phantom” appeared, approximately 350 km from the nominal expected position of Braille. This represented about a 2-sigma error from the recently delivered Braille ephemeris.
- 7/27/99 00:30 UTC ACA – 2 day TCM: In view of this somewhat large apparent ephemeris change, based on suspect data and the fact that the radio solution was indicating that the ACA – 5 day TCM had performed nominally, it was decided to cancel the ACA – 2 day TCM. In other words, aside from the apparent ephemeris error, which was not nearly well enough determined by the “phantom” to act upon, there was no reason to implement the maneuver.
- 7/27/99 03:00 UTC PhotoOp: AutoNav does not detect Braille, but three raw pictures are downlinked.
- 7/27/99 10:00 UTC PhotoOp: AutoNav does not detect Braille, but six pictures are downlinked. With ground analysis of these images, three reliable but very dim images are acquired. The observed position of Braille is consistent with the earlier “phantom.” From these, a design is constructed for the ACA – 1 day TCM. Using the AutoNav software on the ground as would have been onboard if a higher signal had been available from MICAS, a maneuver file is created that includes the TCM. This file is uplinked (see Figure 20).
- 7/27/99 18:30–21:00 UTC ACA – 1 day TCM: Normal execution.
- 7/28/99 00:00–03:00 UTC PhotoOp: 18 pictures of Braille are scheduled and taken. Braille is not yet bright enough for AutoNav to “lock on,” but ground processing extracts another two detections of the downlinked images. These indicate that the spacecraft is sufficiently on target to warrant cancellation of the ACA – 18 hour TCM.
- 7/28/99 10:10–11:30 UTC ACA – 18 hr TCM: Window cancelled.
- 7/28/99 11:33–12:33 UTC PhotoOp: 18 pictures of Braille are scheduled and taken. An unknown number of these images “lock on.” From the three images that were subsequently downlinked, it seems reasonable to assume that many or most of these pictures were successfully processed. After image processing, AutoNav attempted to store the processed images into the OD file. A previously unknown software fault in AutoNav caused the vector of stored planning cycles to be exceeded by 1. This caused a memory write out-of-bounds and a subsequent reboot. Three pictures had, however, been scheduled for downlink.
- 7/28/99 12:33–16:00 Spacecraft Recovery. A series of activities that had normally taken one or two days was accomplished in little more than three hours.
- 7/28/99 16:00–22:25 Data Downlink and Preparation for ACA – 6 hour TCM: With the three pictures received, the AutoNav team completed the operation interrupted onboard, but with much less data. The optical data indicated that the ACA – 1 day TCM had successfully placed the spacecraft within 25 km of Braille, but not on the desired “umbra side.” A maneuver was designed to place the spacecraft on a 15-km impact-parameter trajectory. However, the solution was chosen from the distribution of solutions such that the target point would be biased “to the outside.” In other words, with the 1-sigma variance of solutions at 10 km, it was decided that an extra margin of safety was warranted. This maneuver file was created and uplinked shortly before the spacecraft turned away from Earth for the ACA – 6 hour TCM (see Figure 21)
- 7/28/99 22:25 UTC ACA – 6 hour TCM: This TCM executes nominally.
- 7/29/99 00:00–04:15 UTC (ACA – 30 minutes), three PhotoOps, two ODs: AutoNav takes and processes data normally keeping Braille in field of view (FOV). No Science frames are taken or preserved.
- ACA – 27 minutes RSEN Activated: AutoNav switches to APS sensor. No signal from Braille comes above the AutoNav APS threshold.
- ACA – 20 minutes: An unknown signal (probably a cosmic ray) spoofs AutoNav into a one-quarter APS




Figure 20. Pre-Minus -1 Day TCM, "Flight OD" Braille B-Plane

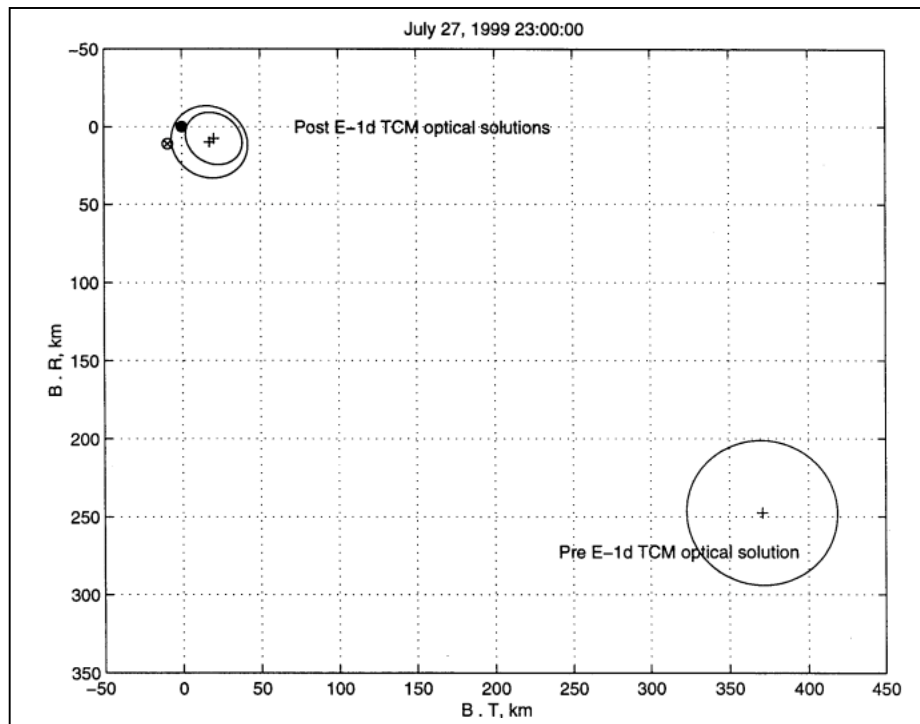


Figure 21. Pre-ACA - 6 hour TCM B-Plane July 27

FOV correction. Braille remains in the APS and CCD fields, but no frames are preserved.

- Down to ACA – 3 minutes: Braille is in APS and CCD fields, but no science frames taken or preserved. Nav activates the first encounter sequence, based on a-priori data. Sequences are scheduled for ACA – 300-, 150-, 90-, and 25-s initiations.
- ACA – 150 s: First CCD science frame taken. Braille is barely out of MICAS CCD FOV due to picture editing, and is outside of all subsequent picture APS and CCD fields.
- Inside 20 s: Braille is imaged in the IR FOV.
- ACA – 10 s: The sequence stops taking Braille pictures inbound.
- ACA + 15 minutes: DS1 is back on the nominal (e.g., pre-flyby ephemeris) Braille track. First successfully taken and returned close-up images of Braille occur here. APS images show an extraordinarily dim image, 10 DN, with 1000 DN expected. CCD images show 400 DN, one-tenth “fullwell,” with expected 1/2 to 1 expected.
- Post-Encounter reconstruction indicates approach Braille images 1 to 2 magnitudes dimmer than outbound, perhaps due to presented geometry of the irregular figure of Braille. Outbound images are also very dim, by factors of 5 to 10 from expectation.

From the above timeline it is apparent that the close-approach events did not proceed according to plan. In review, there was insufficient signal in the APS detector to allow AutoNav to detect Braille. Figure 22 shows diagrammatically the expected and received Braille signal on approach. Because no signal from Braille came above the minimum threshold, RSEN never “locked-on.” One of the principal causes of the lack of detection was the previously poorly characterized non-linearity of the APS detector. This non-linearity in the camera response, is shown in Figure 23. Additionally, a noise-spike, presumed to be a cosmic ray, did penetrate the threshold; AutoNav temporarily locked on to this, causing a deflection in the trajectory. Figure 24

shows the effect of this deflection on the position of Braille in the two visual fields-of-view versus the nominal trajectory that would have been followed if there had not been the cosmic ray event.

3.2.10 Post-Encounter Reconstruction and Performance Analysis—Despite the fact that the performance of the system during the Braille flyby was thwarted, it is nevertheless the case that operability and accuracy of the AutoNav close-approach system had been demonstrated in the testbeds and, more importantly, in-flight during the rehearsal. This was demonstrated using the few acquired images of Braille post-encounter. When these were provided to RSEN, accurate solutions of the spacecraft position were obtained with just one CCD image, leading to the unavoidable conclusion that had this detector been used, the encounter would likely have been very successful. Figure 24 shows the B-plane results of this analysis.

3.2.11 Causes of the Braille-Encounter Failure—There are five principal reasons that the expected high-resolution images of Braille weren’t obtained:

- Problems with the MICAS instrument lead Nav (and the Project) to believe that the CCD was unusable at encounter, requiring Nav’s use of the much less capable and much less understood APS sensor. In the event, the CCD would have been very useable through most (and perhaps all) of the encounter.
- Compounding the first problem, the Science and Nav teams overestimated by a wide margin the expected flux of Braille. Exposures set on the basis of these computations were hopelessly low for Nav and Science. In fact, it is likely that even if RSEN had worked exactly as expected, and kept the target in lock, the scheduled APS images would have had a uselessly low signal on approach due to APS non-linearity. Figure 25 shows a close-up of one of the outbound APS images (0.6 sec exposure) that captured Braille. The smeary figure slightly up and to the left of center is only 10 DN above background, or roughly 1/400 full scale (the white spot

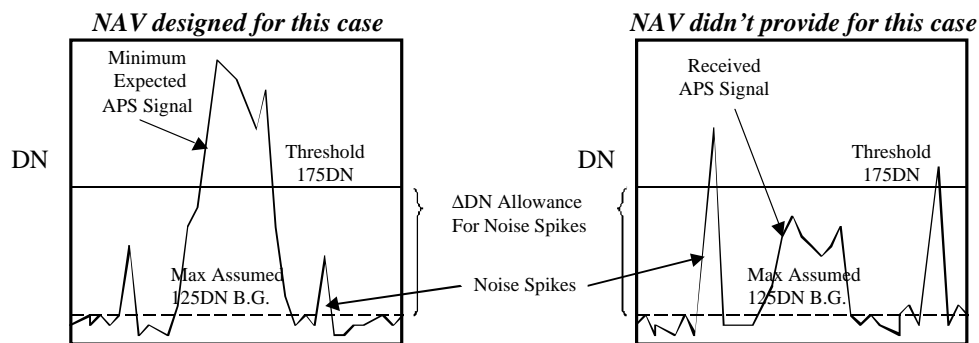


Figure 22. Diagrammatic View of Received RSEN Signal




Figure 23. MICAS APS Channel Non-Linear Signal Response




Figure 24. Encounter Results Using Post-Encounter CCD Braille Pictures




Figure 25. Post-Encounter APS Image of Braille

to the right is a noise spike). Given that the inbound flux from Braille was much lower and that the exposures were similar to this image, and given the non-linear effects of the APS response, the chances of any of the inbound APS frames being successful (even if properly targeted) seem remote. The CCD images, as mentioned above, predicted to be near saturation, were at no greater than one-tenth full-scale outbound, when the target

presented a much higher flux than inbound. A principle contributor to the over-estimation of inbound flux was the failure to realize that the body could present up to a factor of 60 reduction in flux if oblong, highly textured, and presenting itself in an unfavorable geometry—all of which apparently happened.

- The AutoNav RSEN algorithm was simplistic in that it could not distinguish a single-event noise spike (which the system did receive) from a continuously repeatable real signal (which the system did not receive). However, as shown in Figure 26, because of the limited sequence of science frames taken and preserved (discussed below), even if RSEN had not falsely locked, the approach-data return would not likely have improved.
- There was extremely limited space onboard for stored images, but far less than was actually available in terms of RAM. Most of the RAM was dedicated to “packet-space” that was unavailable due to the computational overhead required to turn a picture into packet data. Those few pictures that were taken and preserved were all late in the encounter, during a time when, without orbit updates from RSEN, there was very low probability of successful acquisition. Re-allocation of RAM space might have been possible, but was not undertaken. Taking and preserving earlier, more reliable, but less resolved images was also not undertaken.
- AutoNav code faults caused the spacecraft to safe at encounter – 17 hours. Though the spacecraft was recovered from safe mode in time to re-enter normal encounter operations at encounter – 6 hours, the

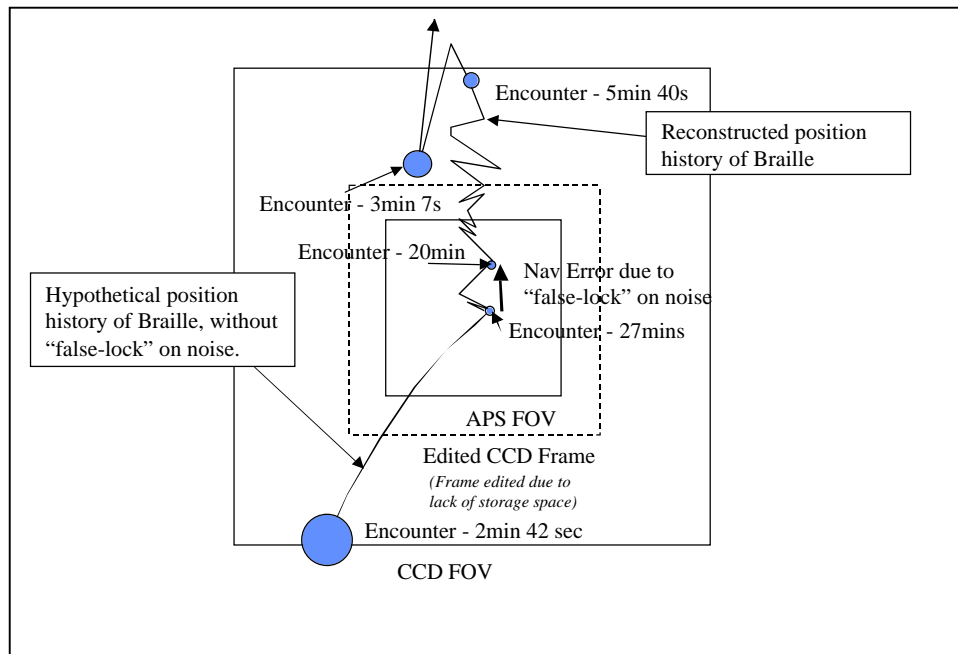


Figure 26. Reconstructed Nominal vs. Perturbed Braille Field-of-View Flight Path

tremendously difficult and intense recovery operations prevented additional data downlink of approach pictures and careful analysis of the apparent low light levels of Braille. However, even if this had happened, it would have been extremely difficult, and probably inadvisable, to alter the entire encounter sequence to lengthen exposure times; in many cases, it would have been impossible. Further, with the knowledge then in hand of the behavior of the APS, it would not have been clear that the approach exposure schedule was in jeopardy. Nevertheless, this software fault was extremely serious; had it occurred in the very next scheduled PhotoOp, the entire encounter activity would have been destroyed. As a result of this concern (prompted also by a similar fault in August), an extensive re-review of the AutoNav code was undertaken by non-Nav Team members. This review revealed only two or three additional problems, none so dramatically serious.

3.2.12 Post-Braille Cruise Operations—Though not formally part of the main mission validation operations, within a few weeks of Braille, navigation events began again in earnest. In order to achieve the targeting requirements for an encounter with comet Wilson-Harrington in January of 2001, it was necessary to start burning the main (IPS) engine within days of closest approach. Fortunately the desired thrust attitude was not too dissimilar to the attitude of the spacecraft with its high-gain antenna oriented on Earth. Therefore, it was possible inside of a week to be burning the main engine and take advantage of the extensive scheduled DSN tracking. Within two weeks of encounter, the first post-Braille Photo-Op navigation event took place, on 8/9. HGA-on-Earth operation of IPS continued, with additional PhotoOps on 8/16 and 8/23. The first two of these PhotoOps were very successful. However, the third revealed another coding flaw in AutoNav, where, due to a dearth of sufficiently bright targets and the need to “double-up” on a single good target at an imaging opportunity, an internal array was overrun, causing the spacecraft to safe. With the real (as opposed to opportunistic HGA-on-Earth) IPS thrusting scheduled to start on that day, a rapid spacecraft recovery took place and the mission burn began early (on 8/25). With the Nav team focussed on accomplishing the next 8 weeks of thrusting and assuring the safety of OpNav events, a one-month hiatus in PhotoOps was declared. Starting on 9/20, PhotoOp events began again; for seven weeks, these were weekly events. There was also a change of strategy. It was decided to simplify AutoNav operations: that picture planning would revert to the original design. That is, that optical frames would be “bore-sighted” on the asteroid target (actually the targets had to be substantially offset from the center of the field in

the CCD, due to large, severely attenuating scratches in the optics at that point) and the system would acquire any available stars. This substantially reduced the “man-handling” of the system and allowed the system to operate in truly autonomous form.

Figure 27 shows the post-fit residuals for this solution, the data-arc extending from 9/27 to 11/1. These residuals make an interesting comparison with Figure 14, showing a factor of 2 to 3 improvement in image-processing performance with a drastic reduction in effort. In fact, the effort was literally reduced to zero; for the period of time shown in Figure 27, the spacecraft was navigating itself, with no updates or changes to its process. This turned out to have substantial advantages: with several critical programs operating (and experiencing navigational problems), the DS1 Radio Nav Team was released to concentrate on these challenges, while DS1 navigated itself. This is perhaps the best characterization of the validation of AutoNav.

4.0 TECHNOLOGY VALIDATION SUMMARY

4.1 Summary Overview

The overarching philosophy behind AutoNav testing was to initially ground test every operation of AutoNav under both normal and a selection of abnormal circumstances. Once in flight operations, the first few events of a given Nav operation were always thoroughly tested on various testbeds. Only after several successful operations under this closely simulated test restriction were the autonomous systems allowed to operate without a very well-tested predict of the expected outcome. The principal difficulty in this strategy was the early, almost complete lack of predictability of the behavior of the scattered light and leakage within the MICAS camera. As discussed in the body of the report, this problem caused general failure of the image-processing algorithms, depriving subsequent functions of data and altering the expected behavior of the AutoNav sessions. In no case, however, was this inability to predict considered to be (nor did it at any time prove to be) a hazard.

The “Fact Sheet” summary table of AutoNav Validation plan and success gives a succinct summary of all of the validation events undertaken. Where applicable, and especially where they were explicitly noted in the Technology Validation Agreement (Appendix F), quantitative goals and achievement levels are listed. In general, there is a range of achievement in these values; where this is so, best and worst values are noted. In the body of the report, especially Section 3, the history and conditions of these variously good and bad results are discussed at length.




Figure 27. Post-Braille AutoNav DataArc and Residuals

4.2 Pre-Flight Validation

4.2.1 Prototype Demonstration—The concept of an autonomous optical navigation system was proved in a MATLAB® simulation of a ballistic mission to an asteroid. This demonstration simulated pictures taken in flight by such a mission, processed those pictures and used the reduced data in an orbit-determination estimation process. Subsequently, maneuvers were computed to control accumulated errors in the simulated orbit due to OD errors, non-gravitational model errors, and perturbations. Finally, the encounter was simulated with late tracking and orbit updates of the target. Results from this simulation gave strong indication that orbit quality of better than 500 km and 0.5 m/s was possible, as well as delivery at the target to better than 10 km.

4.2.2 Development Bench-Testing—As the actual flight system began to develop, tests were on-going, covering a wide range of expected mission-operating conditions. Early in this process, the decision was made to make DS1 a low-thrust mission, requiring a substantial increase in the complexity of AutoNav. Extensive new theoretical development and test was required (see Appendix E). Of a large number of missions considered and partially evaluated, a mission to asteroid McAuliffe, then Mars, followed by a flyby of comet West–Kohoutek–Ikemura was settled upon and extensively evaluated. The extensive cruise phases were simulated and OD performance evaluated, and the ability of the maneuver planner to keep the spacecraft on course was robustly demonstrated. (This mission was

subsequently replaced by the current 1992KD, Wilson–Harrington/Borelly mission, due to a required launch delay.) None of these tests gave performance and capability results in conflict with the prototype demonstration phase.

4.2.3 Software Module Delivery and Version Testing—Each of the elements of AutoNav went through element tests and extensive system tests as part of the delivery process of each new version of the software. The system tests covered various mission phases and all of the interactions and functions of Nav. Additionally, AutoNav systems, particularly the ephemeris services, were required for all other system tests, leading implicitly to additional Nav verification. None of these tests gave performance and capability results in conflict with the prototype demonstration phase.

4.3 In-flight Validation

4.3.1 Early Cruise AutoNav—Upon the first invocation of the higher AutoNav functions in flight, it was obvious that pre-flight performance estimates would not be met; this was almost entirely due to the problems encountered with MICAS. Because of the scattered-light leakage problems, it was impossible to successfully acquire navigational data onboard before extensive AutoNav flight-software modifications were performed. However, even ground processing of the onboard-acquired images revealed problems, keeping the performance of the system (as demonstrated on the ground) above 5000 km and 2 m/s.

4.3.2 Late Cruise AutoNav—By 6/99, all modifications had been made to the cruise AutoNav system, including image processing changes to deal with the scattered light-leakage problems, and severe geometric distortions observed in the field. With these changes and calibrations onboard, the performance of the onboard-cruise navigation on several occasions met the original technology-validation agreement (better than 250 km and 0.5 m/s). However, due to the continuing uncertainty of the geometric distortions, this could not be continuously maintained without hand-editing data on the ground.

4.3.3 Encounter Phase: Rehearsal—As with all previous bench and test-bed testing, when the encounter rehearsal (the final 6 hours of approach operations) was performed onboard, AutoNav met all performance requirements. This included computing and executing a TCM to within 2.5 km of the desired target and keeping the target asteroid (in this case simulated) in the spacecraft field-of-view to within 30 seconds of closest approach, effectively reducing the post-control knowledge error to under 0.5 km in the final field of view. All encounter sequences were started at the appropriated times (within the statistical variation). This performance level, though a rehearsal, was onboard closed-loop autonomous control and met the validation requirements.

4.3.4 Encounter Phase: Actual—Because of an uncorrected electronics fault in the MICAS CCD, it was necessary for AutoNav to switch detectors to the less capable and less well characterized APS channel shortly before encounter. With nearly all of the science and all of the Nav data scheduled from this sensor within 30 minutes of closest approach, the approach sequence was extremely dependent upon models that described the expected brightness of the approaching target. At encounter, the target was far dimmer than expected for at least two reasons. First, the photometric predictions were inaccurate due to the inextendability of the assumed models to the encountered geometry and the lack of allowance for an unfavorable presentation of an oblong object to the approaching spacecraft. Second, the APS sensor exhibited extreme non-linearity at low signal, causing a flux, dimmed by the first phenomenon, to have its signal obliterated. As a consequence, no useable signal was received and close-approach AutoNav did not support the Braille encounter.

5.0 APPLICATION OF AUTO NAV TO FUTURE MISSIONS

5.1 Requirements for Use of AutoNav

Of course, the principal requirement for using an onboard autonomous optical navigation system is a suitable space-science-class imaging instrument. Other requirements include suitable CPU performance and RAM-addressable

program memory and mass-storage (although AutoNav's requirements on the latter two are relatively modest, at about 4.5 and 5 MB, respectively). The CPU performance requirements are somewhat less easy to quantify and will reflect the speed with which the mission requirements call for the "Nav Loop" to be closed. In the case of DS1 at Braille, it was necessary to process pictures from the APS detector in as short a period as 4 s to keep the target "locked" in the field of view as late as possible. AutoNav also depends upon the existence of a very capable and intelligent ACS system, which provides accurate pointing control and knowledge, as well as planning support for turns. The latter includes a predictive ability for computing the expected length of turns. Also necessary is a DS1-like comprehensive ability to protect the spacecraft body under varying circumstances from forbidden orientations and to predict or judge the violation states of certain attitudes. Another ability for which DS1 rests with the ACS system is the capability to vectorize TCMs, as discussed earlier.

5.2 Types of Missions that can Use AutoNav to Advantage

There are various features that have made AutoNav on DS1 advantageous to mission operations and that offer opportunities for future missions. The most basic is the ability of the system to obtain navigational data without the need for Earth-based radio tracking. Another is for AutoNav to make quick "turn-around" closed-loop decisions, without the need for ground intervention. Yet another feature offered by AutoNav was complete automation of intensive Nav-related activities, such as OpNav picture taking, TCM, and IPS mission burns. Such events on all previous missions required extensive sequence, test, and validation activity, most of which was done for DS1 autonomously onboard. These features of AutoNav can, at least potentially, reduce some navigational and other operational costs and improve science return. Depending on the type of mission, the various features can have important or even enabling effects. Missions with severely limited tracking schedules or ability would, for example, not be stressed by the need for navigational tracking. Missions with very complicated dynamics can take advantage of quick-turn-around onboard OD and maneuvers, such as orbital tours of the gas giants. And, clearly, rendezvous missions and flybys (such as DS1) can take advantage of on-site ephemeris updates for improved science return in a way that cannot be duplicated with ground-based processing.

5.3 Adaptations Necessary or Desirable for Future Use

5.3.1 Adaptations for Cameras—Obviously, different missions will have different imaging systems, which will have to be modeled and calibrated, perhaps requiring updates to the distortion model itself, as MICAS did. Parameters applying to the camera and maintained within the AutoNav model include focal length, pixel size, camera sensitivity, and pixel aspect ratio. Different cameras will likely have different means of specifying exposure times

and may have filter specifications, the latter of which MICAS does not have. Some cameras have anti-blooming algorithms, which can clearly be used to advantage (and might have cured or attenuated the bright-object charge bleed problem). Additional channels, or entirely independent cameras on the same spacecraft, could be easily accommodated. Software to automatically compensate for unexpectedly low light levels could be used to advantage during encounters with poorly characterized objects. Use of a scan platform or orientable mirror would require relatively minor model changes to the image processor and OD algorithms.

5.3.2 Dynamic-Model Upgrades—As with the camera, the non-gravitational nature of each spacecraft is different. Although AutoNav’s treatment of the problem is fairly general, modifications for a different spacecraft might be necessary. It should be pointed out that the requirements of optical data on dynamic model accuracy are relatively low. There have been proposals for autonomous navigation systems that use a reversed radio link (i.e., a radio beacon is tracked by the spacecraft; from the onboard interpretation of this signal, the spacecraft state is inferred). On approach to a target, optical navigation can achieve 1-km accuracy with dynamic modeling accuracy of 0.25 km and 0.1 m/s *target-body relative*. To achieve the equivalent accuracy with a radio beacon from *onboard* would require at least 0.005-km and 0.0001-m/s accurate modeling, *Earth tracking station relative*. This is not at all easy and would be very difficult in an onboard autonomous system. Left unsolved with the radio approach is the resolution of unreduced target ephemeris errors.

5.3.3 Ephemeris Extensions—Additional ephemerides for satellites, or the ability to estimate the ephemeris errors of asteroids could enhance the capability of AutoNav. If substantial errors in the ephemerides are expected for the satellites of a planetary target (those satellites being used as navigational targets) then the ability to model and estimate elements of those satellite orbits will be necessary. Again, however, because of the relative insensitivity of optical data to dynamic modeling, the satellite positions need not be described to substantially better than their observability in the camera.

5.3.4 Image Analysis Extensions and Enhancements—For a mission dependent upon extensive imaging and analysis of a large or near body (such as a flyby or rendezvous with a major planet, or a rendezvous and orbit of a small body), *DS1* AutoNav would require upgrades to use appropriate large-object optical data, such as limbs and landmarks. Such algorithms are a standard part of the existing suite of ground optical, navigation tools; such tools are readily adaptable to AutoNav, in the same fashion as other AutoNav capabilities were adapted. The ability to autonomously generate topographic maps onboard is also possible (and in fact

planned) as a future development of the system, which would have substantial benefits to a mission orbiting a poorly characterized object, such as an asteroid. Comets also provide substantial challenges to image analysis. *DS1* AutoNav has only begun to develop some of the autonomous algorithms necessary to deal comprehensively with the variety and severity of the visual environments expected in the near environments expected.

5.3.5 Software and Spacecraft System Adaptations—As is only natural, a change in the underlying VxWorks operating system or support system from that used by *DS1* will force modifications. Principal features of the *DS1* system include the inter-process messaging system and timing services (both updated versions of the Mars Pathfinder systems). As part of the critical software foundation of AutoNav is the structure and nature of the commands available to AutoNav for its work, the most vital of these being the ACS interactions. Other missions may also wish for more substantial interactions with Fault Protection, especially for orbiters where AutoNav may wish to call for an emergency “escape maneuver” during a close-orbiting phase.

5.3.6 Picture Planning Full Automation—One of the least automated features of AutoNav is the picture-planning process. Though requiring only minimal inputs (namely a list of prospective good asteroids), the picture planner is able to resolve all further planning issues, such as turn and timing constraints. Nevertheless, the initial list must still be generated on the ground. Also, AutoNav will not replot or cancel a picture based on positions or paucity of stars, all of which could have been advantageous during *DS1* cruise. However, if a cruise navigation camera has performance similar to that expected originally from MICAS, and none of that instruments faults, a simple “just look at any near-by asteroid” strategy will, in the vast majority of cases, get adequate stars for navigation. Fully automated picture planning will be important, perhaps vital, however, for missions that depend upon landmarks for navigation (e.g., planetary or asteroid orbiters).

5.3.7 Multiple Spacecraft Navigation—For missions with multiple spacecraft performing optical navigation, substantial benefit can be obtained by letting the ships communicate and share their data. This will require some substantial logistical modifications to the OD subsystem, in particular, to allow observations from two uncertain platforms. However, the potential gain is great to obtain independent observations of an approach target from two different inertial references. The Deep Impact mission will likely make use of this capability.

6.0 CONCLUDING REMARKS

The *DS1* mission was rich with remarkable challenges. For those working in the *DS1* development environment, it was

alternately exhilarating and frustrating, with days variously triumphant or terrifying. When working at its best (which it usually did), this small team worked incredibly long hours, heedless of team boundaries, toward the single goal of getting this ensemble of groundbreaking technologies off the ground. The AutoNav team, perhaps more than any other, had the privilege of working in close technical connection with virtually all of the other segments of the mission. In fact, Navigation became something of an integrating factor in the mission operations, intimately connecting Mission Design decisions to flight software, to ACS, to science, and to IPS operations, as well as sequencing Telecom and Testbed operations. This thorough integration into the mission development and operations was unprecedented for the navigation function on JPL deep-space missions and it made the eventual success of the mission overall, and Nav in particular, that much more satisfying for the team. In addition to being well integrated into the overall flight system, AutoNav, more than any other subsystem, was vitally dependent on other technologies and subsystems for its validation, particularly Mission Design, MICAS, ACS, IPS, flight software and Science. With the important exceptions of the problems discussed in the body of this report, the performance of these systems was very good. The working relationship between ACS and Nav, from organizations that according to folklore cannot work together, could not possibly have been better. In fact, it was the maturity and professionalism of all of the teams, especially in the face of what were often staggering obstacles and timelines, that made the working environment of DS1 a good model toward which most projects and individuals could work to their great benefit.

7.0 ACKNOWLEDGMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Special Thanks are due to a number of individuals without whom AutoNav could not have been developed and proven. Marc Rayman, Phil Varghese, Leslie Livesay, and Dave Lehman, as leaders of the DS1 project, were continuously supportive and contributed technically to the project's success. Steve Williams and Pam Chadbourne of the Mission Design team made critical inputs to the navigation effort from the standpoint of mission design and planning. The DS1 software team lead by Dan Eldred helped the AutoNav team mightily as it ventured for the first time into the realm of deep-space flight software. The Testbed team lead by Paula Pingree and Greg Harrison made the extensive testing of the AutoNav systems possible. Outstanding support in the sequencing area was provided by Curt Eggemeyer, Kathy Moyd, and "P.J." Guske. In the Navigation area, many thanks are due to Don Yeomans,

Alan Chamberlin, and Mike Keesey for asteroid ephemeris development. Bill Owen was also of tremendous help in the actual taking of ephemeris observations (especially of Braille), helping with the calibration of MICAS, and providing star catalogs. Finally, and perhaps most importantly, thanks go to the ACS team, lead by Sima Lisman and including Dan Chang, Tony Vanelli, Steve Collins, Sam Sirlin, Guru Singh, and John Esmiller. As mentioned earlier, there is organizational folklore regarding navigation and ACS functions not working well together. For DS1, the working relationship was as one team and could not have been better. To everyone else on DS1 and elsewhere who provided much vital assistance (but who have negligently not been thanked here), the Nav team sends its most heartfelt gratitude. It seems a general feeling on the project that no one in their careers has ever worked so hard, accomplished so much, been so fearful, or had so much fun all at the same time. DS1, in many ways, has been "as good as it gets."

8.0 REFERENCES

- [1] J. K. Campbell, S. P. Synnott, J. E. Riedel, S. Mandell, L. A. Morabito, G. C. Rinker, "Voyager 1 and Voyager 2 Jupiter Encounter Orbit Determination," AIAA Paper 80-0241/AIAA Aerospace Sciences Meeting, Pasadena, California, January 1980.
- [2] J. K. Campbell, R. A. Jacobson, J. E. Riedel, S. P. Synnott, A. H. Taylor, "Voyager 1 and Voyager 2 Saturn Encounter Orbit Determination," AIAA Paper 82-0419/AIAA Aerospace Sciences Meeting, Orlando, Florida, January 1982.
- [3] S. P. Synnott, A. J. Donegan, J. E. Riedel, J. A. Stuve, "Interplanetary Optical Navigation: Voyager Uranus Encounter," AIAA Paper 86-2113/AIAA-AAS Astrodynamics Conference, Williamsburg, Virginia, August 1986.
- [4] J. E. Riedel, W. M. Owen, J. A. Stuve, S. P. Synnott, R. M. Vaughan, "Optical Navigation During the Voyager Neptune Encounter", AIAA paper 90-2877-CP, AIAA/AAS Astrodynamics Conference, Portland Oregon, August 1990.
- [5] R. M. Vaughan, J. E. Riedel, R. P. Davis, W. M. Owen, S. P. Synnott, "Optical Navigation for the Galileo Gaspra Encounter," AIAA paper 92-4522, AIAA/AAS Astro-dynamics Conference, Hilton Head, South Carolina, August 1992.
- [6] V. Coverston-Caroll, Williams, S. N., "Optimal Low Thrust Trajectories Using Differential Inclusion Concepts", AAS/AIAA Spaceflight Mechanics Meeting, Cocoa Beach Florida, February 1994.
- [7] R. J. Koshel, G. Peterson, "Stray Light Analysis of MICAS," BRO Report #3580, Breault Research Organization, Inc, Tucson Arizona, September 1999.

Autonomous Optical Navigation
(AutoNav)
DS1 Technology Validation Final Report

Appendix A
AutoNav Listing of
Telemetry Channels

Below is a list of all of the telemetry channels that the AUTONAV team collects and uses. In addition there is a set of AUTONAV specific files that are downlinked. Also AUTONAV telemetry is contained in apids 17 and 19. (Ed Riedel 10/20/99)

Channel	Mnemonic
N-0101	img_cmplt_st
N-0102	OD_cmplt_st
N-0103	mvr1cmplt_st
N-0104	mvr2cmplt_st
N-0105	setThrsCmplt
N-0106	tcm_type
N-0107	tcm_cmplt_st
N-0108	updtIPSCmplt
N-0109	name_upd_st
N-0110	NAVRT_upd_st
N-0111	ThrsPrsCmplt
N-0116	FileRemaindr
N-0117	append_file
N-0118	ephemRequest
N-0121	OD_CnvergNum
N-0122	FilRecordCnt
N-0123	target_id
N-0124	NumberOfObs
N-0125	PicsProcessd
N-0126	Num_Images
N-0127	EphemReqTotl
N-0128	InvlEpemReq
N-0129	spr_nav_029
N-0141	nav_machine
N-0142	nav_burn_st
N-0143	photo_op_st
N-0144	nav_tcm_st
N-0145	nav_exec1_st
N-0146	nav_exec2_st
N-0147	nav_exec3_st
N-0148	nav_exec4_st
N-0149	maneuver_id
N-0150	thrust_level
N-0151	updateThrust
N-0152	tcm_segments
N-0153	fileID_req
N-0154	change_IMODE
N-0155	thrust_press
N-0156	LinesOfSight
N-0157	numbr_images
N-0158	EphemRecTim

N-0159	IPSdurationT
N-0160	sc_epoch
N-0161	norm_od_xhat
N-0162	vector_X
N-0163	vector_Y
N-0164	vector_Z
N-0165	RCS_dltav_X
N-0166	RCS_dltav_Y
N-0167	RCS_dltav_Z
N-0168	ResidualMean
N-0169	StandrdDev
N-0170	ResidualMin
N-0171	ResidualMax
N-0172	sc_sun_X
N-0173	sc_sun_Y
N-0174	sc_sun_Z
N-0175	sc_sun_Xdot
N-0176	sc_sun_Ydot
N-0177	sc_sun_Zdot
N-0178	IPS_impulseX
N-0179	IPS_impulseY
N-0180	IPS_impulseZ
N-0181	photo_op_tim
N-0182	img_proc_tim
N-0183	preOD_strTim
N-0184	preOD_comTim
N-0185	OD_strt_tim
N-0186	OD_cmplt_Tim
N-0187	OD_perfrmTim
N-0188	man_plan_tim
N-0189	find_mvr1Tim
N-0190	find_mvr2Tim
N-0191	thrustLvlTim
N-0192	tcm_time
N-0193	updtThrstTim
N-0194	BrnDurMsgTim
N-0195	EmergBckTim
N-0196	NAVresetTime
N-0197	thrstPresTim
N-0198	sc_Earth_X
N-0199	sc_Earth_Y
N-0200	sc_Earth_Z
N-3000	ScSunRa
N-3001	ScSunDec
N-3002	ScSunDist
N-3003	SunScVelRa

N-3004	SunScVelDec
N-3005	SunScSpeed
N-3006	ScEarthRa
N-3007	ScEarthDec
N-3008	ScEarthDist
N-3009	HstDvRa
N-3010	HstDvDec
N-3011	HstDvSpeed
N-3012	HstIpsImpRa
N-3013	HstIpsImpDec
N-3014	HstIpsImpls

APIDs 17 and 19

Autonomous Optical Navigation
(AutoNav)
DS1 Technology Validation Final Report

Appendix B
AutoNav Power On Times
of Data Capture

Below is a summary of the AutoNav Activities performed, a detailed description is included in the DS1 AutoNav Technology Validation Report. Starting with the 03/01/99 AutoNav activities, DS1 began a period of 10 weeks of normal operations, which included weekly Photo-Op/OD/ ManPlan sequences, and periods of Mission Burns. (E. Reidel 11/23/99)

Time (UTC)	AutoNav Activity
10/24/98T12:08	Launch
11/06/98	First picture taken with MICAS
11/18/98	First AutoNav Photo-Op session
12/03/98	200 hours of thrusting achieved
12/18/98	First operation of AutoNav NBURN
12/21/98	Second Photo-Op attempt
12/22/98	Second NBURN
01/06/99	NAV File load
01/07/99	Third Photo-Op
01/18/99	Photo-Op
01/20/99	Photo-Op
01/26/99	Photo-Op
02/01/99	Photo-Op
02/08/99	Upgraded AutoNav image-processing software loaded (M4)
02/18/99	First Photo-Op with the M4 software
02/19/99	NAV File load
02/22/99	Photo-Op/OD/ManPlan
02/27/99	Update AutoNav Control Modes
03/01/99	Photo-Op/OD/ManPlan
03/8/99	Photo-Op
03/15/99	Photo-Op
03/16/99	Second part of mission burn with NAV moderated thrusting
03/22/99	Photo-Op/OD
03/29/99	Photo-Op/OD
03/29/99	ManPlan
04/05/99	Photo-Op/OD/ManPlan
04/12/99	Photo-Op/OD/ManPlan
04/19/99	Photo-Op/OD/ManPlan
04/26/99	Photo-Op/OD
05/06/99	Photo-Op/OD
05/24/99	Photo-Op/OD
05/26/99	Photo-Op/OD
05/29/99	Photo-Op/OD
05/31/99	Photo-Op/OD
06/01-06/10/99	Loaded M6 software
06/10/99	Fist Photo-Op/OD/ManPlan with the M6 software.

	ManPlan successfully planned an IPS TCM for 06/14/99.
06/14/99	First IPS TCM
06/16-06/20/99	Photo-Op/OD
06/23/99	Photo-Op/OD
06/29/99	Photo-Op/OD
07/02/99	Photo-Op/OD/ManPlan
07/04/99	Photo-Op/OD/ManPlan
07/06/99	Photo-Op/OD/ManPlan
07/13/99	Asteroid Encounter Rehearsal
07/16/99	Photo-Op/OD
07/18/99	Photo-Op/OD
07/19/99	Photo-Op/OD
07/19/99	Loaded final best-ground determined Braille ephemeris
07/20/99	Photo-Op/OD
07/21/99	Photo-Op/OD
07/22/99	Photo-Op/OD/ManPlan
07/23/99T14:30	-5day IPS TCM
07/24/99	Photo-Op/OD
07/25/99	Photo-Op/OD
07/26/99	Photo-Op/OD
07/27/99T03:00	Photo-Op
07/27/99T10:00	Photo-Op
07/27/99T18:30	-1day TCM
07/28/99T00:00	Photo-Op
07/28/99T11:33	Photo-Op
07/28/99T16:00	Data downlink
07/28/99	-6hr TCM
07/29/99T00:00	3 Photo-Ops 2 ODs
ACA-27 min	RSEN mode activated

Autonomous Optical Navigation
(AutoNav)
DS1 Technology Validation Final Report

Appendix C
Navigation for the New Millennium:
Autonomous Navigation for
Deep Space 1

NAVIGATION FOR THE NEW MILLENNIUM: AUTONOMOUS NAVIGATION FOR *DEEP SPACE 1*

J. E. Riedel, S. Bhaskaran, S. P. Synnott, S. D. Desai,
W. E. Bollman, P. J. Dumont, C. A. Halsell, D. Han, B. M. Kennedy,
G. W. Null, W. M. Owen Jr., R. A. Werner, B. G. Williams

Navigation and Flight Mechanics Section
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California; USA; 91109

Phone: 818-354-8724, Fax: 818-393-6388, Email: jer@radtran.jpl.nasa.gov

ABSTRACT

The first flight of NASA's New Millennium Program, Deep Space 1, will include a new navigational technology: an autonomous optical navigation system. The *DSI* Navigation system will be the first use of autonomous navigation in deep space. The task for this system is to 1) perform interplanetary cruise orbit determination, using images of distant asteroids, 2) control and maintain the orbit of the spacecraft using the ion propulsion system (another technology never before applied to deep space) and conventional thrusters, 3) perform approach orbit determination and control using images of the science targets, 4) perform late knowledge updates of target position during close fast flybys in order to facilitate a high degree of quality data return from 2 targets: asteroid McAuliffe and comet West-Kohoutek-Ikemura. Additionally, an encounter with Mars will probably be performed with possibly a close flyby of one of the Martian moons, Phobos or Deimos. Several functional components are necessary to accomplish these tasks. These include picture planning and image processing, dynamical modeling and integration, planetary ephemeris and star catalog handling, orbit determination data filtering and estimation, maneuver estimation, spacecraft ephemeris updates and maintenance, and general interaction with the other onboard autonomous systems. These systems are described, as is the means of their operation onboard. Finally, performance statistics from trial runs of the system are given.

INTRODUCTION

Autonomous onboard optical navigation will be a necessary component of autonomous spacecraft operations for many future planetary exploration missions. Because of light-travel times, there are experiments and even missions that cannot be performed or have limited data potential unless autonomous navigation systems are incorporated. Close orbits or very fast flybys of small poorly characterized objects are examples of such missions. Reducing operational complexity and costs is another goal of autonomous navigation systems. In the not-too-distant

future, many small robotic missions may be simultaneously exploring the solar system. To increase the efficiency of these missions, the spacecraft must take on more of the responsibilities of their own maintenance, including navigation. Adapting many of the techniques proven for optical navigation for *Voyager* and *Galileo*, the *New Millennium DSI* onboard navigation system must autonomously plan picture sequences, perform image analysis, estimate the trajectory and calculate trajectory corrections using the low-thrust solar-powered ion propulsion system (IPS). *DSI* will be the first planetary exploration mission to autonomously navigate all post-injection phases of its mission. The engineering of such a navigation system poses a number of very significant challenges. An overview of Optical Navigation and how it will be applied to *DSI* is given in Ref. 1.

This first experiment in deep space autonomous navigation will be a closely monitored experiment. As a means of validating the performance of the onboard navigation system, a conventional ground radio-navigation campaign will be maintained. This ground effort offers the further advantage of providing very high quality calibrations of IPS engine performance, something which the flight navigation system (The "Navigator") would not be able to do. Though the Navigator is designed to be capable of fully autonomous operation, with many new technologies been tried on *DSI*, the capability has been maintained to quickly intervene with, and modify the behavior of the system if mission emergencies require.

DSI MISSION ATTRIBUTES

An overview of the New Millennium Program and *DSI* in particular is given in Ref. 2. The *DSI* mission includes a very ambitious and challenging set of mission objectives and activities. Three targets are intended for flyby encounters: asteroid McAuliffe, Mars, with possibly a close flyby of one of the Martian moons, and comet West-Kahoutek-Ikemura (WIKI). Currently, it is anticipated that launch will occur in July of 1998. The McAuliffe encounter

will happen late January of 1999, the Mars flyby in late May of 2000, and the comet encounter about six weeks later. Figure 1 shows a heliocentric view of a likely mission trajectory, with important mission events annotated. The annotations are referenced to Table 1.




Figure 1. *DSI* Mission Design

For the McAuliffe flyby, the *DSI* spacecraft will perform the closest flyby encounter ever attempted in a deep space mission: 10 or perhaps even 5 km from the surface of the asteroid. The encounter parameters of Mars have not yet been determined, but the flyby altitude of the comet will likely be on the order of several hundred kilometers, due to the dangerous environmental conditions near even a relatively inactive comet such as W-K-I.

ID	Time of Event	Description of Event
A	Jul. 1, 1998	<i>DSI</i> Launch
B	Oct. 24, 1998	End of first principal thrust arc
C	Dec. 6, 1998	Beginning of second thrust arc
D	Dec. 27, 1998	End of second thrust arc
E	Jan 16, 1999	McAuliffe encounter
F	Jan 20, 1999	Beginning of third thrust arc
G	Feb. 8, 2000	End of third thrust arc
H	Apr. 26, 2000	Mars encounter
I	Jun. 4, 2000	W-K-I encounter

Table 1. Principal *DSI* Mission Events

The ambitious nature of these encounters is enabled solely by the presence of the autonomous navigation system. Performing navigation functions in a closed-loop sense onboard the spacecraft makes possible very late (before encounter) controls of the spacecraft encounter coordinates, and updates of knowledge about those coordinates.

The objectives of the New Millennium Program (of which *DSI* is the first mission) is to develop and demonstrate new technologies which can enable future space exploration missions. The Autonomous Navigation System is one of these technologies being demonstrated. Another such




Figure 2. *New Millennium DSI* Spacecraft

technology, and one that has a fundamental influence on the nature of the *DSI* mission is its solar electric propulsion system. This system is actually composed of two technologies, a 2.5 kilowatt concentrator-element solar-electric array, known as "SCARLET," and an ion propulsion system (IPS) capable of approximately 100 mNt of thrust, known as "NSTAR". The IPS is principally responsible for making the energetically difficult triple encounter mission possible. However, this propulsion strategy seriously complicates the navigation task. Fig. 2 shows a schematic of the spacecraft, with annotations for the prominent solar arrays, the MICAS camera, and the IPS location on the -Z axis.

MISSION DESIGN IMPACTS ON THE NAVIGATION SYSTEM

Ion Propulsion System

The most challenging aspect of the *DSI* navigation task is the low-continuous-thrust, non-ballistic trajectory. This challenge begins with the design of the mission trajectory, which has been detailed elsewhere (Ref. 3). This highly interactive non-linear process is at the time of this writing, in its final stages for *DSI*. The trajectory is refined almost on a daily basis to reflect changes in the mass of the spacecraft, available power from the solar panels, available launch vehicle capacity and injection conditions, and thrust and efficiency of the engines. Once this design is complete however, it will be made available to the Navigator in the form of polynomial description of engine thrust direction and level as a function of time. A nearly final version of these tables is shown in Figs. 3-5.




Figure 3. IPS Thrust Clock Beam Angle




Figure 4. IPS Thrust Beam Cone Angle

The mission trajectory is divided into segments and sub-segments. The process of searching for the optimum energy

path to the targets places gaps in the thrust arcs, and additional gaps are forced in areas where no thrusting is desired, such as on approach to encounter targets. Additionally, gaps are introduced into the thrust arcs at regular intervals to accomplish OpNav observations and telecommunication.




Figure 5. IPS Thrust Magnitude

The next navigation challenge posed by the presence of the IPS is the need to control the engine. It is not sufficient to guide the engine along the pre-computed polynomial functions. There are error sources in the implementation of the nominal design, with accuracies of between 1 and 2 percent expected. Such errors, when combined with normal statistical navigation errors, could map to millions of kilometers over a seven month trajectory. Thus, the nominal mission design needs to be constantly corrected to account for these errors. Additionally, the presence of the continuous thrust of the IPS requires the Navigator to account for this force and its errors in the dynamic model of the spacecraft's course, and in the treatment of the optical data.

There is substantial uncertainty with regard to the operability and reliability of the IPS and the software managers for it, all being very new technology. This uncertainty must be reflected in the Navigator, which must be designed to cope with inconsistent operation or outages. Such conditions present themselves as gross deviations from the nominal mission design. To the extent possible, the Navigator must use future control authority to correct for unpredictable and statistically anomalous trajectory perturbations. The spacecraft will be instructed to fly the planned thrust profile, representing thrusting at all available times (typically, about 92% of the time.) If outages occur, the Navigator will attempt to correct the trajectory for them. However, if the linear correction algorithm computes a flightpath to the target which is overly energetically disadvantageous to subsequent encounters, the ground will intervene with a redesigned and optimized mission.

In addition to powering the nominal low-thrust trajectory, IPS must be used for dedicated trajectory correction maneuvers during gaps in the mission thrusting, including approach to the encounter targets. The design of these maneuvers is quite different than with the use of

conventional chemical thrusters. Since the IPS thrust is much lower (40 mNt vs. 200 mNt) these maneuvers take much longer. As such, the closer the maneuver takes place to the target, the more non-linear is the process to compute the parameters. Additionally, the *DSI* spacecraft is severely constrained in orientation. Some faces of the spacecraft bus cannot be illuminated by the sun, or may be so at only shallow angles, and/or for short periods of time. Use of the IPS constrains the spacecraft to have the solar panels directly on the sun, with virtually no deviation margin. These and other constraints mean that there are significant regions of the celestial sphere at which the IPS engine cannot point. Fig. 6 shows this constraint space in body-fixed Right Ascension (longi-tude) and Declination (latitude), and Table 1 identifies the particular constraints noted. The result is that through communication with the Attitude Control System (ACS) (Ref. 4), the Navigator must ascertain if the desired maneuver direction is in a forbidden region, and if it is, redesign it to be a vector-decomposed maneuver in two directions that are allowed within the constraint space. This process is known as “vectorization.”




Figure 6. Illumination-Forbidden Regions of Spacecraft Body.

#	Constraint	Cone
1	MICAS Primary Aperture	+/- 10 deg. (+Z)
2	MICAS Optical Bench Radiator	+/- 90 deg. (+Z)
3	MICAS IR Radiator (At all times)	+/- 70 deg. (-X)
4	MICAS IR Radiator (IR in operation)	+/- 90 deg. (-X)
5	MICAS Occultation Port	+/- 1.6 deg.
6	PPU Radiator	+/- 60 deg. (+Z)
7	Star Tracker Boresight	+/- 35 deg.
8	ACS Kinematics Amplification Factor	+/- 30 deg. (+Z and -Z)

Table 2: *DSI* Constraint Space Magnitudes and Directions

Close Encounters

Another large impact on the Navigator from the rest of the system is the very ambitious nature of the mission. Next to the necessity to control the IPS, maintaining the spacecraft position knowledge and pointing through very close and

very fast flybys is the most challenging requirement on the Navigator design. The requirement to keep the encounter target in the camera field of view when possible, created the need to perform the “reduced-state” navigation as discussed below. The close flyby distance of the McAuliffe encounter requires an unprecedented control accuracy, necessitated not only by safety concerns, but also because relatively small perturbations in the flyby asymptote produce serious deviations in target-relative geometry due to the close range, possibly disturbing a carefully constructed observation experiment.

REQUIREMENTS ON OTHER MISSION SYSTEMS IMPOSED BY THE AUTONAV SYSTEM

High Accuracy Imaging Instrument

Potentially, the most obtrusive requirement that the Autonomous Optical Navigation System (AutoNav) places on the spacecraft design is for the presence of a very high quality telescope with which to perform the inter-planetary phase of the navigation task. Some periods of the approach navigation also depend upon high quality astrometry, and therefore require a science-capable telescope. Fortunately, most scientifically sophisticated deep space missions (including *DSI*) carry a camera capable of providing adequate data for the class of astrometry needed by navigation. An overview of requirements posed by AutoNav, and met by MICAS (the Miniature Imaging Camera and Spectrometer) being flown by *DSI* is given here:

- 12-bit digitization. This is required to maintain sufficient dynamic range to image bright extended objects and dim stars.
- 0.6 to 2.0 degree field of view. This is required to maintain adequate resolution for the cruise optical navigation. Typical resolution range is 5 to 40 microradians per pixel.
- 1024 x 1024 pixel array. Such an array size is the minimum standard for quality CCDs, and will determine (via the focal length) the pixel resolution.
- Capability to locate a focused unresolved image to 0.1 pixel or better. Typical focused optics give adequate point-spread-functions to provide this capability without intentional defocusing.
- 80,000 electron (e⁻) “full-well” with 50e⁻ noise. This is a description of the dynamic range and signal quality of the instrument, which is important to define the effective working span of useable brightness.
- Image 12th Magnitude star. This should be possible in a long (smeared) exposure and represents the minimum useable detection of cruise targets, and reflects the presence of accumulated photons/charge from repeated overlays of the drifting image.
- Image 9th Magnitude star. This should be possible in a short (unsmeared) exposure. Such images are the normal mode on approach to a target.

Flight Computer Requirements

The *DSI* flight computer is a RAD6000 based computer system operating at 33MHz. This computer is a radiation hardened version of an IBM-6000 series work-station computer. There are 96Mega-Bytes (MB) of hardened RAM available, which is used as both memory and mass storage. There is 16MB of non-volatile memory from which the computer boots. It is estimated that at least 50MB of RAM will be available for Science and OpNav data storage, and about one-half of the available CPU capacity will be available for Science and OpNav processing during most of the mission.

The computational requirements imposed on the flight computer and data system are relatively modest in most cases. The size of the object code in running configuration, including static variable storage, is about 2 MB. The star catalog, containing about 125,000 stars occupies about 2 MB. The ephemeris file, with the major planets and about 250 minor planets is about 0.5MB, and other miscellaneous files also occupy another 0.5MB. The code and data files will be resident in non-volatile memory (EEPROM). The spacecraft system will load the programs and data from EEPROM into RAM at boot time, and those copies will be used for processing. At least once per day, and more often during critical activities, copies of the current data, including currently best-estimated states, data summaries, and the non-gravitational force histories will be written into EEPROM to protect the data from a system failure with associated CPU reboot. At reboot, the latest stored data is recovered, and the Navigator proceeds in a normal fashion.

Timing and throughput requirements are not stringent during interplanetary cruise; there is ample time during this phase to plan the images and perform the processing. (A detailed description of the operational activities is given below). When the Navigator has an opportunity to take images, the planning process takes only a few seconds. The processing of each cruise image is estimated to take up to a minute, but since each cruise exposure is about 100sec in duration, it is thought that the precision astrometric processing will keep up with the pace of imaging; especially when considering that several minutes (up to 30) will be required to turn the spacecraft from target to target. Nevertheless, there will likely be room available in the RAM-disk space to hold a number of images if the Navigator, for some reason, is delayed in processing. When finished with image processing, the Navigator will delete the images, or select a small subset for compression and downlink, especially in the early portion of the mission. Additional computational leeway is provided from the fact that during the cruise phase, the information content of the data is not changing quickly, and therefore it is only necessary to infrequently process the reduced image data into a solution of the spacecraft state, a process which can take several minutes.

During the encounter phase of the mission, the timing requirements of the Navigator are much more stringent. In the last 5 minutes on approach to the target, a series of up to 5 OPNAV opportunities occur. These are at increasing frequency, to capture the rapidly increasing information available in the images about range to the target, knowledge of which is critical to keep the asteroid in the field of view until the last possible moment. Table 3 shows the image times, ranges, and associated spacecraft state knowledge with each of the late pictures. The timing of these frames is very close, and there is not sufficient time to perform all of the normal processing. Therefore a reduced form of the navigation processing is invoked about 30 minutes from encounter, allowing image processing and orbit determination to complete in 10 to 15 seconds. The spacecraft target-relative ACS held ephemeris is then updated with each image, by means of a simple and quick 3-dimensional bias state change to a previously delivered full 6-d ephemeris. Since these updates occurs in a matter of seconds, the target can be held within the field of view until the ACS can no longer physically accelerate the spacecraft into a turn at a fast enough rate.

Picture Time (sec)	McAullife Range (km)	Downtrack Error (km)	Crosstrack Error (km)
-20	164	0.8	0.5
-40	328	1.6	0.5
-80	656	3.2	0.5
-160	1312	7.5	0.5
-320	2624	15	0.5

Table 3: Near Encounter OpNav Picture Statistics

Interfaces with ACS, IPS and Sequencing Managers:

A number of interfaces with other flight software subsystems have already been alluded to. The most technically intricate of the inter-system interfaces is with the ACS (Attitude Control System). This interface is a set of different queries and responses. The Navigator must ask the ACS for a number of types of information: current attitude of the spacecraft; specifications on turns, such as estimated length of time required to turn from one attitude to another; the validity of a specific attitude for a maneuver; and the accumulated velocity due to general RCS (Reaction Control System - a subsystem of the ACS) activity. ACS, in turn, queries AutoNav for the current mass of the spacecraft; and current spacecraft and planetary ephemeris information. Through an indirect sequencing operation (to be discussed below) the Navigator will request the ACS to perform specific operations; for example, turning to a specific attitude, for image taking or IPS thrusting. ACS will also be asked to execute a Trajectory Correction Maneuver (TCM) with the RCS or execute a TCM with the IPS. AutoNav also maintains an interaction with the IPS manager: IPS reports to AutoNav the currently accumulated thrust while the IPS engine is operating; and AutoNav will, through the sequencing interface, request the IPS to go to a specific

thrust level and burn for a specific duration. The third principal interface that the Navigator maintains is with the Sequencer itself, and this is the simplest major interface. The Navigator will prepare very short sequences (listings of time-ordered commands) to perform specific tasks and ask the Sequencer to start or "launch" them. Additionally, during encounter, the Navigator will be called upon to launch specific encounter sequences at specific encounter-relative times.

Data Uplink and Downlink Requirements:

Necessarily, the Navigator requires a certain level of information transfer both on the uplink and downlink. This is especially so for this the first flight of the system. The early portion of the mission (the first three or four weeks) will see intense use of the telemetry system to downlink dense data sets pertinent to the evaluation of the new technologies. AutoNav will be among these. Principal among the data to be downlinked in this early evaluation period will be the OpNav images themselves. Other data will include processed results from the Navigator, including reduced image data, centers of asteroids and stars in individual frames, computed orbit determination results, and maneuver solutions. It is anticipated that after a short period of evaluation of the dense telemetered navigation data, that the data can then be reduced, compressed or stopped. On approach to the asteroid, the first target, there will again be a short burst (a few days) of dense data, to confirm that the Navigator is initiating approach operations properly.

Again, given normal performance of the AutoNav system, uplink requirements should be fairly modest. The largest sets of information likely to be required sent to the spacecraft are new thrust profiles, reflecting newly redesigned mission trajectories, and asteroid ephemerides. It will likely be necessary to redesign the mission trajectory at several points during the mission. The first such time is shortly after launch when the injection errors are known. Although nominal performance of the Delta 7326 launch vehicle is expected, greater than a one-sigma dispersion of about 100m/s will likely necessitate a redesign of the trajectory. The onboard maneuver computation algorithm will not be able to retarget the spacecraft in a fuel efficient manner in the face of such an injection error. Although the maneuver subsystem is tolerant to a certain degree of uncertainty in the engine performance, if the IPS operation deviates from the schedule by two weeks or more, it is again likely that the mission trajectory and thrust profile will have to be redesigned. Finally, it is expected that immediately after the McAuliffe fly-by that the ground operations Navigation team will redesign and uplink the trajectory and thrust profile. The process of optimizing the flight path for fuel use between two flybys is beyond the current capabilities of the flight *DSI* AutoNav system.

Operational Demands, and Staffing

Despite the expected periodic intervention of ground operations as outlined above, the AutoNav system will exhibit a high degree of autonomy. Operations, such as TCM's and image processing which used to require a significant amount of personnel on navigation and other teams will occur automatically without even the need for the ground to approve the AutoNav system's decisions. Even in the early part of the mission when extensive analysis of the operation of the onboard Navigator will be taking place, the size of the Navigation team will only be between four and five persons, and this includes at least two performing the validating conventional radio navigation task. This bodes well for future missions using versions of the *DSI* AutoNav system. It is estimated that a maximum of three persons would be necessary to fully analyze and maintain the operation of the AutoNav system for future missions at least as ambitious as *DSI*. This compares favorably with the 7 to 10 individuals necessary to perform similar functions for the *Cassini*, *Galileo* and *Voyager* missions.

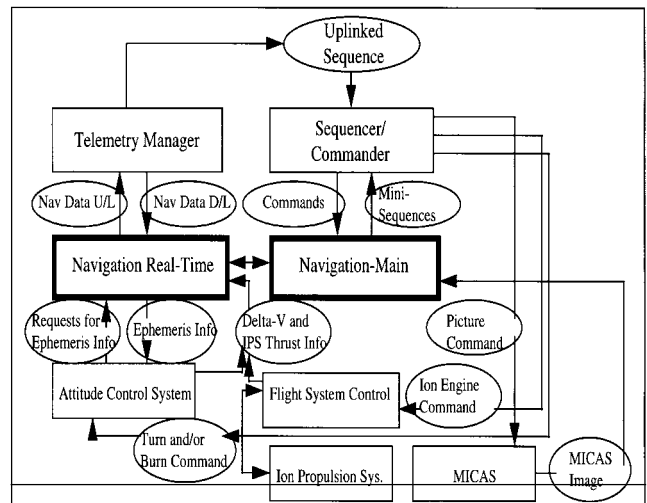


Figure 7: Navigation System Architecture

AUTONOMOUS NAVIGATION SYSTEM DESIGN:

Architecture

The *DSI* software system architecture, emphasizing the navigation system interactions, is shown in Fig. 7. The *DSI* system is based largely on the Mars Pathfinder flight software system. Mars Pathfinder is a conventionally controlled spacecraft, meaning that long series of commands (sequences) are uplinked to the spacecraft for timed execution (Ref. 5). Despite the deterministic nature of the nominal control system, autonomous navigation is still part of the design. This is accomplished by leaving large gaps in the ground-generated stored sequence, in which the AutoNav system is allowed to accomplish autonomous operations; this mode of operations will be discussed in detail below.

The Navigation system is composed of two subsystems, a real-time link, Nav-RT, and a main non-real-time computational link, Nav-Main. The real-time link is responsible for maintaining the ephemeris information for the ACS subsystem and for collecting information about propulsive activity onboard from the ACS and the IPS managers and formatting and relaying it to Nav-Main.

The flow of control through the flight software system and the Navigator is shown in Fig. 7. Normally, commands to the Navigator come via the Sequencer in an uplinked stored sequence. A summary of the possible commands that the Navigator can process is given in Table 4. All requests for action that the Navigator makes, will also be made through sequences, but these will be short and spontaneously generated onboard by the Navigator itself. In addition to the commands received by and issued from the Navigator, there are a limited number of direct calls to the Navigator and returned replies. These were summarized above.

Command	Navigation Action
[NAV-SET-IPS]	Initialize the IPS thrust arc.
[NAV-IPS-UPDATE]	Update the IPS thrust and vector.
[NAV-DO-TCM]	Perform TCM operations.
[NAV-PHOTO-OP]	Plan and take Navigation Pictures
[NAV-START-ENC]	Start an Encounter sequence.
[NAV-DATA-UPDATE]	Update Navigation parameters.
[NAV-DO-OD]	Perform Orbit Determination.
[NAV-PLAN-TCM]	Compute TCM parameters.

Table 4. Navigation Command Summary

Functional Overview

At the most basic level of description, the AutoNav system uses pictures taken by the onboard camera to determine, *via* a batch-sequential stochastic filter, the spacecraft state. After propagating this state to the target body, retargeting parameters are computed and trajectory correction implemented. During the cruise portion of the mission, pictures of asteroids and stars are the principal data, but on approach to a target, images of that target with or without stars are the main navigational data. In the following sections, these functions will be detailed.

Image Planning

The task of the Image Planning subsystem is to provide a schedule of targets for the AutoNav system. These targets include both beacon navigation targets as well as the approach encounter targets. The targets are clustered in time, to enable the planner, when asked, to access a set of viable target-asteroids to use for navigation purposes. The targets are additionally clustered and ordered to minimize attitude changes. Minimizing the cost of the turn sequences is important to minimize fuel usage. Because of the nature of the illumination constraints on the spacecraft, the beacon asteroids cluster into two discrete groups: those in the

“forward” anti-sun half-hemisphere, and those in the “aft” anti-sun half-hemisphere. A fuel and time costly rotation of the spacecraft is necessary to turn from forward to aft, and so at most one such turn is scheduled for each observation opportunity. Within each half-hemisphere, the turns are additionally minimized.

Even though the above considerations are made as part of the ground operations, and possibly even before launch, there is a substantial amount of work for the onboard picture planner to do. Given only a list of asteroid targets, in optimized turn order, the picture planner must assemble a set of specific image requests, including turn commands for exact pointings in inertial space. Additionally, it must predict the locations of the stars to be seen in the field relative to the target at precisely the time the picture is to be taken. This requires accurate storage and evaluation of ephemerides and star positions. The former will be discussed later, but the latter involves the use of accurately built star catalogs and requisite efficient storage of them. For *DSI*, the onboard star catalog will be based on the TYCHO Star Catalog (Ref. 6) and contains about 125,000 stars. The positions on this file are accurate to at least 5 micro-radians, at least factor of two greater than is required to avoid degrading the accuracy of the autonomous OD process.

Image Processing

There are two types of images taken during the mission, long-exposure smeared images of unresolved beacon asteroids, and short-exposure images taken on approach to a target. These latter are pictures of resolved and extended images.

In deep cruise, the need for long exposure images arises from the small size and extreme range of the beacon targets. The consequence of these long exposure times is to cause the ambient motions of the three-axis-inertial stabilized spacecraft to trace the star images over extended parts of the frame. Typical star and asteroid images will be smeared over 20 to 40 pixels. Fig. 8 shows a simulated version of the expected deep space image. Frames such as this have been used to test the algorithms and software. Also, simulations of the expected sort of image have been made using an astrometric observing system at the JPL observatory at Table Mountain. A series of these images, made to simulate the unstable characteristics of the spacecraft, was made by manually slewing the telescope with its joystick controls. These images were then processed by the image processing subsystem of the Navigator. This analysis is documented in Ref. 7.

The processing system for the smeared cruise images was developed for the *Galileo* mission, and is documented elsewhere (Ref. 8) The theoretical basis of the system is a multiple-cross-correlation algorithm, that uses each of the nearly identically smeared star and asteroid images in a

picture as a pattern. Each pattern is then used to locate every other pattern, with the result that extremely complex and often faint patterns can be located relative to one another to high accuracy, usually to 0.1 pixel (picture element) or better.

The actual correlation process can be summarized as a vector inner product. Given a normalized pattern, called a "filter", that is composed of image elements in a matrix $m \times n$ in size denoted as F , and a sample area S , $M \times N$ in size, of which subset regions of $m \times n$ dimensions are extracted, then a function c_{ij} can be maximized:

$$c_{ij} = F \otimes S_{ij} = \sum_{k=1}^m \sum_{l=1}^n F^{kl} \cdot S_{ij}^{kl}$$

The maximum of c_{ij} represents the position of best match between F and the sample region




Figure 8. Simulated Cruise Asteroid Image

When the spacecraft nears one of its targets, and the object becomes resolved, and consequently brightens, the exposure times necessary to image the object necessarily decrease. In fact, the opposite problem faced during the cruise imaging must be dealt with, namely the object becoming too bright to easily image in the same picture with dim stars.

Previous deep space missions depending upon Optical Navigation (Principally Voyager and Galileo) have taken advantage of very accurate position determination of extended images of targets, namely images of the major body and its satellites. For weeks or months such images were available, and with the addition of reasonably good physical constants models (e.g. shape and size), extremely good position determination was possible. For these missions, a tenth to a quarter of a pixel was normal, translating in the final approach images to a few tens of kilometers (Ref. 9). For *DSI* this situation is quite different.

The physical nature of the targets (with the possible exception of Phobos) is poorly known. The resultant uncertainty in the modeled figure contributes to a significantly poorer centerfinding. In compensation, the *DSI* targets do not become resolved, and therefore subject to mismodeling errors, until the spacecraft is quite close.

It is guessed that the uncertainty in the diameter of McAuliffe and W-K-I is at least 50 percent, however the uncertainty of the centerfinding process is not nearly this large. The location of the extended images will be determined by a basic brightness centroiding technique. In general, the region in which the body image is located is predictable to within about one hundred pixels before the picture is taken. Within this vicinity, those areas with brightness greater than background will be used to compute a brightness centroid. The centroid is adjusted for the approach phase angle, *via* the relationship given in the equation:

$$X(\alpha) = \frac{3}{16} \pi R \frac{\sin \alpha (1 + \cos \alpha)}{(\pi - \alpha) \cos \alpha + \sin \alpha}$$

where X is the centroid offset, R is the object radius and α is the solar phase angle. If the approach phase angle were zero, the phase deflection term would be zero, and a brightness centroid measurement of the center of brightness would give an arbitrarily good measure of the geometric center of a body modeled as a sphere. For the two encounters to be flown where there is large uncertainty about physical constants, the phase angles are about 50 and 90 degrees. Differentiating this equation with respect to diameter gives the dependence of the phase correction of a diameter error. This relation evaluated for McAuliffe approach and W-K-I, gives a maximum of less than half a radius, which for both objects is well below a kilometer. As a result this error source does not make a dominant contribution to the overall control and knowledge errors of the AutoNav system. Additional error will occur due to shape and albedo irregularities, but it is expected that these errors are at or below the gross size and phase effects.

For the late encounter knowledge update process (discussed below) the image processing procedure must be very fast, one or two CPU seconds. For this purpose, the precision of the brightness centroid is reduced by a simple process of data compression; the image pixels are merely under-sampled. When the body-image is large, and therefore the relative size error as described above is larger, then the inaccuracies of undersampling do not contribute significantly overall to the navigational errors. Fig 9 shows a simulated version of an approach picture to McAuliffe. Images such as these are being used to test the algorithms and the flight software.

Orbit Determination

One important advantage of an all-optical-data orbit determination system is the insensitivity of the data type to high-frequency velocity perturbations. This is especially true for *DS1* which for the first time will employ a low-continuous-thrust propulsion strategy. Such systems are presumed to have significant time-varying thrust characteristics. With a velocity-measuring data type such as Doppler, this propulsion system poses substantial problems. These problems must be dealt with by the radio navigation that will be performed as part of the *DS1* operations and validation task, but they will not have to be addressed by the onboard AutoNav system.

At the core of the Orbit Determination (OD) subsystem is the modeled representation of the spacecraft flightpath. This representation defines the nature and extent of the parameterization and accuracy possible in the system. The Navigator models the spacecraft motion with a numerical n-body integration, using major solar-system bodies as perturbing forces. Non-gravitational perturbations to the spacecraft trajectory included in the model include a simple spherical body solar-pressure model, a scalar parameter describing IPS engine thrust efficiency, and small accelerations in three spacecraft axes. A spherical-body solar-pressure model is sufficient because for the majority of the time, the spacecraft will have its solar panels oriented toward the sun. Even though the spacecraft can maintain this orientation with any orientation of the bus-body about the panel yoke axis, the panel orientation by-far dominates the solar pressure effect.




Figure 9: Simulated Asteroid Approach Frame

During the cruise phase, the optical system is typically capable of taking 250km measurements, depending on the available set of beacon asteroids. Over one week's time, that represents the capability of measuring velocity to about 0.4 m/sec, or accelerations to about 1.3 mm/sec². The IPS

engine is capable of delivering a maximum of about 0.1Nt thrust, but on average will only be capable of half of that during the mission due to power restrictions. *DS1* has a mass of about 420kg, and therefore a typical inflight acceleration is about 120 mm/sec². The IPS engine thrust is believed to be predictable to about one percent, or about 1.2 mm/sec². It is clear then, that long-frequency signatures in the IPS performance will be barely perceptible to the optical system in one week's time. These errors must be modeled. The capability of the Navigator IPS thrust noise model will not nearly meet the requirements of the ground radio navigation system, which has a 0.1 mm/sec velocity sensitivity, and a comparable acceleration sensitivity. However, coping with the noise in the engine performance will still be the single most complicating factor in the flight OD algorithms.

The OD filtering strategy is an epoch-state, batch sequential stochastic filter. With the time-constant of the sensitivity to the expected engine performance errors on the order of a week, data batches of a maximum of a week are used. This is especially sensible since for much of the cruise periods, there will likely be only one OpNav observing period per week. The latter limitation is to reduce the on-off cycling of the engine. The data arc will typically be composed of 4 one-week data batches. The spacecraft state at the beginning of the first batch is the principal estimable parameter. Over each batch a random variation in the thrust magnitude is estimated, as well as small random accelerations. A term proportional to the solar-pressure is also an estimable parameter.




Figure 10. Schematic of Orbit Determination Data-arc Structure

Fig. 10 shows the subdivision of the data arc into batches over which an estimate parameter set is constant. $\mathbf{X}(t_0)$ is the spacecraft state at the start of the data arc, $\mathbf{X}(t_1)$ at the start of the second batch, etc. ρ_n is a scalar parameter describing a proportionality factor on the nominal IPS thrust magnitude in the spacecraft +Z direction. For any observation made at time t within batch one, the filter must integrate the state $\mathbf{X}(t)$, and the state transition matrix. The later has two components, for the state itself: $\partial\mathbf{X}(t)/\partial\mathbf{X}(t_0)$ and for the dynamic force parameters: $\partial\mathbf{X}(t)/\partial\mathbf{X}(\rho_1, \mathbf{S})$ where \mathbf{S} is a vector of other force models, including solar pressure and small bias accelerations active across the data arc; these latter model the small components of the thrust error which project in the cross directions from Z. For this

observation at time t , and for subsequent observations a measurement matrix \mathbf{A} can be formed:

$$\mathbf{A} = \begin{bmatrix} \partial O_1 / \partial \mathbf{q} \\ \partial O_2 / \partial \mathbf{q} \\ \vdots \\ \partial O_n / \partial \mathbf{q} \end{bmatrix}; \quad \frac{\partial O_{2 \times 1}}{\partial \mathbf{q}} = \frac{\partial O_{2 \times 1}}{\partial \mathbf{X}} \cdot \frac{\partial \mathbf{X}}{\partial \mathbf{q}}$$

\mathbf{O}_n is the observation vector for observation n , and is a 2×1 vector, (pixel and line). The formulation of $\partial \mathbf{O} / \partial \mathbf{X}$ is documented elsewhere (Refs. 9,10). \mathbf{q} is a vector of estimable parameters, and for batch 1, $\mathbf{q} = [\mathbf{X}(t_0), \rho_1, \mathbf{S}]$. \mathbf{A} is combined into a covariance matrix referenced to t_0 , Γ_{t_0} , via a **UD** factored orthogonalization procedure (Ref. 11) an example of which is known as the Householder transformation. To process data in batch 2, an additional parameter must be added to the estimate vector, namely ρ_2 the thrust proportionality error for batch 2. Thus for batch 2, $\mathbf{q}_2 = [\mathbf{X}(t_0), \rho_1, \rho_2, \mathbf{S}]$ and the filter will integrate \mathbf{X} from t_1 to t_2 , as well as $\partial \mathbf{X}(t) / \partial \mathbf{X}(t_1)$ and $\partial \mathbf{X}(t) / \partial \mathbf{X}(\rho_2, \mathbf{S})$. The state partials for a time t in batch 2 relative to the solve-for epoch t_0 and those with respect to ρ_2 are given by:

$$\frac{\partial X(t)}{\partial X(t_0)} = \frac{\partial X(t)}{\partial X(t_1)} \cdot \frac{\partial X(t_1)}{\partial X(t_0)},$$

$$\frac{\partial X(t)}{\partial \rho_1} = \frac{\partial X(t_1)}{\partial \rho_1} \cdot \frac{\partial X(t)}{\partial X(t_1)}.$$

And in general, for batch n , where $\mathbf{q}_n = [\mathbf{X}(t_0), \rho_1, \rho_2, \dots, \rho_n, \mathbf{S}]$:

$$\frac{\partial X(t)}{\partial X(t_0)} = \frac{\partial X(t)}{\partial X(t_{n-1})} \cdot \frac{\partial X(t_{n-1})}{\partial X(t_0)}, \text{ and}$$

$$\frac{\partial X(t)}{\partial \rho_m} = \frac{\partial X(t_{n-1})}{\partial \rho_m} \cdot \frac{\partial X(t)}{\partial X(t_{n-1})}.$$

where ρ_m is an arbitrary thrust error vector from an earlier batch. When all of the data from all of the batches is combined into \mathbf{A} and Γ_{t_0} , an estimate of the parameters can be made:

$$\begin{bmatrix} X_{t_0} \\ \bar{\rho} \\ \mathbf{S} \end{bmatrix} = \Gamma_{t_0}^{-1} \mathbf{A}' \mathbf{W} \Delta \mathbf{y},$$

where,

$$\Delta \mathbf{y}_{1 \times 2N} = \mathbf{O}_{2 \times N} - \mathbf{C}_{2 \times N}$$

where $\Delta \mathbf{y}$ is the residual vector formed as the difference between the observation vector \mathbf{O} and the computed predicted value \mathbf{C} . \mathbf{W} is the observation weighting matrix. N is the total number of frames taken, and $2N$ is the number of data (pixel and line for each). Iterations are performed on this solution, repeating the solution one or more times with the improved integrated ephemeris and force models from the previous solution. When the solution is converged, the elements of ρ are not equally well determined; ρ_1 is the best determined from a covariance standpoint, as all of the data in the data arc influence a measurement of ρ_1 , whereas ρ_n is the poorest, as only the last batch has an influence on its solution. To get the covariance to start the next solution cycle the covariance at t_0 must be mapped forward in time:

$$\Gamma_{t_{n/2}} = \mathbf{D} \Phi_{t_0}^{t_{n/2}} \Gamma_{t_0} \Phi_{t_0}^{t_{n/2}'}.$$

where $\Phi(t_0, t_{n/2})$ is the state transition matrix from t_0 to the midpoint of the data arc. \mathbf{D} is a de weighting matrix to allow for errors accrued due to unmodeled perturbations.

The decision has been made to entirely reinitialize the solution process for each data arc. Operationally, this process typically has the following events:

- ◆ A solution is performed for a four batch data-arc spanning typically 28 days, with an epoch-state at the beginning of the first batch. This solution uses effectively no *a priori* constraint, relying on the data arc for virtually the complete state determination.
- 2) Data is accumulated beyond the last batch, into what is the “new” batch.
- 3) The estimated state from step 1 is integrated to the beginning of the second batch. This integrated state becomes the reference or epoch-state for the next solution.
- 4) A solution is made using the data in the new batch, but excluding the old (original “first”) batch. The process repeats starting at step 1.

In this approach, the rationale for completely redetermining the state using the data arc only, without any pre-constraint, or forwarding of information from previous solutions is two-fold. First, there is sufficient information in a month’s worth of optical data (four typical batches) to sufficiently determine the position and velocity of the spacecraft. Second, the earlier data (earlier than about a month) are sufficiently decoupled from the current data arc *via* the random non-gravitational accelerations so as to contribute little or no information to the solution.

Integration and Ephemeris Services

The characteristics of the spacecraft dynamic models are discussed above, but the actual mechanism used to perform the integration is a separate issue, as is the representation of the spacecraft integrated trajectory, and the ephemerides of

the major and minor planets, including the encounter targets.

The numerical integrator used is a Runge-Kutta 8th-order. This integration algorithm, while not computationally the most efficient available, represents the best compromise between speed and accuracy (Ref. 12). The heritage of the algorithms chosen to be incorporated into the flight Navigator was an important aspect of that decision. The coded version of the RK-8 actually used has a history of use in diverse orbital applications of more than twenty years. This integrator has a manually set maximum and minimum integration step size, and automatically ranges between them based on the current level of dynamic perturbation. The accuracy achieved when operating under flight conditions, is several tens of meters over a seven-month ballistic cruise, with full dynamic perturbations in force. This comparison is against the JPL Orbit Determination Program (ODP) principal integration routine (Ref. 13) which sets the standard for deep space navigation accuracy. The RK-8 subroutine will be used to integrate the spacecraft position and the partial derivative equations for purposes of state and parameter estimation.

As stated earlier, *DS1* is a complex mission from the standpoint of expected dynamic perturbations. In order for the trajectory integrator to provide sufficient accuracy to the system, information about actual onboard propulsive activity is provided to the Navigator. This information comes from two sources, the IPS manager and the ACS. From the IPS device-manager comes a constant tally of accumulated thrust time and thrust level. By monitoring voltages and currents in the ion engine, the IPS manager is able to compute an estimated thrust magnitude. Over a span of about a minute, the IPS manager tallies this thrust, and then reports to the Navigator the accumulated thrust and time since the last message. This process continues whenever the IPS is in operation and thrusting.

The ACS also reports all propulsive activity to the Navigator, in a somewhat different manner. The ACS is constantly inducing propulsive events, but of varying magnitude compared to the IPS. In the maintenance of the spacecraft attitude, the ACS is inducing small limit-cycling turns with a frequency of roughly ten seconds when doing precision imaging (e.g. navigation observations) or tens of minutes during ballistic cruise. Additionally, ACS is responsible for implementing TCM's. These can implement several m/sec of velocity change in a matter of minutes. Every turn of the spacecraft is a propulsive event, since only in one axis (the roll -Z- axis) are the thrusters balanced, and each turn can impart roughly a mm/sec of velocity to the spacecraft. Attitude maintenance maneuvers will approximately average to zero delta-v, due to their short extent; asymmetries in the thruster performance will not however, nor will large turns. Even a few mm/sec when

accumulated and mapped over a one month-long data arc is many kilometers of spacecraft displacement. This is very observable to the Navigator, and therefore must be tallied. During all periods of operation therefore, the ACS Velocity Estimator is monitoring ACS activity and computing accumulated velocity. When an accumulation of more than a mm/sec is achieved in any of the three inertial directions, a report is sent to the Navigator. If some fixed time, (usually 10 minutes) passes without the minimum accumulation, a report is sent nevertheless. The Navigator accumulates both types of information, and condenses it into a record of propulsive activity over the past. This record is kept for approximately five weeks, more than enough to cover the past integration history over the longest expected data arc. The trajectory integrator then reads this record to integrate an accurate propulsive history from the epoch-state to the end of the data arc.

The planet, asteroid and spacecraft ephemerides are represented as Chebyshev function polynomials of varying order. This follows the standard representation of the planetary ephemerides in the ground navigation software. The accuracy of the stored planetary and asteroid ephemerides (relative to their generating values) is .01km, using a 10-30 coefficient model, effective over about 5 days. The spacecraft ephemeris, with a similar representation accuracy, uses 25 coefficient representation over 1-2 day intervals.

IPS Control, Maneuver and TCM Design

Perhaps the most crucial function of the Navigator is the control of the IPS. A deep space mission has never been flown whose trajectory was not composed of long ballistic cruise segments, punctuated by planetary gravitational assists and virtually instantaneous velocity changes. This, the first deep space low thrust mission, compounds the challenge, by requiring control of the ion engine to be performed autonomously.

The design of a low-thrust mission is a specialized technology of its own (Ref. 13), independent of the navigation function. And clearly this design process proceeds well in advance of the stage of the mission requiring autonomous navigation. The results of the design are provided to the Navigator in the form of a time-history of thrust level and direction (Figs. 3-5). The form of storage onboard of the direction profiles is by first order polynomial in time, with each week having a separate set of coefficients. The thrust levels are stored as discrete integer levels for each week.

As will be discussed below, during typical cruise operations, the Navigator will be called upon to perform weekly determinations of the thrust profile. Part of this evaluation will be to use the current best estimated state to determine what changes to the upcoming week's thrust profile are

necessary to return the spacecraft to an intersecting trajectory with the target. As discussed earlier, the changes that are possible to the designed mission trajectory are limited, due to constraints of spacecraft body orientation. Also, there is limited time to implement the mission thrust-arcs, and the existing design already uses most of the time available on the first leg, to McAuliffe. Therefore, the corrections that are possible are constrained, and represent relatively small and linear (or nearly so) corrections to the nominal designed mission.

The strategy to be used for updating the thrust profile is to treat one or more of the upcoming weekly thrust periods as an individual maneuver. Corrections to the nominal thrust polynomial can be considered the parameters of a maneuver to be estimated. Details of the algorithm used to accomplish these corrections are recorded elsewhere (Ref. 14). Briefly, it is based on a linear estimate of control parameters, \mathbf{s} which have varying dimension, depending on the number of adjacent control segments being adjusted. A trajectory miss vector $\Delta\mathbf{X}$ is computed in the 3-dimensional encounter asymptotic coordinates. The parameters \mathbf{s} are small changes in direction in each segment, and a change in duration of the overall burn arc. In order to obtain the solution that minimizes the corrections to the nominal thrust arc, the minimum-normal solution for \mathbf{s} , is formed *via* the equations:

$$\Delta\mathbf{s} = \mathbf{K}' (\mathbf{K}\mathbf{K}')^{-1} \Delta\mathbf{X},$$

where,

$$\mathbf{K} = \begin{bmatrix} \frac{\partial X}{\partial \alpha_1} & \frac{\partial X}{\partial \delta_1} & \frac{\partial X}{\partial \alpha_2} & \frac{\partial X}{\partial \delta_2} & \cdots & \frac{\partial X}{\partial \alpha_m} & \frac{\partial X}{\partial \delta_m} & \frac{\partial X}{\partial \tau} \end{bmatrix},$$

$$\Delta\mathbf{s}' = [\Delta\alpha_1 \quad \Delta\delta_1 \quad \Delta\alpha_2 \quad \Delta\delta_2 \quad \cdots \quad \Delta\alpha_m \quad \Delta\delta_m \quad \Delta\tau],$$

and,

$$\Delta\mathbf{X}' = [\Delta\mathbf{B} \cdot \mathbf{R} \quad \Delta\mathbf{B} \cdot \mathbf{T} \quad \Delta t_{of}]$$

$\Delta[\mathbf{B} \cdot \mathbf{R}, \mathbf{B} \cdot \mathbf{T}, t_{of}]$ are the target relative asymptotic coordinates, representing two cross-track directions, and the along-track direction at closest approach. The solve-for parameters, $\Delta\alpha_n$, $\Delta\delta_n$, and $\Delta\tau$ are changes in a series of n thrust segment directions, and the end time of the final thrust arc. This solution is performed iteratively until converged. In this way, the solution process is actually a non-linear one, but will only succeed if a solution exists near the linear region.

As the IPS thrust arc progresses, and variations in engine performance and minor (or major) outages in thrust time relative to the nominal plan occur, the spacecraft trajectory will deviate from the designed-to nominal trajectory. The

targeting strategy outlined above will return the spacecraft to the specified target conditions, but in so doing, will alter the velocity vector of the encounter asymptote. Enough of a change in this vector could cause a potential problem in maintaining the next legs of the mission to potentially Mars and WKI. If it is determined that sufficient changes to the asymptote have occurred, the trajectory will be reoptimized on the ground, and the corresponding thrust profiles will be uplinked to the spacecraft. With a redesigned mission will be a new projected mass-usage profile, associated with propellant consumption. The accuracy of this profile will effect the dynamics of the onboard integration, and therefore will be uplinked with the thrust profile.

During periods of non-thrusting, and in the twenty days before encounter conventional TCM's will be performed. These will use the IPS with the exception of the final 2 maneuvers, which will be executed using the hydrazine thrusters of the ACS. Table 5 shows the TCM schedule, with expected and associated OD errors mapped to encounter at each TCM for the final 20 days of approach to McAuliffe. The algorithm used to compute these maneuvers is the same as used for the IPS control algorithm. Necessarily however, the maneuver solution is for only three parameters: the three components of delta-velocity. Another important difference between a RCS TCM and an IPS control, is that the former occurs in a relatively short period of time; whereas IPS controls can take hours or days. In most cases the applied maneuvers are expected to be small, on the order of one m/s or less, which for the IPS will take less than two hours.

Time to Encounter	Range to McAuliffe (km)	Downtrack Error (km)	Crosstrack Error (km)
-20d	12.6E6	570	660
-10d	6.3E6	138	27.3
-5d	3.1E6	69	5.5
-2.5d	1.6E6	54	2.5
-1.5d	0.9E6	44	1.5
-1.0d	630E3	42	1.2
-12h	315E3	40.2	0.89
-6h	157E3	40.1	0.55
-3h	72E3	40.1	0.50

Table 5: Approach TCM Schedule with Associated OD Performance Statistics

The nature of the bus-body illumination constraints has been discussed earlier, as has the need to constrain the direction of TCMs accordingly. The need to perform maneuvers in any direction of the sky persists however, as statistical variations in the orbit determination process do not observe the constraints of onboard instruments. Any direction of propulsive maneuver (using either RCS or IPS) can be accomplished by vectorally splitting the maneuver into two parts, whose vector sum equals the original design (Fig. 11). Through interaction with ACS, the Navigator determines if a particular maneuver request is allowed, and if not,

decomposes the TCM into two parts. The precise nature of the interaction necessary to accomplish this will be discussed below.




Figure 11. TCM Vector-decomposition

There is substantial uncertainty about the size of asteroid McAuliffe (even more about comet WKI), and complete lack of information regarding the shape of this asteroid, and its rotational axes. As a result, the desire to fly past this target at a small integer multiple of nominal radii presents a small but still substantial risk to the spacecraft. To cope with this safety issue, the nominal aim point will be 10km from the asteroid surface. From about 6 hours to 3 hours before closest approach, the Navigator will make determinations of the McAuliffe's size. The process used will be a combination of simple triangulation and area analysis. If, in this 3 hour period, there is no indication of an anomalously large size, an E-3 hour "Bold-Encounter" Deflection maneuver will be performed, to take the spacecraft in to the very near aimpoint. Along with this maneuver, the spacecraft will be directed to use a somewhat different encounter sequence (discussed below) to correspond to those conditions.

Late Knowledge Update

The final control of the spacecraft trajectory will occur at about 6 hours prior to encounter. Subsequent to that maneuver, the full navigation picture processing and OD estimation process will be in force. But at approximately 30 minutes from closest approach, normal navigation operations will cease. Because of the very short timescale of activities at encounter, the Navigator must initiate simplified processes. The principal technical feature that enables the simplified processes is the fact that for the final few minutes of the approach, the Navigator can acquire no additional useful information about the velocity of the spacecraft. This being the case, the data filter reduces dramatically to a 3-state estimate of instantaneous spacecraft position only. The estimates occur from picture to picture, and each solution is conditioned by the covariance obtained from the previous picture. Over so short a time-span, the absence of any process noise, or other attenuation of the accumulating information does not cause a substantial error due to mismodeling. This is due to the rapidly increasing power of the data as the spacecraft approaches; any modeling errors in previous images would be overwhelmed

by the increased power of the later pictures. The picture processing used during this final stage of the approach has been discussed above.

OPERATION OF THE NAVIGATION SYSTEM:

The operation of the Navigator, though largely an autonomous function, is managed in a gross sense by ground commands. These commands are imbedded in a conventional stored sequence. Typically, a ground directive is given to the Navigator, followed by a period of uncommitted time in which the Navigator is allowed to perform autonomous action. Following are detailed descriptions of the major Navigation actions.

Navigation Imaging Opportunity

The simplest period of activity during the mission is ballistic cruise (non-powered cruise). During this period of time, the only regular navigation operations that occur are the taking and processing of navigation frames. Such an event is triggered by a Nav-Photo-Op spacecraft command. Though this operation happens during all phases of the mission, it will be discussed here in the context of a non-thrusting (ballistic) portion of the trajectory. For most of the mission, this operation will occur once per week. At one point in the sequence, a Nav-Photo-Op directive is issued to the Navigator by the ground-generated stored sequence. Associated with this command, is a period of time allocated to the Nav function to accomplish picture planning, execution and processing. Even though the Photo-Op opportunity is triggered by a ground command, very little planning is required on the ground, other than the specification of the length of the opportunity window.

Before the Photo-Op session begins, it is the ground system's responsibility to put the spacecraft in a state that is possible to command turning and imaging operations. This preparation activity includes turning the camera on, and changing whatever camera states are necessary, and doing so with sufficient lead time to insure readiness when the Photo-Op begins. If any ACS states need setting, this must also be done. Additionally, the ground must insure that no operations occur which conflict with imaging and turning commands during the extent of the Photo-Op.

Very little information is necessary to pass to the AutoNav system with this directive, but it is necessary to inform Nav how much time is available to obtain its images. When the "Nav-Photo-Op" directive is issued, the following operations take place:

- 1) Nav determines what the current attitude of the spacecraft body is, in order to be able to return to that attitude after imaging if requested. Otherwise, ground operations can specify a different terminal attitude.
- 2) AutoNav identifies the set of navigational targets that are appropriate for the current time of the mission.

- 3) A target is selected, in order, from the list starting at the beginning of this period. Each of the lists has been optimized so as to minimize the extent of the turns between targets.
- 4) Nav determines from ACS how long a turn from the current attitude to the requested attitude will take. Additionally, The ACS planning expert is asked how long it will take to turn from the target attitude to the *a priori* attitude. If the sum of these is less than the time remaining in the AutoNav session, then the sequence of operations continues, other wise a branch to the end procedure (step 10) commences.
- 5) AutoNav prepares a small file onboard which contains a "mini-sequence." This sequence requests ACS to turn to the specified target
- 6) AutoNav launches the ACS-turn mini-sequence, using one of the eight available sequence strings.
- 7) AutoNav waits for a "Turn Complete" message.
- 8) On receipt of the "Turn Complete" message, AutoNav builds and launches a mini-sequence to take the MICAS image, with automatic notification of "Image Complete" being sent to AutoNav.
- 9) With receipt of the "Image Complete" notification from the launched sequence, the main Photo-Op events continue, with a branch back to event 3) and a selection of the next target in the list.
- 10) Begin the termination process for the Photo-Op, with the construction of a minisequence to turn the spacecraft back to the starting or other requested attitude.
- 11) Launch of the final turn mini-sequence, and this marks the end of Photo-Op.

IPS Control:

During the months of continuous thrusting, there are periods of time when the IPS must be shut down for short periods. These interruptions include time for navigation data taking, for downlink of data, and possibly for technology validation experiments. Also, on a regular basis, perhaps once per day, the direction of the engine thrust must be updated by the AutoNav system.

As with the Nav-Photo-Op directive, use of the commands to enable the AutoNav system to operate the IPS, require the ground operating system to prepare the spacecraft for the autonomous operation of the navigation system. In the case of a "NAV-SET-IPS" command, the ground generated sequence turns on and otherwise conditions the IPS engine. From a cold start, there is a considerable amount of preparation necessary, taking up to an hour. However, since these activities are well known, repetitive, and well calibrated in terms of time required, the mission operations team uses a fixed sequence, called a "block" and as part of normal invocation of the Navigator, this will be routinely done.

To begin autonomous IPS operations then, the ground first issues the "IPS-PREPARATION" block command leaving the ion engine in a state ready for the AutoNav system to issue a simple "thrust-on" command. Then, after leaving sufficient time in the sequence to complete the preparation cycle, the sequence issues a "NAV-SET-IPS" command. In

response to this command, the AutoNav system begins a series of tasks:

- 1) A computation is made of the necessary thrusting over the next day. The direction of engine is determined, as is the duration of the burn.
- 2) The ACS planning expert (APE) is queried to determine the length of time required to turn the spacecraft to the desired position.
- 3) A mini-sequence is constructed to accomplish several tasks:
 - Turn the spacecraft to the desired direction
 - A delay necessary to guarantee completion of the turn.
 - A directive to the IPS manager to turn on the thrust grids of the ion engine, and to leave the thrust on for a maximum of 1 day, or for a shorter duration if specified.
- 4) The mini-sequence is launched.

The duration specified for each IPS SET or UPDATE command is the duration of the mission thrust arc, which can be several months. This is clearly longer than the time-span to the next SET or UPDATE command, at which time the duration will be reset to a span reflecting recent IPS activity. To accomplish the necessary updates to the thrust vector, the ground-generated sequence will include periodic requests of AutoNav to update the direction. Although it would be possible for the AutoNav system to autonomously provide update vectors, in order to do so, AutoNav would have to become aware of other scheduled events on board the spacecraft which would cause a change in the status of the engine, such as telecommunication events. Since it causes little impact on the ground system to issue the NAV-UPDATE-IPS command, AutoNav will rely on this method. On receipt of this command, the Navigator will construct and launch a new minisequence to update the thrust direction and duration. These directives will go to the ACS attitude commander and IPS manager respectively.

At the end of a mission-thrust segment, the navigator will, in response to an UPDATE command, issue a directive to the IPS manager with a thrust duration of less than the expected time to the next SET or UPDATE command. The IPS manager will keep track of the amount of time that the IPS has been thrusting since a SET or UPDATE directive, and if this duration is met, the manager will shut down the IPS.

As stated earlier, the timings of events that shut down the IPS, such as navigation picture taking and telecom sessions is not known *a priori* onboard by the Navigator, being carefully scheduled by the ground. Therefore, the AutoNav system must cope with the otherwise unscheduled shut-down of the engines at any time. This is accomplished *via* the design of the IPS control software, involving continued monitoring of the accumulated thrust from the engine. At any time, the Navigator is prepared to evaluate the thrust

accumulated thus far, and to thereby reevaluate the necessary duration of thrust given to the IPS manager in a command. Therefore, the ground control system may shut down the engines at any time, and the Navigator will adjust to the circumstance.

Such a shutdown is simply implemented. The ground-generated sequence commands the thrust to turn off, then commands the engine to whatever shut-down state is required. The Navigator is made aware of the shutdown implicitly via the lack of "engine-on" status messages from the IPS manager.

Trajectory Correction Maneuvers:

With conventionally navigated spacecraft, the implementation of a TCM required a major effort for the ground control team. With the AutoNav system, ground control is relieved of all responsibility for the TCMs except for scheduling. Much as with the OpNav image taking, the ground merely schedules a time-gap in the sequence in which the AutoNav system may place its autonomous operations. In this case, the operations are to turn the spacecraft and operate the engines: either the RCS thrusters or the IPS.

During an extended mission-thrust period, no dedicated TCM's are necessary, as continuous corrected control is taking place. However, after a mission burn, during a ballistic cruise, and especially on approach to an encounter target, dedicated opportunities to correct the trajectory are required. These can be scheduled frequently with no additional ground costs. For *DS1*, it is anticipated that the spacecraft travel no more than a month between TCM opportunities, and that they occur much more frequently on approach to a target, as has been discussed earlier.

The ground implementation of a TCM is as follows. Prior to issuing any command to the Navigator, ground operations must insure the readiness of the RCS system or the IPS (or both), depending on which is to be mandated to be used, or if the navigator will be given the option of using either. Such preparations might include turning on the IPS, or activating the TCM RCS thruster heaters. When the preparations are complete the ground-generated sequence issues a NAV-PERFORM-TCM command. This begins a series of activities:

- 1) The Navigator will refer to an orbit determination calculation (recently performed in response to a stored-sequence directive) based on the latest data, to determine the current spacecraft state and its propagation to the encounter target.
- 2) The velocity change necessary to take the spacecraft to the target is computed.
- 3) The ACS vectorizer is queried as to whether this TCM needs vectorization, and if so, what are the components into which it can be broken down. (Fig. 11).

- 4) The APE is consulted as to the extent of time required to implement the turn(s).
- 5) The Navigator constructs a mini-sequence to accomplish a series of tasks:
 - A: Direct ACS to turn the spacecraft to the requested attitude,
 - B: Wait the required amount of time to implement the turn,
 - C: Direct ACS to implement the delta-v.
 - D: If an unvectorized turn, proceed to E, otherwise, complete steps A through C for the second leg of the TCM,
 - E: Direct ACS to turn back to the *a priori* attitude, or a requested terminal attitude.
- 6) The Navigator then starts the mini-sequence, to accomplish the above activities, and this completes the implementation of a TCM.

These activities are constrained to take place in a given amount of time. This constraint is enforced by two methods, first by a hard limit in the total length of time provided in the sequence. If the Navigator hits this limit in constructing its mini-sequence, this constitutes an error. To prevent this error from occurring, the Navigator is initially constrained from implementing TCMs of greater than a certain magnitude. The magnitude of this limit will correspond to a 3-sigma maximum expectation value of statistical delta-v. If this limit is surpassed, the Navigator will implement the maximum magnitude in the computed direction. The allocated sequence time will correspond with this expected maximum time with some additional appropriate buffer.

Orbit Determination

In response to a NAV-DO-OD command, the navigator will take a number of important actions:

- 1) Update the data arc to a pre-specified length (usually 28 days) deleting older data from the data file.
- 2) Update the estimable epoch-state, to be positioned at the beginning of the newly truncated data-arc.
- 3) Perform orbit determination on the edited data arc, computing a new epoch-state estimate.
- 4) If control opportunities exist in the next planning segment (usually 7 days, but getting progressively shorter on approach to encounter) compute the retargeting parameters for this control. These parameters will be used in response to IPS control or TCM commands to the Navigator.
- 5) Write a spacecraft ephemeris file based on the new estimates and controls for use by the NAV-RT ephemeris server.

Though for *DS1* operations, NAV-DO-OD will be a ground-sequence issued command, this need not be so. This command could as easily be issued by the Navigator as a self-induced command. This mode of operation was decided against for various non-navigational reasons.

Encounter Operations:

The activities of the *DS1* encounter will be determined well in advance of the encounter itself. These operations will be

encoded into a series of sequences stored onboard the spacecraft, and triggered into operation by the Navigator.

At least for the McAuliffe encounter, the dependence of the scheduled sequences upon the high accuracy knowledge of the location of the spacecraft relative to the target does not become strong until the last five minutes of approach. The important dimensional dependence is upon the down-track dimension, as this direction remains poorly determined until very late. Consequently, the final approach sequence is subdivided into 4 short sub-sequences, each with increasing sensitivity to time-of-flight (down track position) errors, and each positionable with greater accuracy by the Navigator.

For the approximately five hours following the final TCM, prior to the start of the McAuliffe-Encounter operations, images are being taken by the spacecraft and passed to the Navigator for processing. Throughout this “Far Encounter” period, the Navigator is updating its estimate of the spacecraft encounter coordinates, including the time of closest approach (TCA.) Since the timing of these events is not dependent upon an accurate determination of TCA, these can be scheduled in the sequence in a completely deterministic way.

The first of the asteroid encounter sequences (AE1) begins 260 seconds before closest approach at a range of about 2000 km. The first action of this sequence is to take an OPNAV image, at E-240sec. This image is immediately sent to AutoNav for processing. As the science activities of the encounter sequence proceed, the AutoNav system is reducing the data and obtaining a new encounter state estimate. The science activities of AE1 will include infrared and ultra-violet observations of McAuliffe. Since the combined processes of data readout, image analysis, and state estimation take approximately 12 to 15 seconds, there is time in AE1 for the Navigator to process several pictures if the science sequence allows. Each update of the target-relative ephemeris is automatically reflected in improved pointing accuracy. This is so because the ACS system is regularly querying the Nav system for the latest ephemeris information. All science observations are specified as target relative (vs. absolute inertial directions) and thus are improved in accuracy whenever the Navigator improves the accuracy of the ephemeris. It should be emphasized again however, that once the sequence is started, the time of a specified event is deterministic and cannot change. AE1 will end at E-175sec.

The second encounter sequence (AE2) will begin at about 160 seconds before closest approach. As with AE1, the first action of the sequence will be to take an OPNAV image, in this case, at about E-155 seconds. There is a gap of about 15 seconds between AE1 and AE2 which will allow the Navigator to move the start point of AE2 to correspond to updated estimates of the time of closest approach. As with

AE1, there will be opportunities for multiple OpNav pictures to be taken and processed, and the estimated spacecraft ephemeris updated before the end of AE2 at E-90 seconds.

The third encounter sequence (AE3) will begin at E-80 seconds, and as previously, the first activity is to take an OPNAV image at E-75 seconds. Additional OPNAV images may be taken in AE3 using the other visual frequency imaging system, the APS (Active Pixel Sensor), before the sequence ends at E-40 seconds.

The final encounter sequence (AE4) begins at E-35 seconds. The final OPNAV image is taken with the CCD sensor at E-33 seconds, and the final target-relative ephemeris is made available to ACS at about E-23 seconds. From this time until the spacecraft can no longer accelerate its slew-rate to keep the target tracked, at about E-15 seconds, science images with the APS and CCD will be taken. Even when this limit is reached, several images may still be taken over the next few seconds, as the asteroid (then over three CCD fields of view in apparent diameter) sweeps out of view. AE4 will continue taking IR images of the asteroid as it sweeps out of view, and turn the spacecraft to view the retreating asteroid on departure. This turn should be complete within about a minute, whereupon science imaging (but no OPNAV imaging) will continue, until AE4 ends at E+240seconds.

The above sequence describes the activities for the 10km flyby. As discussed earlier, if the Navigator senses that McAuliffe is of nominal size, a “Bold-Encounter” deflection maneuver will take place at E-3hr to send the spacecraft to a 5km above the surface flyby. In this case, the Navigator will direct a somewhat different AE4 sequence in which the last OpNav image will likely be at E-20sec, and the final science image at E-7sec, with a range of about 50km.

Following AE4, conventional deterministic sequencing will resume, with final science views of the asteroid. Within five days or so, AutoNav operations will also resume, with periodic beacon-asteroid images, and autonomous control of the IPS.

PRELIMINARY SIMULATION RESULTS

Although the development of the navigation flight system is not yet complete, preliminary simulations have been run with the software to assess its performance. This simulation uses the current baseline trajectory obtained from mission design, which assumes a launch on July 1, 1998 and flyby of the asteroid McAuliffe on January 16, 1999. Covariance analysis was performed on the last 30 days of this cruise prior to asteroid encounter to determine OD performance in both an interplanetary cruise and small body flyby scenario. The analysis assumes no a-priori knowledge on the state at

the E-30 day epoch. Data scheduling during this time frame is shown in Table 6. Note that up to around E-12 hours, observations are taken of multiple beacon asteroids to fix the heliocentric spacecraft trajectory. Subsequent observations up to the encounter are solely of the target asteroid to accurately determine the target-relative spacecraft state, in particular, the time-of-flight or downtrack component.

The resulting performance is graphically displayed in Figures 12 and 13. These show the semimajor and semiminor axes of the 3-dimensional positional uncertainty ellipse mapped to the encounter as a function of time. Figure 12 shows the dramatic improvement in position knowledge in all three dimensions gained from the data from E-30 to about E-7 days. The largest dimension of the ellipse has a value of about 70-80 km at this time, and represents the best knowledge of the downtrack uncertainty of the spacecraft position relative to the target obtainable from the beacon and target asteroids. The two other dimensions of the ellipse however, have about the same values and are an order of magnitude better than the largest component. This is due to excellent cross-track information obtained from observing the target asteroid with optical data. By the time of encounter, these components will be known to the 100-200 m level.

Figure 13 shows an expanded view of the last hour prior to encounter. Note that the semimajor axis of the uncertainty ellipse (representing the downtrack error) which had not shown much improvement from E-7 days has a sudden dramatic drop at about E-1 hour. This is caused by the changing geometry as the spacecraft flies by the asteroid. The cross line-of-sight measure of the spacecraft position relative to the target is rotated into the downtrack direction, thereby improving the estimate of this component. This clearly illustrates the need for late observations of the target, and why it would be impractical to process this important data on the ground due to light-time considerations. Only by processing this information onboard can the improved knowledge from late observations be taken advantage of for science purposes.

Table 6: Observation Scheduling for 30 Days Prior to Asteroid Encounter		
Time to Encounter (days)	# of observations	IAU Catalog # of asteroids used
29	13	5,15,46,126,132,163,183,270,313,398,696,1036,3352
22	13	5,15,46,126,132,163,183,270,313,398,696,1036,3352
15	12	5,15,126,132,163,180,183,270,313,398,1036,3352
13	13	5,15,126,132,163,180,183,270,313,398,1036,3352
10	12	5,15,126,132,163,180,183,270,313,398,1036,3352

8	12	5,126,132,163,180,183,270,313,347,398,1036,3352
6	12	5,126,132,163,180,183,270,313,347,398,1036,3352
4	12	5,126,132,163,180,183,270,313,347,398,1036,3352
3	12	5,126,132,163,180,183,270,313,347,398,1036,3352
2	12	5,126,132,163,180,183,270,313,347,398,1036,3352
1	12	5,126,132,163,180,183,270,313,347,398,1036,3352
0.4	12	5,126,132,163,180,183,270,313,347,398,1036,3352
0.4 - 0.0	39	3352




Figure 12: Autonomous Navigation System Orbit Determination Performance, Far Encounter




Figure 13: Autonomous Navigation System Orbit Determination Performance, Near Encounter

FIRST USE OF THE NAVIGATOR

Though in theory the *DSI* Navigator could be run with no ground interaction, in order to provide a well documented validation of its first use, to provide some optimization of the navigation function, and to allow for sensible safety margins, much ground analysis will be taking place during operations. The MICAS images will be extensively analyzed to provide calibrations of the camera itself, and of the pointing accuracy of the ACS system. This information will be relayed back to the Navigator in form of improved camera models. It has been mentioned that the Navigator's maneuver estimator is reasonably robust to deviations in the planned thrust schedule. But such deviations might induce fuel-costly changes to one or more encounter geometries if left uncorrected. For this and other reasons, periodic opportunities to re-optimize the mission trajectory and thrust-arcs will be present. Finally, the need to carefully gauge the performance of both the Navigator and the IPS engine requires a comprehensive and unprecedented ground radio navigation campaign (Ref. 15). The extent of this ground analysis though providing a large measure of confidence and safety for *DSI* operations, does imply that the cost savings of navigational autonomy will not be seen on *DSI*. Once demonstrated however, this technology will provide future projects with capable and economical systems with which to navigate difficult but rewarding planetary missions.

ACKNOWLEDGMENT:

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Special thanks are due to a number of individuals at JPL: Marc Rayman, the *DSI* Mission Engineer, under whose auspices much of the *DSI* Navigation development has occurred, Abdullah Aljabri who lead the *DSI* Flight Software Team during its most important formative stages, Steve Williams and Dennis Byrnes for leading the mission search and design process and Ted Drain for consultation and assistance in implementation details of the flight system, especially the ephemeris server subsystem.

REFERENCES:

- 1) Riedel, J.E., Bhaskaran, S., Synnott, S. P., Bollman, W. E., Null, G. W.. "An Autonomous Optical Navigation and Control System for Interplanetary Exploration Missions", IAA paper IAA-L-0506, Proceedings, Second IAA International Conference on Low-Cost Planetary Missions, April 16-19, 1996 Laurel, Maryland, USA.
- 2) Rayman, M. D., Lehman, D. H., "NASA's First New Millennium Deep Space Technology Validation Flight," IAA paper IAA-L-0502, Proceedings, Second IAA International Conference on Low-Cost Planetary Missions, April 16-19, 1996 Laurel, Maryland, USA
- 3) Coverstone-Carroll V., Williams, S. N., "Optimal Low Thrust Trajectories Using Differential Inclusion Concepts", AAS/AIAA Spaceflight Mechanics Meeting, Cocoa Beach Florida, USA, Feb. 1994
- 4) Lisman, S., Chang, D, Hadaeagh, F., "Autonomous Guidance and Control for the *New Millennium DSI* Spacecraft" AIAA paper 96-3817, July 29-31 1996, San Diego, Ca., USA.
- 5) Stolper, S. A., "Embedded Systems on Mars: The Mars Pathfinder Attitude and Information Management (AIM) Subsystem", Proceedings, Embedded Systems, Conference East, Vol. II, Miller Freeman, Inc., San Francisco, 1997.
- 6) Hog, E.; "The TYCHO Project on-board the HIPPARCOS Satellite," IAU Symposium No. 141, "Inertial Coordinate System on the Sky", pg. 307-310
- 7) Bhaskaran, S., Riedel, J. E., and Synnott, S. P., "Autonomous Optical Navigation for Interplanetary Missions", Space Sciencecraft Control and Tracking in the New Millennium, E. Kane Casani, Mark A. Vander Does, Editors, Proc. SPIE pp. 32-43 (1996).
- 8) Vaughan, R. M., Riedel, J. E., Davis, R. P., Owen, W. M., Synnott, S. P., "Optical Navigation for the Galileo Gaspra Encounter," AIAA paper 92-4522, AIAA/AAS Astrodynamics Conference, Aug. 1992, Hilton Head, SC.
- 9) Riedel, J. E., Owen, W. M., Stuve, J. A., Synnott, S. P., Vaughan, R. M., "Optical Navigation During the Voyager Neptune Encounter", AIAA paper 90-2877-CP, AIAA/AAS Astrodynamics Conference, Aug 1990, Portland, OR.
- 10) Synnott, S. P., Donegan, A. J., Riedel, J. E., Stuve, J. A. "Interplanetary Optical Navigation: Voyager Uranus Encounter," AIAA paper 86-2113, AIAA conference, Aug. 1986, Williamsburg Va.
- 11) Bierman, G. J., Factorization Methods for Discrete Sequential Estimation (Academic Press, New York), 1976
- 12) Battin, R.; "Resolution of Runge-Kutta-Nystrom Condition Equations through Eighth Order", AIAA Journal, Vol. 14, August 1976, pp 1012-1021.
- 13) Ekelund, J; "The JPL Orbit Determination Software System", AAS paper 79-111, AAS/AIAA Astrodynamics Specialist Conference, Provincetown, Mass, USA, June 25-27, 1979
- 14) Desai, S.; Bhaskaran, S.; Bollman, W.; Halsell, C.; Riedel, S.; "The Autonomous Navigation System: Autonomous Control of Low Thrust Propulsion Systems", AIAA paper 97-3819, AIAA Guidance, Navigation and Control Conference, Aug. 1997, New Orleans, USA.
- 15) Wood, L. J., Hast, S. L., "Navigation System Design for a Halley Flyby/Tempel 2 Rendezvous Mission Using Ion Drive", Paper 79-110, AAS/AIAA Astrodynamics Specialist Conference, Provincetown, Mass., June 25-27 1979.

Autonomous Optical Navigation
(AutoNav)
DS1 Technology Validation Final Report

Appendix D
Orbit Determination Performance
Evaluation of Deep Space 1
Autonomous Navigation System

Orbit Determination Performance Evaluation of the Deep Space 1 Autonomous Navigation System

S. Bhaskaran, S. D. Desai, P. J. Dumont,
B. M. Kennedy, G. W. Null, W. M. Owen Jr.,
J. E. Riedel, S. P. Synnott, R. A. Werner

Navigation and Flight Mechanics Section
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Abstract

NASA's New Millennium Program consists of a series of missions whose primary purpose is to demonstrate the feasibility of new technologies for spaceflight. Deep Space 1, the first mission in the New Millennium Program, will demonstrate an Ion Propulsion System to provide thrust and an autonomous onboard navigation system to guide the spacecraft. The mission plan is to fly by an asteroid, Mars, and a comet using these and other new technologies.

The onboard navigation system, in order to be as self-contained as possible, uses images of asteroids taken by the spacecraft's camera as its sole data type in determining the spacecraft's trajectory. These images are clustered at intervals varying from hours to a week depending on the phase of the mission, with up to 12 different asteroids sighted per cluster. The images are then incorporated into a least-squares filter at periodic intervals to estimate spacecraft orbit parameters. The orbit determination solutions are in turn used by the navigation system to compute maneuvers required to guide the spacecraft to its targets. Since this navigation strategy has never before been used in flight, it is important to perform pre-launch assessments of its performance. This is accomplished by the use of Monte Carlo simulations which drive the navigation software with a truth model of the spacecraft trajectory and the observables. The truth model simulates realistic errors which are expected in flight, and individual realizations of these errors are drawn from random samplings of the errors with provided statistics. This technique is used to analyze the first leg of the mission, the flyby of the asteroid McAuliffe. The results indicate that, under nominal conditions, the combined orbit determination/maneuver computation strategy is capable of navigating the spacecraft to a safe flyby. In addition, the propulsive events required are within the abilities of the hardware.

INTRODUCTION

Standard navigation techniques for interplanetary spacecraft involve the use of a combination of radio (two-way coherent Doppler and ranging) data, obtained by tracking the spacecraft using antennas at JPL's Deep Space Network (DSN) tracking stations, augmented by optical data from an onboard camera during encounters. This combination of data is very accurate and has been used successfully to navigate spacecraft to all planets in the solar system except Pluto, and to three asteroids. However, in order to fully realize NASA's vision of the future of deep-space exploration, with multiple small, inexpensive spacecraft roaming the solar system, it is desirable to automate some or all of the processes required for interplanetary missions, including navigation. It is possible

to fully automate the navigation process by eliminating the radio data and using an onboard camera to triangulate the spacecraft's position by observing multiple solar system bodies. In this system, the data would be processed by an onboard filter to obtain the complete spacecraft ephemeris, from which maneuvers could be planned and performed to achieve the desired targeting conditions. Such a system is being developed for JPL's Deep Space 1 (DS-1) asteroid/comet flyby mission, the first in the New Millennium Program series of missions. The purpose of this paper is to analyze the performance of the orbit determination (OD) and maneuver targeting links of the DS-1 autonomous navigation system. Specifically, the ability of the system to deliver the spacecraft to its first target is assessed.

THE MISSION

The New Millennium Program is a recent program instituted by NASA with the primary purpose of demonstrating new technologies for future space missions. Its ambitious goal is to fly a series of missions, both Earth orbiting and interplanetary, each testing technologies which have not been proven in flight conditions and which have dramatic potential of enabling missions which could not be flown previously or of lowering the cost of space flight. The hope is that the missions will prove these technologies so that future science oriented missions can use them without incurring the cost or risk of flying a new technology. More information on the New Millennium Program can be found on its web site at (<http://nmp.jpl.nasa.gov>).

DS-1 is the first of the interplanetary missions of the New Millennium Program. In addition to autonomous navigation, other primary technologies being demonstrated include the first use of an ion propulsion system for trajectory control, an advanced solar array for power, and low-mass imaging system named MICAS (Miniature Integrated Camera and Spectrometer) (see Ref. 1 for a more detailed description of all the technologies to be validated, and Ref. 2 for an overview of all aspects of autonomous navigation). The mission itself will be launched onboard a Delta 7326 rocket between July 1 and July 31, 1998, perform a close (less than 20 km) flyby of the asteroid 3352 McAuliffe on January 20, 1999, receive a gravity assist from the planet Mars on April 2000, and then finally rendezvous with comet West-Kohoutek-Ikemura (W-K-I) in early June 2000 at a distance of about 500 km. The main science return will come from high resolution imaging of the asteroid and comet during their respective flybys using the MICAS camera.

ION PROPULSION SYSTEM

Perhaps the most important aspect of the DS-1 mission in terms of its impact on navigation is the use of an Ion Propulsion System (IPS) engine. Unlike chemical propulsion systems which burn for short periods of time at very high thrust, the IPS produces very little thrust but is capable of burning for very long periods of time. Ionized xenon is accelerated by passing it through a charged grid before exiting out of the nozzle. The resulting thrust is on the order of millinewtons, with specific impulses reaching values in the thousands of seconds (as compared to 200-400 seconds for chemical rockets). The thrust can be throttled by varying the voltage on the grids; for DS-1, the IPS has about 100 throttle levels, with a thrust range of 20 to 90 mN. Since the power is generated from the solar arrays, the maximum achievable thrust depends on the distance to the sun.

The characteristics of an IPS trajectory are different from those using chemical engines. Trajectories using chemical engines have long coast periods punctuated by near-instantaneous velocity changes at given times to achieve course corrections. IPS trajectories, on the other hand, are characterized by long thrusting periods of weeks to months, interspersed with coast arcs when the IPS is shut off. For DS-1, the thrusting periods have the dual purpose of providing enough energy to the spacecraft to reach its targets, and correcting launch injection, OD, and maneuver execution errors to achieve the desired targeting conditions. More details on the latter will be described below.

Designing the low-thrust reference trajectory for DS-1 is a complicated process. Briefly, the first step is to compute an optimal trajectory which takes the spacecraft from its launch injection

conditions to the targets. The trajectory is optimized by finding the set of control parameters (the right ascension α and declination δ of the thrust pointing vector, and the duration of thrusting) which achieves the targets with a minimum amount of fuel usage. Since this process is dependent on several factors, including the launch date and available power from the solar array, the nominal trajectory is constantly being revised as new data (especially about solar array performance) is received. To analyze the OD performance, we used a single reference trajectory whose characteristics should not deviate greatly from the final one flown. For the current design of the mission trajectory, the nominal IPS thrusting period begins on July 16 (15 days after launch), and ends on September 4, 1998. Prior to this period, the IPS will be used primarily for calibrating engines during its initial checkout phase. After this period, the nominal thrusting phase, or mission burn period, is over, and the IPS will only be turned on for trajectory correction maneuvers (TCMs).

If the launch injection were perfect and the IPS thrusting in exactly the designed direction and magnitude, then the mission burn would be sufficient to achieve the targets and no TCMs would be needed. In reality, of course, errors in these and other factors cause trajectory deviations, and corrections are necessary. Thus, the onboard navigation system will be used to periodically check the position and velocity of the spacecraft and correct the thrust parameters as needed. This is accomplished in the following manner. At seven-day intervals during cruise, the IPS is shut down for a period of about 12-16 hours while the spacecraft slews to take sightings of up to 12 asteroids (each of these thrust/shutdown segments is referred to as a planning cycle). These observations are used to compute an OD solution to get the current spacecraft state. This state is mapped forward to the next encounter, and if the deviation from the desired encounter condition is large enough, a linearized course correction consisting of adjustments to the α and δ of the thrust vector during subsequent planning cycles, and the duration of the final mission burn segment, is computed. After the mission burn is over, the OD solution at the end of each planning cycle will be used to support TCM opportunities every few weeks. These TCMs will consist of a single IPS burn at a computed direction and duration. In the final 30 days prior to asteroid encounter, the planning cycles will have shorter durations of variable length, and the final 4 TCMs will be performed using the hydrazine based reaction control system (RCS) thrusters. These thrusters are normally used for attitude control, but due to the short time remaining before encounter, it was decided that IPS burns may require too much time to implement. Table 1 lists the times and types of maneuver opportunities for this reference trajectory. Note that both the IPS and RCS TCMs come in pairs several hours apart. This is to allow for vectorization of the maneuver, whereby if a computed thrust vector is in a direction which violates a spacecraft attitude constraint, it is broken into two segments in allowable directions whose vector sum is equal to the original. A complete description of the linear correction strategy used to correct the mission burns and compute TCMs is given in Ref. 3. Assuming that the IPS performs reasonably close to its specifications, the linear correction strategy will suffice. However, if there are very large deviations in the IPS performance from its design, or if frequent outages occur during mission burns, a redesign of the reference trajectory will be done on the ground and uplinked to the spacecraft.

ORBIT DETERMINATION

Orbit determination is the process by which the spacecraft's state (position and velocity) and other parameters relevant to the trajectory, such as nongravitational accelerations acting on the spacecraft, are estimated. In order to keep this process as self-contained onboard the spacecraft as possible, the only data used to obtain an OD solution are images taken of solar system bodies (asteroids in this case) by the MICAS camera. In principle, the procedure to obtain a simple position fix of the spacecraft in heliocentric space using asteroid sightings is extraordinarily simple. A single sighting of an asteroid places the spacecraft along the line of sight (LOS) to that asteroid. Observing a second asteroid at the same time will deterministically fix the three-dimensional heliocentric position of the spacecraft, provided the ephemerides of the sighted asteroids and the inertial pointing direction of the camera are known. In practice, however, two simultaneous sightings are not practical with one camera, and instead, a series of LOS fixes are taken of several asteroids. For DS-1, the number

Table 1: Maneuver Schedule for Nominal DS-1 Trajectory

Maneuver ID	Maneuver Type	Date	Time to Asteroid Encounter
0	Mission Burn	July 16, 1998 15:00:00	188 days
1	Mission Burn	July 23, 1998 15:00:00	181 days
2	Mission Burn	August 1, 1998 15:00:00	172 days
3	Mission Burn	August 8, 1998 15:00:00	165 days
4	Mission Burn	August 15, 1998 15:00:00	158 days
5	Mission Burn	August 22, 1998 15:00:00	151 days
6	Mission Burn	August 29, 1998 15:00:00	144 days
7	Mission Burn	September 5, 1998 15:00:00	137 days
8	Mission Burn	September 12, 1998 15:00:00	130 days
9	IPS TCM	September 19, 1998 15:00:00	123 days
10	IPS TCM	September 19, 1998 22:00:00	123 days
13	IPS TCM	October 10, 1998 15:00:00	102 days
14	IPS TCM	October 10, 1998 22:00:00	102 days
18	IPS TCM	November 7, 1998 15:00:00	74 days
19	IPS TCM	November 7, 1998 22:00:00	74 days
23	IPS TCM	December 5, 1998 15:00:00	46 days
24	IPS TCM	December 5, 1998 22:00:00	46 days
28	IPS TCM	December 31, 1998 20:53:46	20 days
29	IPS TCM	January 1, 1999 03:53:46	19 days, 17 hours
31	IPS TCM	January 1, 1999 03:53:46	19 days
31	IPS TCM	January 10, 1999 20:53:46	10 days
32	IPS TCM	January 11, 1999 03:53:46	9 days, 17 hours
33	IPS TCM	January 15, 1999 20:53:46	5 days
34	IPS TCM	January 16, 1999 03:53:46	4 days, 17 hours
35	IPS TCM	January 18, 1999 20:53:46	2 days
36	IPS TCM	January 19, 1999 03:53:46	1 days, 17 hours
37	RCS TCM	January 19, 1999 20:53:46	1 day
38	RCS TCM	January 19, 1999 21:13:46	1 days, 23 hours, 40 minutes
39	RCS TCM	January 20, 1999 08:53:46	12 hours
40	RCS TCM	January 20, 1999 09:13:46	11 hours, 40 minutes
41	RCS TCM	January 20, 1999 14:53:46	6 hours
42	RCS TCM	January 20, 1999 15:13:46	5 hours, 40 minutes
43	RCS TCM	January 20, 1999 17:53:46	3 hours
44	RCS TCM	January 20, 1999 18:13:46	2 hours, 40 minutes

of sightings taken during a given observation window of opportunity is limited by the amount of time it takes to slew the spacecraft from one asteroid to another; an upper limit of 12 is anticipated. Several clusters of sightings are then incorporated into a least-squares filter to obtain an OD solution. The accuracy of this type of data is dependent on several factors, including the angular separation, brightness, and distance to the imaged asteroids, the resolution of the camera, the ability to pinpoint the location of the asteroid in the camera frame (centerfinding), the accuracy of the camera pointing information, and the knowledge of the asteroid ephemerides. These factors will be addressed in the following sections. For clarity, the term “beacons” is used to denote the asteroids used solely for triangulation, while “target” refers to the objects being encountered (asteroid McAuliffe and comet W-K-1 for DS-1).

The Camera System

The MICAS camera system actually has two imaging devices, one a standard charge-coupled device (CCD), and the other an experimental active pixel sensor (APS) array. Of these, it is anticipated that the autonomous navigation (autonav) system will primarily use the CCD because of its larger field-of-view (FOV). Use of the APS by the autonav system will be limited to the final 30 minutes prior to encounter when the CCD image will be oversaturated. Both are connected to a 677 mm focal length telescope. The CCD has a 1024×1024 pixel array, giving a total FOV of 0.8° , or about 14 mrad. Each pixel therefore has an angular resolution of $13 \mu\text{rad}$.

Image Processing

The image processing link forms the core of the autonav system. Its primary purpose is to predict the locations of beacons at given times, determine the center of the asteroid in the camera frame, and compute the associated pointing of the camera boresight. The ability of the navigation system to perform autonomously hinges on its ability to accurately perform the centerfinding and ensuring that bad data do not corrupt the solution.

Computing predicts of beacon asteroids is the simplest of these procedures. A list of beacon asteroids to observe as a function of time for the entire mission is stored onboard the spacecraft, along with ephemerides of all the beacons (more will be said about the choice of beacons later). At predetermined times, the current spacecraft trajectory is differenced with the nominal ephemeris of given beacon to get the relative pointing vector. This information is then passed to the spacecraft attitude control system (ACS) which slews the spacecraft to the correct orientation at the correct time and shutters the picture with the provided exposure length.

Because of its importance, the centerfinding algorithms (and the associated pointing solution) used during cruise when asteroids are distant point sources have had the most testing. The details of these procedures have been documented in Refs. 4 and 5; only a brief description will be given here. The algorithms are a modification of similar ones used for the Galileo mission, both onboard the spacecraft and on the ground. They use a pattern matching technique to filter out unwanted bright spots and locate the asteroid and known stars in the camera FOV. From experimental results (see Refs. 4 and 5), the algorithms are capable of determining the location of the asteroid relative to the stars to a precision of 0.1 pixels.

For computing the pointing direction of the camera boresight, an initial guess of the values are needed. This is provided by the ACS system, which uses a wide FOV star tracker for attitude knowledge and control. The accuracy of the pointing available from ACS is about 0.3 mrad prior to alignment calibrations, and 0.1 mrad after. If at least three stars are visible in the CCD image, however, the pointing information can be improved by computing a least-squares fit to the pointing (α and δ of the boresight, and the twist around the boresight) using the ACS values as an initial guess. Assuming 0.1 pixel centerfinding ability, the pointing can be determined to a few μrad . If fewer than three stars are available, then the accuracy is degraded. Analysis has shown that three or more stars will be available during cruise. Encounter navigation requires new data types because the extended target body is very bright (usually about magnitude 2-3 per pixel) and because very

near encounter the target image fills the camera field of view. When stars are in the field but the contrast between the stars and target body exceeds the camera dynamic range, then “flash-mode” observations are made by alternating short exposures of the target body and longer exposures to bring up stars; camera pointing is determined from the star exposures and interpolated to the time of the target body exposure. When the range to the target is sufficiently small (an hour from closest approach for McAuliffe), then “starless” observations of the target are processed using the camera pointing values obtained from the star tracker.

The star catalog used by autonav contains 221,594 stars that lie within 30° of the ecliptic and have a catalogued visual magnitude of 10.50 or brighter. The positional data for the stars are taken from the highly accurate catalogs produced by the European Space Agency’s Hipparcos satellite.

For purposes of evaluating the OD, an observation uncertainty, σ_o , of 0.1 pixel was used for the beacon observations, which represents the current best estimate of the centerfinding accuracy for distant, unresolved asteroids. As the spacecraft nears encounter, however, the target asteroid becomes resolved and the pattern matching centerfinding algorithm cannot be used. Instead, a simple brightness centroiding on the asteroid is done. Because the asteroid has an unknown shape, this method can only determine the brightness center, and the true center is unknown. The error is potentially as large as the radius of the asteroid, so, the data are deweighted accordingly. The uncertainty used is the angular extent of the body in the camera FOV, converted to pixels:

$$\sigma_o = \frac{\tan^{-1}(R/\rho)}{13 \times 10^{-6}}, \quad (1)$$

where

- R = the assumed radius of the asteroid,
- ρ = the range to the asteroid.

Asteroid Ephemerides

An implicit assumption in the use of triangulation asteroids for orbit determination is that the heliocentric positions of the asteroids at the time of the observation is known exactly. In fact, this is not really the case; the orbits of the 5,000 or so numbered asteroids are known to different accuracies. The larger and/or brighter objects which have been tracked for longer periods of time have orbital accuracies in the tens of km, while the smaller and dimmer objects which have not been observed as much are known to within only several hundreds of km.

To properly account for the ephemeris errors, the orbits of the asteroids used for triangulation would have to be estimated along with the spacecraft trajectory in the OD filter. However, this would greatly increase the complexity of the filter since there are over 80 beacons. To keep the onboard OD algorithm simple, therefore, asteroid ephemeris errors are ignored. Instead, by using up to 12 asteroids per observation set, we rely on simple averaging to remove these errors during cruise.

Encounter presents a special case. For a given camera and centerfinding ability, the accuracy of an observation is directly proportional to the distance to the asteroid. During encounter, the target is several orders of magnitude closer than the beacons so the power of its observations overwhelms the information provided by the beacons. The result that the spacecraft’s target-relative state is accurate to the level of the data, but its heliocentric state estimate is skewed by an amount roughly equivalent to the ephemeris error present in the target’s orbit. This is an acceptable consequence though, since it is the the target-relative, not the heliocentric, state which is important for targeting and visually tracking the object during the flyby.

In order to minimize the adverse effects of ignoring the asteroid ephemerides, a ground campaign is underway to improve the orbits of some asteroids. About 80 asteroids have been identified as probable beacons for the current DS-1 trajectory; 30 of these are being observed from JPL’s Table Mountain Observatory with the expectation that their orbits can be improved by a factor of 3 or

4. Of particular importance are the flyby targets, McAuliffe and W-K-I. The current prediction is that, assuming the observations are successful, McAuliffe's orbit uncertainty can be improved from its current value of 127 km, 50 km, and 60 km in the radial, transverse, and normal directions, respectively, by about a factor of 3.

Beacon Asteroid Selection

One non-autonomous portion of the navigation function is the selection of the beacons used for triangulation. This procedure, referred to as the picture planner, is done on the ground and the results stored onboard before launch. The picture planner propagates the spacecraft state and asteroid states using either conic elements or numerical integration. For each planned weekly triangulation session, it searches for acceptable observing opportunities by examining observation characteristics for the lowest-numbered 5000 asteroids and selecting the subset of asteroids which produce the best combined accuracy in the local instantaneous spacecraft state determination. These computations take into account camera sensitivity, full well, system noise, and dynamic range. Observation geometry conditions constraining beacon selection include beacon brightness, beacon distance, solar phase angle, spacecraft pointing constraints, camera measurement accuracy, star background (at least two suitably bright stars are required), and star-relative smear of the beacon during the computed exposure time (the cross-correlation can tolerate only 1-2 pixels of star-relative smear). Closer asteroids provide better observation accuracy provided that the star-relative smear is acceptably small. Attitude control performance parameters such as absolute pointing accuracy (about 0.1 of the CCD field) and expected limit cycle "kick velocity" (about 3 pixels/sec) are also used in the picture planning computations. Camera exposure time and pointing can be adjusted to provide the best astrometric measurement accuracy for each observation opportunity. For each selected asteroid the output includes observation epoch, asteroid identification, exposure time, and the few-hour effective span for which the prediction is valid. The trajectory file for the beacon asteroids will typically contain 100-200 asteroids. For encounter, the picture planner output is referenced to the encounter time and the onboard navigator then updates the absolute observation times using its latest encounter time determination.

Dynamical Equations and Filtering

In general, the process of determining orbital state parameters of an interplanetary spacecraft is a nonlinear one. However, the process can be considerably simplified by linearizing the problem, which amounts to solving for deviations of the orbit parameters about a reference trajectory rather than the orbit parameters themselves. This allows powerful methods of linear estimation theory to be applied, resulting in more stable solutions. This does require, though, that initial guesses to the state parameters be available to generate the reference orbit.

The second-order equations of motion used to generate a reference trajectory can be written as two first-order equations:

$$\dot{\mathbf{r}} = \mathbf{v} \quad (2)$$

$$\dot{\mathbf{v}} = -\frac{\mu_s}{r^3}\mathbf{r} + \sum_{i=1}^{n_p} \mu_i \left[\frac{\mathbf{r}_{ri}}{r_{ri}^3} - \frac{\mathbf{r}_{pi}}{r_{pi}^3} \right] + \frac{AG}{mr^3}\mathbf{r} + \frac{k}{m}\mathbf{T} + \mathbf{a}, \quad (3)$$

where

- \mathbf{r} = the heliocentric cartesian position vector of the spacecraft,
- \mathbf{v} = the heliocentric cartesian velocity vector of the spacecraft,
- \mathbf{r}_{pi} = the heliocentric cartesian position of the *i*th perturbing planetary body
- \mathbf{r}_{ri} = the position of the spacecraft relative to the *i*th perturbing body, i.e., $\mathbf{r}_{ri} = \mathbf{r}_{pi} - \mathbf{r}$
- μ_s = the gravitational constant, GM , of the sun,

- μ_i = the gravitational constant of the i th perturbing planet,
- n_p = the number of perturbing planets,
- A = the cross-sectional area of the spacecraft,
- G = the solar flux constant,
- \mathbf{T} = the thrust vector from the IPS, in newtons,
- k = the thrust scale factor, with value approximately 1,
- m = the spacecraft mass, and
- \mathbf{a} = miscellaneous accelerations acting on the spacecraft.

In Eq. 3, the first term on the right hand side represents the central body gravitational acceleration from the sun, the second term is the sum of the third body gravitational acceleration from the planets (all except Pluto are used), the third is the solar radiation pressure, the fourth is the acceleration due to thrusting from the IPS, and the final term accounts for any additional unmodeled accelerations acting on the spacecraft.

The two gravitational acceleration contributions are straightforward, but the non-gravitational forces acting on the spacecraft deserve some discussion. With regard to solar radiation pressure, it is obvious from Eq. 3 that a simple spherical model for the spacecraft was used. In reality, the spacecraft's cross-sectional area is dominated by the two solar array panels, with the spacecraft bus contributing a much smaller proportion. During flight, the panels will almost always be pointed at the sun, with the bus rotating to provide thrust vector control, camera pointing, etc. Since the dominant effect is from the panels which remain more or less fixed relative to the sun, it was decided that the complexity of using a more accurate model was not needed.

Thrusting events on the spacecraft come from two sources: the IPS for mission burns and TCMs, and the RCS for attitude control and late TCMs. IPS events are explicitly accounted for in the filter via the fourth term in Eq. 3, but are handled differently depending on whether the integration is performed from a past time to the present for OD purposes (the data arc), or for predicting the state of the spacecraft at some future time (predicts). For the former, the actual thrust achieved by the IPS is not measured directly (such as with an accelerometer), but is instead indirectly computed based on measured voltages across the ion acceleration grid. At preset intervals varying from seconds to minutes, the voltage is read out for computing the magnitude of the thrust, and the spacecraft attitude at the corresponding time is also obtained from the ACS to get the thrust direction. This information is passed to the navigation system which accumulates the high rate data and, when a certain threshold in either the thrust magnitude or change in direction is reached, prints a record to a history file containing an averaged thrust magnitude and direction over that time span. This averaging minimizes the storage required to maintain history information over a long data arc. Since thrust is not directly measured, the value of thrust computed will have some uncertainty associated with it. The characteristics of the measurement error is somewhat uncertain at this time, but is expected to be within $\pm 1.5\%$ of the true value. The scale factor k in the fourth term of Eq. 3 is used to account for this measurement error and will be an estimated parameter in the filter.

RCS thrusters are used primarily for attitude control, but will also be used for TCMs near encounter. Once again, the way they are handled in the integration depends on whether the integration is over a past time or for predicts. For history information, onboard ACS software sends out thruster activity reports in terms of the velocity change, or $\Delta \mathbf{V}$, accumulated over a time span, with the minimum time span presently set to 1 second. The navigation software receives these high rate messages and compresses the data by waiting until a minimum $\Delta \mathbf{V}$ threshold is reached, after which a record of the total $\Delta \mathbf{V}$ vector at that time is written to the same history file which stores the IPS activity. Additional records are also written if a time threshold is passed, so that small $\Delta \mathbf{V}$ s which do not reach the threshold will be properly time tagged. Finally, prior to obtaining an OD solution, all remaining $\Delta \mathbf{V}$ s which have not reached either the magnitude or time threshold are written to the history file. RCS activity for attitude control and TCMs are handled in this manner without any distinction being made between the two. Unlike the IPS, though, each individual RCS event is not modeled explicitly in the filter. Instead, the fifth term in Eq. 3., the general acceleration term,

is used to estimate the averaged acceleration errors over a given span caused by mismodeling of the ΔV s caused by RCS firings.

At the time when an OD solution is needed, the integration over the data arc proceeds as follows. The history file is sorted so that the IPS and RCS propulsive activity records are time ordered. Starting from the beginning of the data arc, the integrator will proceed through the time span, stopping and restarting at each thruster event. For IPS events, the acceleration contribution of the thruster is interpolated by computing the thrust magnitude and direction over the time span in which it is active, then dividing the thrust magnitude by the current spacecraft mass to get acceleration. For the RCS, the instantaneous ΔV contribution at a given time is added to the current velocity vector, and the new state is propagated forward. Although this method of stopping and starting the integrator at thrust discontinuities is time consuming, the accuracy gained is substantial and necessary to prevent filter divergence.

For predicts, the thrust value during IPS mission burns will be computed as a function of the available power from the solar arrays, which in turn is a function of the distance to the sun. During IPS TCMs, the thrust is nominally zero but will be adjusted by the maneuver software for retargeting. The adjusted value and associated duration are written to a maneuver file; this file is read by the integration routine to obtain the appropriate thrust information. The scale factor k for IPS predicts will always take a value of 1. RCS TCMs also are nominally zero, adjusted during retargeting, and written on the maneuver file. As with the history integration, the integrator for predicts will stop at RCS events on the maneuver file, add the instantaneous computed ΔV , and restart the integration. The accuracy of this method remains only if the RCS ΔV s are small (on the order of m/s) and therefore take only a short time for the thrusters to achieve; large RCS ΔV s would incur an integration error penalty. Current analysis indicates that RCS TCMs should indeed be fairly small, so this is not currently a cause for concern. Finally, attitude control events in the future are not predictable and presumably average to zero over the course of the mission. For this reason, they are not modeled and the general acceleration term in Eq. 3 is ignored for predicts.

Filter

Once the reference trajectory for the data arc is generated, the solution of the state parameters, which are corrections to the nominal values used to generate the reference, can be obtained using the techniques of epoch state batch filtering from linear estimation theory [Ref. 6]. If we define the adjustable parameters of the nominal trajectory, $\mathbf{q}^*(t)$, as

$$\mathbf{q}^*(t) = [X^*(t) \ Y^*(t) \ Z^*(t) \ \dot{X}^*(t) \ \dot{Y}^*(t) \ \dot{Z}^*(t) \ k_1^* \dots k_n^* \ a_x^* \ a_y^* \ a_z^*]^T, \quad (4)$$

where

- X^*, Y^*, Z^* = the cartesian position components,
- $\dot{X}^*, \dot{Y}^*, \dot{Z}^*$ = the cartesian velocity components,
- $k_1^* \dots k_n^*$ = thrust scale factors, with a different scale factor estimated for each planning cycle in a data arc, and
- a_x^*, a_y^*, a_z^* = the components of the general acceleration vector,

then the updated trajectory, $\mathbf{q}'(t)$, is

$$\mathbf{q}'(t) = \mathbf{q}^*(t) + \Delta \mathbf{q}(t), \quad (5)$$

where $\Delta \mathbf{q}(t)$ is the vector of estimated corrections (henceforth, the Δ will be eliminated in the notation for the correction vector $\Delta \mathbf{q}$). If the nominal values are reasonably close to the truth, then the corrections should be linear over the batch time span, and the corrections at the epoch state, $\mathbf{q}(t_0)$, can be linearly mapped to any other time t using the state transition matrix, Φ , as

$$\mathbf{q}(t) = \Phi(t)\mathbf{q}(t_0), \quad (6)$$

where

$$\Phi(t) = \frac{\partial \mathbf{q}^*(t)}{\partial \mathbf{q}^*(t_0)}. \quad (7)$$

To get the state transition matrix values at a given time t , note that

$$\dot{\Phi} = \frac{\partial \dot{\mathbf{q}}^*(t)}{\partial \mathbf{q}^*(t)} \frac{\partial \mathbf{q}^*(t)}{\partial \mathbf{q}^*(t_0)} = \mathbf{A} \Phi(t). \quad (8)$$

This matrix differential equation represents a set of $(9 + N_{pc})^2$ scalar first-order equations, where N_{pc} is the number of planning cycles in a data arc. The initial condition is $\Phi(t_0) = \mathbf{I}$, the identity matrix. The partial derivative matrix \mathbf{A} is computed analytically. Many of the elements of \mathbf{A} are zero, so that only $7N_{pc} + 63$ equations are needed. These equations are integrated along with the nominal trajectory to get Φ as a function of time.

To set up the equations for the epoch state batch filter, the partial derivatives of the observations with respect to the state are needed. The observables in this case are the pixel p and line l coordinates of the beacon or target asteroid centers obtained using the centerfinding techniques described earlier. Thus, at the time of the observation, the partials matrix \mathbf{H} is

$$\mathbf{H} = \begin{bmatrix} \partial p / \partial X & \partial p / \partial Y & \partial p / \partial Z & \mathbf{0}_{2 \times (6 + N_{pc})} \\ \partial l / \partial X & \partial l / \partial Y & \partial l / \partial Z & \mathbf{0}_{2 \times (6 + N_{pc})} \end{bmatrix} \quad (9)$$

The observed (p, l) depends only on the spacecraft's position relative to the beacons at the instant the image is taken; hence the partials with respect to the velocity and acceleration components are zero. The non-zero values of \mathbf{H} can be computed analytically, and the equations for these partials are given in Ref. 7. To map these partials back to the epoch, the state transition matrix is used:

$$\tilde{\mathbf{H}} = \mathbf{H} \Phi, \quad (10)$$

where $\tilde{\mathbf{H}}$ is the observation partial matrix at epoch. Given the *a priori* covariance matrix, \mathbf{P}_0 , the observation weighting matrix, \mathbf{W} (a diagonal matrix whose elements are $1/\sigma_o^2$, with σ_o being the observation uncertainties from Eq. 1), and a residual vector, \mathbf{Y} , which are the differences between of the observed centroid values and the predicted ones computed from the nominal spacecraft trajectory, the original epoch state batch filter equations for the solution vector $\hat{\mathbf{q}}$ and the formal covariance \mathbf{P} are:

$$\hat{\mathbf{q}} = [\mathbf{P}_0 + \tilde{\mathbf{H}}^T \mathbf{W} \tilde{\mathbf{H}}]^{-1} \tilde{\mathbf{H}}^T \mathbf{W} \mathbf{Y} \quad (11)$$

and

$$\mathbf{P} = [\mathbf{P}_0 + \tilde{\mathbf{H}}^T \mathbf{W} \tilde{\mathbf{H}}]^{-1}. \quad (12)$$

In practice, however, the equivalent *U-D* factorized method is used [Ref. 8]. In this method, which was adopted to minimize round-off error and ensure stability, \mathbf{P} is expressed as the product $\mathbf{U} \mathbf{D} \mathbf{U}^T$, where \mathbf{U} is upper triangular with ones on the diagonal and \mathbf{D} is diagonal.

After an initial testing phase, the OD solution strategy to be adopted is as follows. After the first 28 days of cruise during which autonav is enabled, an OD solution is performed. Nominally, this means that four planning cycles are incorporated with 12 observations in each planning cycle, resulting in 48 observations. The *a priori* covariance matrix, \mathbf{P}_0 , for the solution is set such that the position and velocity components are effectively unconstrained, with values of 10^8 km and 100 m/s used for the 1σ uncertainties in position and velocity, respectively. The nominal scenario calls for thrusting during this period, so four thrust scale factors, corresponding to each of the 7-day planning cycles, are also estimated, with *a priori* uncertainties for each set to 5%. Finally, the *a priori* sigmas on the components of the general acceleration term are set to 3×10^{-9} km/s. These values allow the filter to freely adjust the spacecraft's initial position and velocity while constraining the thrust and accelerations to be within reasonable bounds.

Following this initial solution, solutions are performed at 7-day intervals during cruise by dropping the data from the earliest planning cycle and adding the data from the planning cycle just completed. Thus, the OD is performed over a sliding window of a constant 28-day length. The number of




Figure 1: 1σ Uncertainties in Position and Velocity vs Data Arc Length

planning cycles in this window will vary during encounter when planning cycle lengths shorten to less than 7 days, but the total number of days is always kept constant (the amount of data in each of these 28-day batches will also vary during encounter). The same values for \mathbf{P}_0 are used every time a solution is done, so effectively, each batch solution has no “memory” of a previous solution, and data information older than 28 days is lost. The nominal starting trajectory to which corrections are made, however, is the latest in that the starting values for position and velocity at a given batch epoch are the mapped values from the previous OD solution.

The rationale for using this solution strategy can be seen from the plot in Fig. 1. On the figure, the mean position and velocity formal sigmas are plotted as a function of batch length. It can be seen that the uncertainties make noticeable improvements when data from 14, 21, and 28 day batches are used, but they quickly level off afterwards. This is due to the relatively large non-gravitational accelerations acting on the spacecraft, primarily from the IPS thrusters. The noise in these accelerations hinders the mapping of information from one time to the next so that after 28 days or so, the data add little information to solve for the epoch state. For this reason, 28 days was chosen to be the batch length, providing enough information to obtain a reasonable solution but not cluttering the filter with useless data. However, by using the sliding batch window approach and updating the nominal trajectory at each OD solution, the nominal trajectory is implicitly computed with information older than 28 days, after the first OD solution.

The formal uncertainty plots in Fig. 1 show that, for a typical cruise data arc, the filter can determine the spacecraft position to about 130 km in position and 0.7 m/s in velocity. It is also instructive to see how well the filter can estimate the thrust scale factors. Fig. 2 plots the formal uncertainties in the estimate of four thrust scale factors in a typical 28 day arc. In this run, the *a priori* uncertainty on the scale factor was set to 5%. It is clear from Fig. 2 that the first scale factor is poorly determined, with no improvement from the *a priori*, while the fourth is best determined (to about 2%— a little less than half of the *a priori*). In general, the later scale factor estimates will be better than the earlier, although for this particular thrust profile, the second scale factor is better determined than the third due to the orientation of the thrust vectors. The reason for the first scale factor being so poorly determined is twofold: first, the first set of observation data is taken 7 days after the epoch, and second, the epoch state is unconstrained. This results in all errors being absorbed by the state, with nothing attributed to the thrust. Conversely, the fourth planning




Figure 2: 1σ Uncertainties in Thrust Scale Factor Estimates

cycle has data at both ends thereby tightly constraining the position, with the result that remaining errors have to be absorbed by the scale factors. Even then, the improvement is not dramatic, so some care must be taken in interpreting the values of the scale factors estimated by the filter.

The complete set of dynamics and filter will be used to obtain OD solutions throughout the cruise and up to 30 minutes prior to the nominal encounter time. After this, it is expected that the processing time required with the onboard computer resources is not sufficient to permit rapid turnaround of the OD result to update pointing predicts during the flyby. For this reason, the target observations taken after Encounter (E) – 30 minutes will be brightness centroided and passed to a fast, compact 3-state filter (named the Reduced State Encounter Navigation filter, or RSEN). A version of the RSEN filter has already been developed for a similar flyby of a comet for the STARDUST mission, and a description of the algorithm and its performance is given in Ref. 9. The observations are used to update the target relative position only; the target relative velocity has been well determined at this point. The initial state for RSEN is provided at the E–30 minute point from the main navigation module. RSEN is then used primarily to maintain visual lock on the asteroid during the period surrounding closest approach; it will not be used to support further TCMs.

MONTE CARLO SIMULATION AND RESULTS

If the dynamic equations used in the filter accurately modeled the true forces acting on the spacecraft and the errors in the observations were also correctly represented, then the formal covariance obtained after filtering would accurately represent the statistics of the estimated values. This is clearly not the case however, as we have deliberately simplified the nongravitational acceleration terms and ignored some of the errors which affect the data. For this reason, Monte Carlo simulations are needed to assess the true filter performance and compare it with the formal uncertainties. For the simulations, a “truth” model of the trajectory and observations are generated and provided to the filter. For a given run, the truth model represents a single realization from a random sampling of the errors which affect that model. One hundred runs are performed, and the results evaluated by computing statistics on the difference between the known truth and the estimated values computed by the filter. The details of this process will now be described.

Trajectory Model

The trajectory model used for the truth integration is the same as in Eq. 3, with a couple of additions. These modifications are used to simulate errors in the true thrust output by the engines, and to simulate errors in the measurement of the thrust provided to the autonav system. These simulated differences are modeled as sinusoid functions of time. If T and T' represent the commanded and true thrust magnitudes, respectively, then the magnitude for the thrust term in Eq. 3 used for the truth integration is

$$T' = T + A_T \cos[2\pi(t - t_0)/\tau_T], \quad (13)$$

where

- A_T = the amplitude of the additional thrust magnitude,
- t, t_0 = the current and epoch time, and
- τ_T = the time constant of the magnitude variation.

The right ascension α and declination δ of the truth thrust vector are similarly modeled.

The measured thrust magnitudes and directions are taken as a variations on the true values, *i.e.*,

$$T^* = T' + B_T \cos[2\pi(t - t_0)/\tau_T] \quad (14)$$

and similarly for α^* and δ^* . These measured quantities are passed to the OD filter. The measured thrust vector is broken into time ordered segments and sent to the autonav routines to be placed into the history file. Thus, the information used to integrate the trajectory in the filter is different from the truth integration used to generate observables. This best mimics what will happen onboard the real spacecraft where the true thrust produced by the IPS will not be known to the filter.

For a given set of Monte Carlo runs, the amplitude terms in Eqs. 13 and 14 are random samples with zero mean and given standard deviation. The time constants, however, are kept constant for a particular set of runs. The amplitude of the thrust magnitude variations is expressed as a percentage of the nominal thrust value, and the direction amplitudes are in degrees.

In addition to thrust, the initial state (position and velocity) is varied. Each run of the Monte Carlo simulation is started with a random sample of zero mean and assumed standard deviation around the nominal initial state. In general, this is the largest error source for which the filter must solve.

Observation Model

During flight, the observables will be taken from centroiding on images of asteroids. Although the capability exists to generate simulated images to centroid, the time it takes to generate a single image precludes their use in a 100-sample Monte Carlo run. Thus, the observable generation was simplified to taking samples of the expected statistics of the observations. The process used is as follows. The true spacecraft-to-beacon vector is computed using the truth spacecraft trajectory and truth asteroid ephemerides. This vector is converted into camera coordinates, and random noise is added, with the noise having zero mean and a given standard deviation. The resulting pixel and line values are passed to the filter as the observations.

As mentioned earlier, the ephemerides of the asteroids are not perfectly known, and the error is not accounted for in the filter. This effect is simulated in the Monte Carlo runs by using a different ephemeris for the truth observable generation as compared to the nominal ephemerides used by the filter to get the computed observables. To get a precise representation of this error would require that the covariance of the ephemerides of each beacon asteroid be sampled, and this value added to the nominal ephemerides for the truth. This process is time consuming, however, so a simpler solution was used. A single number representing a crude mean of the ephemeris errors of all the beacons is used, and a random sample for the three-dimensional position error of each beacon is drawn using this value as the standard deviation and added to the nominal to get the truth. For the flyby target asteroid, though, a separate value for the uncertainty in the radial, transverse, and

normal components of the ephemeris error is sampled to get the truth. Thus, the target asteroid, being a special case for evaluation, has a more realistic representation of its ephemeris error.

Evaluation of Results

The evaluation of the filter performance is done by differencing the known truth values with values obtained by the filter, and then collapsing the results for the 100 samples by computing statistics on the differences. The values used for evaluation depend on the mission phase. During cruise, maneuvers are computed using the epoch state value estimates mapped to the current time, which represents the best knowledge of the trajectory from which to plan course corrections. Thus, the cruise performance is evaluated by comparing the mapped heliocentric cartesian state from the filter with the concurrent true state. During encounter, however, the increasing power of the target asteroid data will cause the heliocentric trajectory to be adjusted to fit the target-relative data. If no target ephemeris errors were present, then the heliocentric and target-relative path would be the same. Since the simulation (and reality) will have this error, the estimated trajectory is adjusted by the filter so that it is correct relative to the target, but not necessarily in heliocentric space. For evaluation of the encounter results then, we use the true spacecraft state relative to the true target state, differenced with the estimated spacecraft position state to the nominal target state. In addition, the target-relative states are transformed into the so-called “B-plane” encounter coordinates. The B-plane is an imaginary plane centered on the flyby target and perpendicular to the incoming trajectory asymptote. It is defined by three mutually orthogonal unit vectors: \mathbf{S} , parallel to the relative incoming trajectory asymptote and normal to the B-plane; \mathbf{T} , in the B-plane and parallel to the program reference plane (Earth Mean Equator of J2000.0); and $\mathbf{R} \equiv \mathbf{S} \times \mathbf{T}$, also in the B-plane. The intersection of the incoming asymptote with the B-plane defines the vector \mathbf{B} , whose components are denoted as $\mathbf{B} \cdot \mathbf{R}$ and $\mathbf{B} \cdot \mathbf{T}$. Finally, distances in the \mathbf{S} direction are usually converted into equivalent times of flight by dividing by the hyperbolic excess velocity.

Another criterion used for evaluation is the additional $\Delta \mathbf{V}$ needed to achieve the target beyond the nominal thrusting. Recall that the nominal thrusting includes only the mission IPS burns early in the cruise; IPS and RCS TCMs are nominally zero. The combination of launch injection errors and OD and maneuver execution errors during the course of the mission cause deviations from the nominal trajectory which need to be corrected by the TCMs. For the IPS, corrections in the direction do not require additional fuel, but corrections to the duration do. Thus, the amount of change in the IPS durations required to correct the errors is a measure of performance. Similarly, statistics on the required $\Delta \mathbf{V}$ for the RCS burns are also computed and presented.

Results

The results for the nominal case assume the current best estimates for baseline error values which affect the trajectory and the observations. The following uncertainty values were used (all values are 1σ):

- Initial launch+15 day injection errors of 5000 km in position, 0.5 m/s in velocity.
- IPS thrust magnitude execution errors of 2% of the nominal.
- IPS thrust direction execution errors of 1.0° in α and δ .
- Time constant for execution errors in magnitude and direction of ∞ (in other words, the error is a bias across the mission duration).
- IPS thrust magnitude measurement errors of 1.5% of the nominal.
- IPS thrust direction measurement errors of 0.05° .
- Time constant for measurement errors of ∞ .
- Data noise of 0.1 pixel.

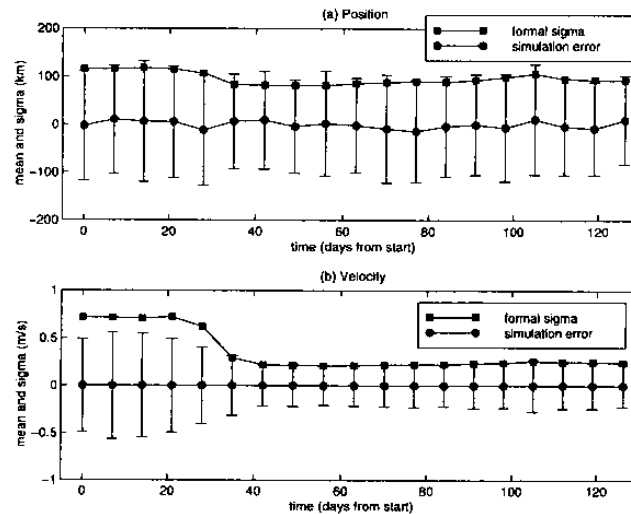


Figure 3: 1σ Formal Uncertainty from Filter and Actual Statistics from Simulations for Cruise

- Beacon asteroid ephemeris errors of 100 km.
- Target asteroid ephemeris errors of 40 km, 16 km, and 20 km in the radial, transverse, and normal directions, respectively.

The cruise results are shown in Fig. 3. Plotted are the formal covariance sigmas obtained by the filter as well as the mean and standard deviation of the actual errors from the 100 Monte Carlo simulation runs. Overall, the statistics of the actual errors matched the predicted uncertainty from the filter. In position, the errors in the early and late parts of the cruise came fairly close to the formal sigma, while in the middle, the standard deviation was roughly 1.5σ . In velocity, the standard deviations never exceeded the formal sigmas and the time history of simulation statistics almost exactly matched that of the formal sigmas. In addition, since the mean of the errors was near zero, the implication is that ignoring asteroid ephemeris errors did not introduce significant biases into the estimates, and that these errors were sufficiently averaged out. The effect of the nongravitational accelerations is shown by the fact that the estimates did not improve markedly over the course of the mission. The initial position determination was good to about 120 km, and this improved to only about 95 km. Slightly more improvement was seen in the velocity error, which decreased from about 0.5 m/s to 0.2 m/s.

Although these results for the heliocentric spacecraft trajectory are not as accurate as those achievable by standard Doppler and range tracking, the advantage of using optical data becomes obvious when examining its capability of delivering the spacecraft to its target. Fig. 4 shows a plot of the mean and standard deviation of the truth minus estimated errors in the target B-plane coordinates, along with the expected 1σ uncertainty from the filter, for the final 2 days before encounter. In this case, the mean values show a bias of about 0.5 km and 0.2 km in $\mathbf{B} \cdot \mathbf{R}$ and $\mathbf{B} \cdot \mathbf{T}$, and about 1.6 second in TOF. This is caused by the systematic error of the center-of-brightness to center-of-mass offset in the observations of the extended body. Because the object is expected to be small, however, this bias is not a critical factor in choosing the flyby aimpoint. The standard deviation of the errors about the mean are similar in magnitude in the crosstrack components, and about 1.6σ in TOF.




Figure 4: 1σ Formal Uncertainty from Filter and Actual Statistics from Simulations for Encounter

This plot clearly indicates the ability of the optical data to determine the crosstrack target-relative position of the spacecraft. Both the expected and actual errors shrink rapidly from several km at E-2 days to sub-kilometer levels at E-3 hours. In the TOF, or downtrack, component however, there is little improvement after E-2 days. For this reason, it was decided that the final four RCS TCMs only control the crosstrack errors in the B-plane, and accept the TOF control provided by the last IPS TCM. As will be described shortly, this results in considerably smaller maneuvers required by the RCS at almost no cost in delivery accuracy.

The spacecraft delivery to its flyby aimpoint is shown in Figs. 5 and 6. Fig. 5 plots the target B-plane, overlain with the expected size of the asteroid, the flyby aimpoint, and the 3σ ellipsoid defining the expected uncertainty from the filter of the delivery. The scatter of dots shows the true flyby location after the E-3 hour targeting maneuver from the Monte Carlo simulation runs. Even with the half km bias in the OD results, it can be seen that the predicted subkilometer level control of the flyby aimpoint was met in about 85% of the cases. The rms of the errors was 0.8 km, and the maximum was 1.7 km. In no case was there a danger of impacting the asteroid.

The errors in the downtrack, or TOF direction, is shown as a histogram in Fig. 6. The two vertical dashed lines in this plot show the 3σ formal sigma in the TOF axis, and the histogram plots the number of samples out of the 100 which fell into a particular time bin. Once again, the majority of cases are within the formal error bounds, with only a few cases exceeding it. The maximum values are 43 seconds on the late side and 38 seconds on the early side. Overall, larger error values and sigmas are seen in the TOF axis as opposed to those in the B-plane itself due to the lack of direct information about this axis from the optical data. The errors in the TOF can be reduced only very near to encounter, when the LOS direction to the asteroid rotates perpendicular to the downtrack direction.

Clearly, the flyby results for the nominal case in all three axes are acceptable in terms of safe delivery to the target. For the primary science goal of imaging the asteroid during closest approach, however, improvements are needed. In particular, the TOF uncertainty would preclude keeping sight of the asteroid with a 0.6° FOV camera during the flyby. Thus, the RSEN filter described earlier will be used to update the pointing information. The uncertainties in the OD after the last targeting maneuver will be reduced by RSEN during the terminal approach.




Figure 5: True B-plane Delivery Locations from Monte Carlo Simulations




Figure 6: True Time-of-flight Errors from Monte Carlo Simulations

Table 2: IPS Duration Change Statistics

Minimum Duration Change	-14.0 hours
Maximum Duration Change	57.0 hours
Mean Duration Change	22.9 ± 14.4 hours

Table 3: RCS ΔV Statistics

Minimum ΔV	0 m/s
Maximum ΔV	0.25 m/s
Mean ΔV	0.09 ± 0.04 m/s

The amount of change in the IPS and RCS thrust profiles needed to achieve the target conditions in the presence of nominal errors is given in Tables 2 and 3. In Table 2, the sum of all the duration changes for each sample run was tallied, and the statistics on the 100 sums were computed. The minimum duration change is negative because, in 7 samples, the final mission burn had to be shortened from its nominal 6 day duration, and the sum of the remaining IPS TCM durations did not exceed this decrement. At first glance, this would appear to be a benefit since less fuel is expended to reach the first target, leaving more ΔV capability for the remainder of the mission. However, since the thrust profile is optimized for the entire mission assuming a certain spacecraft mass, the heavier spacecraft may not be able to reach its second target using the nominal profile, which may prompt a redesign of the trajectory.

Table 3 shows similar statistics on the ΔV magnitude sums using the RCS engines. Here, the minimum is zero because in 1 sample, the targeting using IPS was accurate enough such that the RCS was never used for maneuvering. The worst case is only 0.25 m/s; this is easily achievable by the RCS thrusters, which have the capability of providing close to 2 m/s of ΔV.

Fig. 7 plots the ΔV statistics for each TCM. For comparison purposes, the IPS durations were converted to ΔV by applying an approximate scale factor of 10 m/s per day of IPS thrusting (in other words, an IPS duration of one day results in a ΔV of 10 m/s). The mean ΔV and its standard deviation for the 100 samples is plotted. As expected, the largest value occurred at the first IPS TCM, which made an average correction of 5 ± 2.6 m/s. In general, the earlier TCMs make larger corrections, and they are used more often. In this case for example, maneuver 9, the first TCM, was required in 92 samples, as compared to the E-1 day TCM being used in 61 samples, the E-12 hour TCM in 29 samples, the E-6 hour TCM in 19 samples, and the E-3 hour in only 5 samples.

The results from the nominal case validate the maneuver strategy of not controlling the TOF using the RCS thrusters. As a comparison, a set of Monte Carlo runs were made where all three components were targeted in the final four TCMs. These results showed an order of magnitude increase in the ΔV, with the mean value jumping from less than 0.1 m/s to over 1.6 m/s. The maximum value in several instances hit a software limit of 2.5 m/s. The delivery in the B-plane was almost identical, and the TOF miss went from from an rms value of 18.3 seconds down to 16.1 seconds. This marginal improvement in the TOF control obviously does not justify the increased fuel expenditure needed to achieve it.

CONCLUSIONS

The simulations described in this paper are the first results of performance testing on the DS-1 autonomous navigation system. This test validates the basic concept of using onboard optical sightings as the sole data type, and proves that, under certain assumptions, the system is capable of navigating a spacecraft safely to a close flyby of an asteroid. In addition, a statistical look at the




Figure 7: ΔV Statistics per Maneuver

additional ΔV required from the IPS and RCS engines under these assumptions was accomplished, and this revealed that the values obtained are within the capabilities of the current hardware. Finally, the simulations also served the purpose of functional testing of the components of the navigation system.

The testing is far from over, however, and many more simulations need to be run before full confidence in the system can be established. The performance in the presence of variations in the error sources, including worst case scenarios, needs to be analyzed. In a similar vein, the software needs to be stressed to its limit to find out when and under what conditions it fails. Since an autonomous flight system needs to be exceptionally robust, these failure modes need to be identified and handled gracefully to avoid loss of the spacecraft. In addition to preparing the software, the ground testing will also prepare the analysts to handle problems and contingencies during the flight of a revolutionary method of navigation.

Acknowledgment

The work described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

References

1. M. D. Rayman, D. H. Lehman, "NASA's First New Millennium Deep Space Technology Validation Flight", IAA Paper IAA-L-0502, Proceedings of the Second IAA International Conference on Low-Cost Planetary Missions, Laurel, MD, April 1996.
2. J. E. Riedel, S. Bhaskaran, S. P. Synnott, S. D. Desai, W. E. Bollman, P. J. Dumont, C. A. Halsell, D. Han, B. M. Kennedy, G. W. Null, W. M. Owen Jr., R. A. Werner, B. G. Williams, "Navigation for the New Millennium: Autonomous Navigation for Deep Space 1", Proceedings of the 12th International Symposium on Space Flight Dynamics, Darmstadt, Germany, June 1997.

3. S. D. Desai, S. Bhaskaran, W. E. Bollman, C. A. Halsell, J. E. Riedel, S. P. Synnott, "The DS-1 Autonomous Navigation System: Autonomous Control of Low-Thrust Propulsions Systems", AIAA Paper 97-38819, AIAA Guidance, Navigation and Control Conference, New Orleans, LA, August 1997.
4. R. M. Vaughan, J. E. Riedel, R. P. Davis, W. M. Owen Jr., S. P. Synnott, "Optical Navigation for the Galileo Gaspra Encounter", AIAA Paper 92-4522, AIAA/AAS Astrodynamics Conference, Hilton Head, S.C., August 1992.
5. S. Bhaskaran, J. E. Riedel, S. P. Synnott, "Autonomous Optical Navigation for Interplanetary Missions", *Space Sciencecraft Control and Tracking in the New Millennium*, E. Kane Casani, Mark A. Vander Does, Editors, Proc. SPIE, pp. 32-43 (1996).
6. P. B. Liebelt, *An Introduction to Optimal Estimation*, Addison-Wesley, 1967.
7. W. M. Owen Jr., R. M. Vaughan, "Optical Navigation Program Mathematical Models", JPL Internal Document JPL-EM 314-513, August 9, 1991.
8. G. J. Bierman, "Measurement Updating Using the U-D Factorization", Proc. 1975 Conf. on Decision and Control, Houston, TX, pp. 337-446.
9. S. Bhaskaran, J. E. Riedel, S. P. Synnott, "Autonomous Nucleus Tracking for Comet/Asteroid Encounters: The STARDUST Example", Paper AAS 97-628, AAS/AIAA Astrodynamics Conference, Sun Valley, ID, August 1997.

Autonomous Optical Navigation
(AutoNav)
DS1 Technology Validation Final Report

Appendix E
The DS1 Autonomous Navigation System:
Autonomous Control of Low Thrust
Propulsion System

Copyright © AIAA 1997 – Printed with permission.

THE DS-1 AUTONOMOUS NAVIGATION SYSTEM: AUTONOMOUS CONTROL OF LOW THRUST PROPULSION SYSTEM

S. D. Desai, S. Bhaskaran, W. E. Bollman, C. A. Halsell,
J. E. Riedel, S. P. Synnott

Navigation and Flight Mechanics Section, Jet Propulsion Laboratory,
California Institute of Technology, Pasadena, California

Abstract

The Deep Space-1 (DS-1) mission to be launched in 1998 will use an autonomous navigation system to guide the spacecraft on a low thrust trajectory to flybys of an asteroid and a comet. The ion propulsion system to be validated on DS-1 will provide low thrust solar electric propulsion to the spacecraft and presents additional challenges to the development of the autonomous navigation system. In order to maintain a trajectory to the designated mission target bodies, the autonomous navigation system must autonomously determine the orbit of the spacecraft, and adjust the thrust profile to be implemented by the ion propulsion system to correct any deviations from the nominal spacecraft trajectory. A detailed description of the component of the autonomous navigation system that controls the low thrust profile of the ion propulsion system is presented, and examples of some tests of this system are used to illustrate its capabilities.

Introduction

The first of NASA's New Millennium technology validation missions, the Deep Space-1 (DS-1) mission¹, will be used to demonstrate and validate the first completely autonomous navigation system ever used by an interplanetary mission. Among the various technologies to be validated on the DS-1 mission, the most important is the use of an ion propulsion system (IPS) as the primary propulsion system of the spacecraft. The IPS provides solar electric propulsion (SEP) by accelerating ionized xenon gas

Copyright ©1997 by the American Institute of Aeronautics and Astronautics, Inc. No copyright is asserted in the United States under Title 17, U.S. Code. The U. S. Government has a royalty-free license to exercise all rights under the copyright claimed herein for government purposes. All other rights are reserved by the copyright owner.

through a large potential. Historically, spacecraft have usually been powered by "chemically" powered engines, but the total impulse available from these engines has been limited by the mass of propellant that the spacecraft can carry. The few maneuvers that are implemented with the conventional chemical engines have usually been limited to durations that are each as short as a few minutes. In contrast, SEP has the capacity to provide continuous low thrust to the spacecraft, of the order of tens of millinewtons, for durations that are as long as many months. SEP is especially beneficial to high energy interplanetary missions where large changes in the energy of the orbit of the spacecraft can be achieved with considerably less mass than a chemical propulsion system.

The low thrust provided by the IPS is the largest nongravitational force acting on the spacecraft, and errors in the pointing angle, duration, and magnitude of the thrust applied by the IPS on DS-1 are likely to be the largest cause for deviations from the nominal spacecraft trajectory. The implementation of the nominal design of the SEP thrust profile on DS-1 is expected to have accuracies of the order of 1-2%. Continuous monitoring of the IPS and regular updates of the thrust pointing angles and thrust durations will be necessary to correct for deviations from the designed SEP thrust profile and spacecraft trajectory. Although redesigns of the SEP thrust profile could be computed on the ground, it would be much more efficient and advantageous to compute corrections to the designed SEP thrust profile on the spacecraft itself since these updates are expected to occur frequently. Autonomous control of the IPS on DS-1 is an integral part of the autonomous navigation system.

The DS-1 autonomous navigation system will use autonomous optical navigation (OPNAV) to determine the best estimated orbit of the spacecraft. This

best estimate of the spacecraft state will then be used to compute the corrections to the designed SEP thrust profile that are necessary to maintain a spacecraft trajectory to the designated targets. The OPNAV system uses a camera onboard the spacecraft to take images of the relative positions of asteroids with respect to the spacecraft. The beacon asteroids are then used to triangulate for the spacecraft position using precise orbit determination techniques. More details of the DS-1 autonomous navigation system and the OPNAV system are described elsewhere^{2,3,4}. This paper is devoted to describing the current strategies and algorithms that will be used by the autonomous guidance and control component of the DS-1 autonomous navigation system to adjust the designed SEP thrust profile to be implemented by the IPS in order to achieve the specific target conditions. The results from some tests used to validate this low thrust trajectory guidance and control system are also discussed.

Definition of the Designed Thrust Profile

The nominal SEP thrust profile for the low thrust trajectory of DS-1 is designed prior to launch as a completely independent process to the autonomous navigation system⁵. At present, the DS-1 trajectory is being designed for an encounter with the asteroid McAuliffe, a flyby of Mars, and an encounter with the asteroid West-Kahoutek-Ikemoura (WKI). The DS-1 autonomous control system will be responsible for computing updates and small changes to the designed SEP profile. However, if the corrected SEP thrust profile becomes energetically disadvantageous for subsequent encounters, or if there are significant outages in the IPS, the ground navigation team will have opportunities to redesign the SEP profile for uplink to the DS-1 autonomous navigation system. It is likely that early redesigns will occur immediately after launch to account for orbit injection errors, and after the IPS has been calibrated.

In order to simplify the design and control of the DS-1 trajectory, the designed SEP thrust profile will be split into successive planning cycles. The majority of the planning cycles will have a duration of 7 days, while plans on approach to the target encounter time will become successively shorter. This allows the autonomous navigation system to prepare, or plan, the SEP profile for upcoming plans by computing the precise orbit of the spacecraft before computing the adjusted SEP profile for the future plans that occur before encounter time. Figure 1 provides a heliocentric view in the equatorial plane of a sample DS-1 low thrust trajectory to encounters




Figure 1: Sample DS-1 Trajectory to McAuliffe and West-Kahoutek-Ikemoura

with McAuliffe and WKI. The launch date for this trajectory is July 1, 1998, and the encounters with McAuliffe and WKI are on January 17, 1999 and June 6, 2000, respectively.

The SEP profile for each planning cycle k , for $k = 0$ to K , will be defined by a constant thrust magnitude T_k and consequently a constant mass flow rate, and a duration τ_k that the IPS applies this thrust during each plan. The IPS thrust pointing vector in each plan is specified by the time dependent pointing angles of right ascension $\alpha(t)$, and declination $\delta(t)$, which are each defined by first order polynomials of time in each plan.

$$\alpha(t) = \alpha_k + \dot{\alpha}_k(t - t_k); t_k \leq t < t_k + \tau_k \quad (1)$$

$$\delta(t) = \delta_k + \dot{\delta}_k(t - t_k); t_k \leq t < t_k + \tau_k \quad (2)$$

In addition, a particular *duty cycle* D is imposed on the SEP profile of the low thrust trajectory when it is designed, where the duty cycle specifies the maximum duration that the IPS is permitted to thrust in each planning cycle. A constant duty cycle is usually defined for the entire SEP thrust profile. Here, reference will also be made to SEP *segments*, where an individual SEP segment refers to the combination of SEP plans where the IPS is thrusting continuously except for the time at the end of a SEP plan where the IPS is not thrusting only because of the imposed duty cycle limitations. This means that all of the plans except for the last plan in any particular SEP segment will have a thrust duration that is exactly at the specified duty cycle limit. Only the last plan κ of each SEP segment is permitted to have a thrust duration that is free to range from zero duration to

the duration available from the specified duty cycle limit. Given the start time t_k of each planning cycle k in a SEP segment, the implicit constraint on the durations that the IPS is permitted to thrust in each plan of a particular SEP segment is as follows.

$$\tau_k = D(t_{k+1} - t_k) \text{ when } k \neq \kappa \quad (3)$$

$$0 \leq \tau_\kappa < D(t_{\kappa+1} - t_\kappa) \quad (4)$$

All SEP plans that are not part of a SEP thrusting segment will have a thrust duration of $\tau_k = 0$.

The nominal DS-1 SEP profile is designed to allow approximately 8% of the duration in each planning cycle to be devoted to telecommunications with ground operations, and to taking the images of the asteroids that are used as beacons by the OPNAV system for the autonomous orbit determination of the spacecraft. Due to attitude constraints on the spacecraft the IPS cannot be operating during either of these procedures. The remaining 92% of the duration in each planning cycle is available for thrusting by the IPS. For the actual DS-1 flight the SEP profile will be designed such that the IPS will have a 92% duty cycle. However, for the purposes of testing the autonomous navigation system, and especially the autonomous control system, trajectories with a suboptimal 85% duty cycle are currently being used. This approach is taken to ensure that trajectories with suboptimal performance from the IPS are available for the McAuliffe and WKI encounters, but also to ensure that the autonomous control system is capable of controlling the DS-1 trajectory if the IPS does not perform to the specified 92% duty cycle specifications.

The DS-1 trajectory shown in Figure 1 is designed to an 85% duty cycle, and the associated SEP profile between launch and the McAuliffe encounter, is shown in Figure 2. The pointing angles in each SEP segment could be considered to be continuous except for the time during the SEP plans when the IPS is not thrusting because of the specified duty cycle limit. The SEP profile for the McAuliffe encounter, shown in Figure 2, has two SEP segments. The first SEP segment begins 15 days after launch, is approximately 10 days long, and contains 2 SEP plans. The second SEP segment begins 31 days after launch, is 100 days long, and contains 16 SEP plans. The first segment at the beginning of the mission is specifically designed to be used to test and calibrate the IPS.

It should be noted that the right ascension and declination of the SEP thrust pointing vector from the last plan in each segment, α_κ and δ_κ , are extrapolated using the rates $\dot{\alpha}_\kappa$ and $\dot{\delta}_\kappa$, to the few subsequent plans which have zero IPS thrust durations.




Figure 2: Right Ascension (a), Declination (b), and Magnitude (c) of SEP Thrust Pointing Vector for DS-1 Trajectory to McAuliffe

The reason for this is to provide nominal design values of the IPS thrust pointing angles in these zero duration plans to allow the autonomous control system to include these plans as part of the individual SEP segments if it becomes necessary for the IPS to thrust during these plans. For example, there are six SEP plans at the end of the second SEP segment which are nominally designed with zero duration, but which could become part of the second SEP segment and be used to thrust by the IPS if necessary.

Those SEP plans that are not needed to correct the designed SEP profile then become available for trajectory control maneuvers (TCMs) for the ballistic phase of the trajectory before encounter. TCMs will be performed either by the IPS or the hydrazine engines on DS-1, and should usually have durations of less than 12 hours if the IPS is used to perform these maneuvers.

The DS-1 spacecraft is severely constrained in orientation, because certain faces of the spacecraft cannot be illuminated by the Sun, and because use of the IPS requires that the solar panels face directly into the Sun. These constraints in orientation translate into constraints on the pointing angle of the IPS thrust vector. When the SEP profile of the DS-1 mission is designed these angular constraints on the IPS thrust vector are specified in each plan by angles θ_k for each plan k .

Define the pointing vectors \vec{p}' and \vec{p} to be the thrust pointing vectors at the beginning of each plan of the designed SEP profile, and the corrected SEP profile, respectively.

$$\vec{p}' = [\cos \delta'_k \cos \alpha'_k \quad \cos \delta'_k \sin \alpha'_k \quad \sin \delta'_k] (5)$$

$$\vec{p} = [\cos \delta_k \cos \alpha_k \quad \cos \delta_k \sin \alpha_k \quad \sin \delta_k] (6)$$

The primes (') are used here to indicate that the pointing vectors and angles are from the designed SEP profile. The constraint angles θ_k then define the maximum angular correction that can be applied to the IPS thrust pointing vector specified at the beginning of each plan of the designed SEP profile.

$$F_k(\alpha_k, \delta_k) = \vec{p}' \cdot \vec{p} \quad (7)$$

$$\cos^{-1}(F_k(\alpha_k, \delta_k)) \leq \theta_k \quad (8)$$

The SEP thrust profile for the DS-1 autonomous navigation system is then defined by a table of t_k , T_k , τ_k , α_k , δ_k , α'_k , δ'_k , and θ_k for each of the planning cycles between launch and encounter, with the last three parameters used only to check that corrected SEP profiles do not violate the angular constraints imposed on the designed SEP profiles.

Linear Control Equation for SEP Profile

If the angular rates, $\dot{\alpha}_k$ and $\dot{\delta}_k$, and the thrust magnitudes T_k specified in the designed SEP profile are assumed to be fixed, then the remaining independent variables which provide control authority for the thrust vector from the spacecraft IPS are the pointing angles at the beginning of each plan, α_k and δ_k , and the thrust durations τ_k only from the last plan in each SEP segment, since these durations are the only durations of plans within a SEP segment that are not set at the duty cycle limits. However, the last plan defined for each SEP segment, or the value of κ , is permitted to change (increase or decrease) as it becomes necessary.

It is assumed that the autonomous control system will only be used to update the SEP profile to correct for small deviations from the nominal trajectory, while any significant deviations from the nominal trajectory will require a complete redesign of the DS-1 trajectory and SEP profile. As such, a simple linear targeting approach seems adequate for the autonomous control system. Also, the control system will be restricted to using only those plans within a single SEP segment to correct the SEP profile at any time.

The autonomous orbit determination system computes the current best estimate of the spacecraft state at some time t , and this is integrated forward in time to provide a spacecraft state at the specified encounter time t_e using the currently available SEP profile. This present course encounter state $X_e(\alpha_k, \delta_k, \tau_k)$ is a function of,

$$(\alpha_k, \delta_k, \tau_k) \text{ for } k_1 \leq k \leq \kappa, \text{ and } t < t_{k_1} < t_e$$

where the plan k_1 is the first complete plan after the time t where the best known spacecraft state has been computed by the autonomous orbit determination system. If the difference between the present course and desired encounter time spacecraft states is not below a specified tolerance threshold ϵ , then adjustments to the parameters τ_k , α_k , and δ_k for $k = k_1$ to $k = \kappa$, a total of $2(\kappa - k_1 + 1) + 1$ parameters, can be used to guide the spacecraft to the required target state. The desired target state $\bar{X}_e(\bar{\alpha}_k, \bar{\delta}_k, \bar{\tau}_k)$ is a function of,

$$(\bar{\alpha}_k, \bar{\delta}_k, \bar{\tau}_k) \text{ for } k = k_1 \text{ to } k = \kappa,$$

where the overbars ($\bar{\quad}$) are used to indicate the adjusted SEP profile variables that are necessary to achieve the required target state. It is these variables, $(\bar{\alpha}_k, \bar{\delta}_k, \bar{\tau}_k)$ that must be determined by the autonomous control system.

For small deviations from the nominal trajectory it should not be necessary to use all of the available pointing angles to guide the spacecraft to the target state, and a subset of the pointing angles from plans $k = k_1$ to $k = \kappa$ could be used. If a strategy that attempts to correct the low thrust trajectory as soon as possible is adopted then the DS-1 control system will be restricted to using the pointing angles from all plans from plan k_1 to plan k_2 to provide control authority to the IPS, where k_2 is restricted as follows.

$$k_1 \leq k_2 \leq \kappa \quad (9)$$

The required target state can be expanded into a Taylor series expansion about the present course encounter state and SEP profile as defined by the independent variables α_k , δ_k and τ_κ . Assuming that a target trajectory SEP profile only has small deviations from the present course trajectory SEP profile, then retaining only the linear terms from the Taylor series expansion provides the linear control equation for the DS-1 SEP profile.

$$\Delta X_e = K \Delta s \quad (10)$$

The vector ΔX_e is the difference between the desired target state and the spacecraft state at encounter time computed from the current SEP profile.

$$\Delta X_e = \bar{X}_e(\bar{\alpha}_k, \bar{\delta}_k, \bar{\tau}_\kappa) - X_e(\alpha_k, \delta_k, \tau_\kappa) \quad (11)$$

The matrix $K(\alpha_k, \delta_k, \tau_\kappa)$ contains the first order partial derivatives of the control variables, and should be evaluated from the present course SEP profile used to compute $X_e(\alpha_k, \delta_k, \tau_\kappa)$.

$$K^T(\alpha_k, \delta_k, \tau_\kappa) = \begin{bmatrix} (\partial X_e / \partial \alpha_{k_1}) \\ (\partial X_e / \partial \delta_{k_1}) \\ (\partial X_e / \partial \alpha_{k_1+1}) \\ (\partial X_e / \partial \delta_{k_1+1}) \\ \vdots \\ (\partial X_e / \partial \alpha_{k_2}) \\ (\partial X_e / \partial \delta_{k_2}) \\ (\partial X_e / \partial \tau_\kappa) \end{bmatrix} \quad (12)$$

The operator $[\cdot]^T$ denotes the transpose of the matrix $[\cdot]$. The partial derivatives in the K matrix are numerically computed using finite central differences. An example is given below.

$$\frac{\partial X_e}{\partial \alpha_{k_1}} = \frac{X_e(\alpha_k, \delta_k, \tau_\kappa) |_{\alpha_{k_1+\epsilon}} - X_e(\alpha_k, \delta_k, \tau_\kappa) |_{\alpha_{k_1-\epsilon}}}{2\epsilon} \quad (13)$$

The control vector Δs contains the first order corrections to the control variables of the SEP profile.

$$\Delta s = \begin{bmatrix} \bar{\alpha}_{k_1} - \alpha_{k_1} \\ \bar{\delta}_{k_1} - \delta_{k_1} \\ \bar{\alpha}_{k_1+1} - \alpha_{k_1+1} \\ \bar{\delta}_{k_1+1} - \delta_{k_1+1} \\ \vdots \\ \bar{\alpha}_{k_2} - \alpha_{k_2} \\ \bar{\delta}_{k_2} - \delta_{k_2} \\ \bar{\tau}_\kappa - \tau_\kappa \end{bmatrix} \quad (14)$$

A total of $M = 2(k_2 - k_1 + 1) + 1$ variables provide control authority for the DS-1 low thrust trajectory, and Δs is a vector of dimension M . The two pointing angles from at least the first available plan k_1 in a SEP segment, and the duration from the last plan κ of that SEP segment are always included in the search for an updated SEP profile, and $M \geq 3$ always. If N is used to denote the dimension of the target vector ΔX_e , then K is a matrix of dimension $N \times M$. The target vector is defined either by the three dimensional position coordinates at encounter time, or by the six dimensional state including position and velocity, so that $N = 3$ or $N = 6$ always. When targeting to the three dimensional position, the residual target vector ΔX_e is always specified in terms of target relative asymptotic coordinates in plane of the trajectory.

$$\Delta X_e^T = [\Delta B \cdot R \quad \Delta B \cdot T \quad \Delta TOF] \quad (15)$$

The target relative coordinates $B \cdot R$ and $B \cdot T$ define positions in the two crosstrack directions, and TOF defines the along track position in terms of a time of flight with respect to the point of closest approach.

The corrections to the SEP profile that are needed to guide the spacecraft to the target state are solved through iterative solutions of Equation (10) for Δs . In the first iteration, the present course trajectory SEP profile is used to compute the matrix $K(\alpha_k, \delta_k, \tau_\kappa)$ and the encounter time state $X_e(\alpha_k, \delta_k, \tau_\kappa)$, which then provides a first order solution of the corrections Δs and an updated SEP profile defined by $(\bar{\alpha}_k, \bar{\delta}_k, \bar{\tau}_\kappa)$. The updated SEP profile then becomes the present course trajectory SEP profile in the next iteration, $(\alpha_k, \delta_k, \tau_\kappa) = (\bar{\alpha}_k, \bar{\delta}_k, \bar{\tau}_\kappa)$, from which the next set of SEP profile corrections are computed. If the corrected duration of the last plan extends past its boundaries, as specified in Equation (4), the value of κ is increased or decreased as becomes necessary. This procedure is repeated until the norm of the residual between the target state and the encounter state is within the specified threshold e .

$$|\Delta X_e| < e \quad (16)$$

A convergence criteria of 1 km in position and 10^{-5} km/s in velocity is usually sufficient.

Equation (10) is a linearized equation, and convergence of the iterations required to solve this equation are not guaranteed. However, tests have shown that when the iterative solution to the linear control equation does not converge there is usually an insufficient number of control parameters in the control vector Δs . As such, when the iterative procedure does not converge within a specified finite number of iterations, more parameters are added to the control vector. More specifically, k_2 is incremented in steps of 1, and the dimension of the control vector is increased in steps of 2, by sequentially adding the two pointing angles of consecutive SEP plans in steps of one plan at a time, until a converged solution is found. An obvious failure mode of the control system then arises when $k_2 > \kappa$ and there are no more control parameters available to find a converged solution, and the ground navigation system would then be notified to redesign the SEP profile.

Solution Strategies of Control Equation

The method used to solve Equation (10) is dependent on the dimension M of the control vector Δs with respect to the dimension N of the residual encounter state vector ΔX_e . This results with three cases which each require different solution methods. Similar solution methods are also used when the angular constraints are imposed.

Case 1. $N = M$

This is the simplest case where the number of equations and control parameters are identical. For each iteration, a unique solution of Δs from the control equation is computed from a simple inversion of the matrix K .

$$\Delta s = K^{-1} \Delta X_e \quad (17)$$

Case 2. $N > M$

In the case where there are fewer control parameters than equations, the corrections Δs are computed from least squares solutions to Equation (10) at each iteration. That is, the corrections to the SEP profile are chosen to be the vector Δs that minimizes the following performance index J .

$$J = \frac{1}{2} (\Delta X_e - K \Delta s)^T (\Delta X_e - K \Delta s) \quad (18)$$

The least squares solution to the control equation is found by minimizing J with respect to Δs .

$$\Delta s = (K^T K)^{-1} K^T \Delta X_e \quad (19)$$

Note that since $N = 3$ or $N = 6$, and $M \geq 3$ always, the least squares solution is only used when targeting to a position and velocity at encounter time with the angles of fewer than 3 planning cycles. The converged least squares solutions only provide a minimum to the performance index and the residual encounter state ΔX_e , and the iterative search ends when this minimum is reached even though it does not necessarily lie within the threshold limit ϵ .

Case 3. $N < M$

When there are more control parameters than the dimension of the target state, the solution to the control equation is chosen to be the solution that minimizes the corrections Δs subject to the constraint $\Delta X_e = K \Delta s$. The performance index is:

$$J(\Delta s, \lambda) = \frac{1}{2} (\Delta s^T \Delta s) + \lambda (\Delta X_e - K \Delta s) \quad (20)$$

where the constraint has been adjoined with the Lagrange multiplier λ . The first variation of $J(\Delta s, \lambda)$ with respect to Δs and λ is given as δJ below.

$$\begin{aligned} \delta J &= \frac{1}{2} (\delta \Delta s^T \Delta s + \Delta s^T \delta \Delta s) \\ &\quad - \lambda K \delta \Delta s + \delta \lambda (\Delta X_e - K \Delta s) \end{aligned} \quad (21)$$

Note that $\delta \Delta s^T \Delta s = \Delta s^T \delta \Delta s$. For a minimum of $J(\Delta s, \lambda)$, the first variation δJ must vanish for arbitrary $\delta \Delta s$ and $\delta \lambda$, and the following two equations must be satisfied to have $\delta J = 0$.

$$\Delta s^T - \lambda K = 0 \quad (22)$$

$$\Delta X_e - K \Delta s = 0 \quad (23)$$

Inserting the transpose of Equation (22) into Equation (23) provides a solution for λ which can be inserted into the transpose of Equation (22) for a solution for Δs .

$$\lambda^T = (K K^T)^{-1} \Delta X_e \quad (24)$$

$$\Delta s = K^T (K K^T)^{-1} \Delta X_e \quad (25)$$

Equation (25) involves an inversion of an $N \times N$ matrix whose dimension is completely independent of the number of control parameters M in Δs , and therefore never exceeds a dimension of 6.

With Angular Constraints

After a converged solution for an updated SEP profile is computed from one of the above three solution methods it then becomes the new present course SEP profile. This new present course SEP profile

is then checked to ensure that the thrust pointing vector at the beginning of each plan satisfies the angular constraint requirements from Equation (8). If the initial thrust pointing vector of any plan in this new SEP profile violates the angular constraint then corrections to this new SEP profile are computed, but by imposing an angular constraint equality to the pointing angles of all of the plans that violate the constraints. If the pointing angles from all of the plans from k_1 to k_2 that were included into the control vector used to compute this new SEP profile violate their respective angular constraints then in addition to applying the angular constraint equality to all of these plans, k_2 is incremented by 1 to include the pointing angles of the next consecutive SEP plan to the control vector but without any angular constraint applied to this additional plan. As before, this procedure is repeated until a converged solution of an updated SEP profile where all the plans satisfy the angular constraints is found. When $k_2 > \kappa$ and no more plans are available to add to the iterative search, the ground navigation system is notified to redesign the SEP profile.

The angular constraint equality imposed on all of the plans which violate the constraint requirement in Equation (8) is as follows.

$$F_k(\bar{\alpha}_k, \bar{\delta}_k) = \hat{p}' \cdot \hat{p} = \cos \theta_k \quad (26)$$

A first approximation of this constraint equality is made by defining an updated SEP profile which resets the pointing angles of the initial pointing vector \hat{p} of all of the violating SEP plans in the present course SEP profile to a pointing vector \hat{p} that satisfies the constraint equality in Equation (26), that lies in the plane defined by \hat{p} and the initial pointing vector of the design trajectory \hat{p}' , and that lies in between \hat{p} and \hat{p}' .

$$(\hat{p}' \times \hat{p}) \cdot \hat{p} = 0 \quad (27)$$

$$\hat{p} \cdot \hat{p} = \cos [\theta_k - \cos^{-1}(\hat{p}' \cdot \hat{p})] \quad (28)$$

This first approximation of the updated SEP profile becomes the new present course SEP profile and although it now satisfies the constraint equality, the residual encounter state vector ΔX_e is usually no longer within the specified threshold ϵ . Further iterations are necessary to search for an updated SEP profile which both satisfies the constraint equality and provides a residual encounter state that is within the threshold limits.

The additional iterations are performed in a similar manner to the three methods already described above, except with additional equations that define the angular constraint equality. The linearized form

of the angular constraint equality for an arbitrary plan k is found by expanding Equation (26) into a Taylor series about the new present course trajectory and retaining only the linear terms.

$$\begin{aligned} \Delta F_k &= F_k(\bar{\alpha}_k, \bar{\delta}_k) - F_k(\alpha_k, \delta_k) \\ &= A_k \Delta s \end{aligned} \quad (29)$$

The only nonzero elements elements of the vector A_k are those that correspond to the elements of Δs with right ascension and declination corrections for SEP plan k .

$$[A_k]_i = \begin{cases} \frac{\partial F_k}{\partial \alpha_k} & i = 2(k - k_1) + 1 \\ \frac{\partial F_k}{\partial \delta_k} & i = 2(k - k_1) + 2 \\ 0 & \text{all other } i \end{cases} \quad (30)$$

An expression like Equation (29) is necessary for all those plans that had violated the angular constraint in any of the prior converged solutions for a SEP profile. The partial derivatives are evaluated from the present course SEP profile, and are analytically represented as follows.

$$\frac{\partial F_k}{\partial \alpha_k} = \cos \delta'_k \cos \delta_k (\sin \alpha'_k \cos \alpha_k - \cos \alpha'_k \sin \alpha_k) \quad (31)$$

$$\begin{aligned} \frac{\partial F_k}{\partial \delta_k} &= \sin \delta'_k \cos \delta_k \\ &- \cos \delta'_k \sin \delta_k (\cos \alpha'_k \cos \alpha_k + \sin \alpha'_k \sin \alpha_k) \end{aligned} \quad (32)$$

It is important to note that both of these partial derivatives are equal to zero when the pointing angles are from the designed SEP profile, with $\alpha_k = \alpha'_k$ and $\delta_k = \delta'_k$, and the matrix A_k is then singular. However, the first approximation of the angular constraint which was computed from Equations (26) to (28), already satisfies the constraint defined in Equation (26), and subsequent iterations for the updated pointing vectors will not approach the design trajectory pointing vectors since the angular constraint equality would no longer be satisfied.

The linear control equation with angular constraints can then be considered to be a combination of Equations (10) and (29).

$$\Delta Y = K_A \Delta s \quad (33)$$

$$\Delta Y = \begin{bmatrix} \Delta X_e \\ \Delta F_k \\ \vdots \end{bmatrix}, \quad K_A = \begin{bmatrix} K \\ A_k \\ \vdots \end{bmatrix} \quad (34)$$

The vector ΔY and the matrix K_A include the residuals ΔF_k and the corresponding vectors A_k , respectively, for all of the plans k that have violated the angular constraint. If there were N_A plans that violated the angular constraints, then the dimension of the vector ΔY is $(N + N_A)$, and the dimension of K_A is $(N + N_A) \times M$.

In this case, the method chosen to solve Equation (33) is now dependent on the relationship of the dimension $(N + N_A)$ to the number of the parameters M , which result with three solution methods, say Cases 1A, 2A and 3A, which are analogous to Cases 1, 2, and 3 described above.

Case 1A: $(N + N_A) = M$

$$\Delta s = K_A^{-1} \Delta Y \quad (35)$$

Case 2A: $(N + N_A) > M$

$$\Delta s = (K_A^T K_A)^{-1} K_A^T \Delta Y \quad (36)$$

Case 3A: $(N + N_A) < M$

$$\Delta s = K_A^T (K_A K_A^T)^{-1} \Delta Y \quad (37)$$

As will be mentioned later, the DS-1 autonomous control system will usually be restricted to targeting only to the three dimensional coordinates in position that are required at the encounter time. As such, the minimum norm solution described in Case 3A is always used once angular constraints are included into the iterative search for the updated SEP profile.

Simulations of Targeting to a Position Only

Examples of some tests of the linear targeting strategy to a three dimensional position at encounter time for the DS-1 trajectory to McAuliffe using the 85% duty cycle SEP profile shown in Figure (2) as the designed SEP profile are shown below. The second SEP segment to McAuliffe will probably be re-designed after the IPS has been calibrated during the first SEP segment, so the tests are restricted to simulating errors and computing updated SEP profiles only for the second SEP segment before the McAuliffe encounter. The second segment of the design trajectory begins at SEP plan $k = 3$ and ends at SEP plan $k = 18$. It is assumed that the orbit determination system provides a perfect observation of the spacecraft state at any opportunity to update the SEP profile. The actual operation of the autonomous navigation system on DS-1 is simulated by considering the planning cycles as a time line of the DS-1 trajectory. The tests step through this

time line starting with SEP plan $k = 3$, and assumes that the IPS has actually implemented a thrust in all prior SEP plans of the second segment that is equivalent to a duty cycle that is lower than the designed 85% duty cycle that would have guided the spacecraft to McAuliffe.

So, if the spacecraft is simulated to be at the beginning of plan k_1 , the lower duty cycle is imposed on all plans of the updated SEP profile from $k = 3$ to $k = k_1 - 1$, and the autonomous control system is provided with an opportunity to update the SEP profile in as many future SEP plans with $k \geq k_1$ as is necessary. For example, when $k_1 = 3$, the SEP profile is exactly as designed and no corrections are applied. When $k_1 = 4$, an error in the duty cycle of plan $k = 3$ has been applied and plans with $k \geq 4$ are used to correct this error to maintain a trajectory that has an encounter with McAuliffe. Then, when $k_1 = 5$, in addition to the error already applied to plan 3, an identical error in the duty cycle of plan $k = 4$ of the SEP profile that was updated when $k_1 = 4$ is also applied, and SEP plans with $k_1 \geq 5$ are used to correct these errors. This process is repeated to the end of the second SEP segment.

Four specific examples are shown to illustrate how changing the minimum number of plans included in each solution affects the angular and duration corrections to the designed SEP profile, and how applying the angular constraint affects these corrections. The first three examples do not impose the angular constraint. The angular and duration corrections of the updated SEP profile with respect to the designed SEP profile from the first example are shown in Figure 3. These corrections are those computed by the autonomous control system when the search for an updated SEP profile is started with only 1 SEP plan, $k_2 = k_1$. The percentages labeled on each curve indicate the duty cycle that was actually applied by the IPS in the SEP plans with $3 \leq k < k_1$. Although the iterative search is started with the angles of the first available SEP plan, a converged solution is not always found with only one plan. For example, at least two plans ($k_2 = k_1 + 1$) are necessary to find converged solutions when $k_1 = 4, 5$, and 6, and the applied duty cycles are less than 83%. As the applied duty cycle is reduced further more solution opportunities require at least two plans to find a converged solution. The extreme example is when the duty cycle applied to prior plans was 79%, and converged solutions required the used of three plans when $k_1 = 4, 5, 6, 7$, and 8, and two plans when $k_1 = 9$, and 10.

Similarly, as the applied duty cycle is reduced the number of plans in the second segment gradually increases with the value of κ increasing to the point







Figure 3: Angular corrections (a) and duration corrections (b) with respect to the designed SEP profile using a minimum of 1 SEP plan to correct prior errors in the SEP profile. No angular constraints are applied to the corrections.

Figure 4: Same as Figure 3, but using all available SEP plans to correct prior errors in the SEP profile. No angular constraints are applied to the corrections.

where $\kappa = 23$ by the end of the simulation which applied 79% duty cycles on all prior plans. When prior plans had a duty cycle of 78% a converged solution for all of the SEP plans in the second segment could not be found because the durations eventually extended beyond plan $k = 24$ where no nominal pointing angles were specified in the designed SEP profile.

Figure 4 is similar to Figure 3 except that all available plans were used to correct any prior errors in the duty cycle and $k_2 = \kappa$ always. In this example, converged solutions were also found for all of the

plans when the duty cycle applied to prior plans was 78%. This is because the duration corrections were much smaller, almost by a factor of 2, than the duration corrections when a minimum of 1 plan was used to correct errors in the duty cycle. For example, when a duty cycle of 79% was applied to prior plans, the last plan of the second segment was changed from the design value of $\kappa = 18$ to $\kappa = 23$ for the example shown in Figure 3, and to $\kappa = 20$ for the example shown in Figure 4. However, reducing the duration correction also had the effect of delaying angular corrections to the plans at the end of the SEP segment, as they accumulate through each







Figure 5: Same as Figure 3, but using a minimum of 3 SEP plans to correct prior errors in the SEP profile. No angular constraints are applied to the corrections.

Figure 6: Same as Figure 5, but with an angular constraint of 10 degrees applied to the corrections.

update of the SEP profile from $k_1 = 3$ to $k_1 = \kappa$.

Figure 5 shows the angular corrections and duration corrections when the angles from a minimum of three segments, $k_2 = k_1 + 2$, are used to correct any errors in the duty cycle of prior SEP plans. The most significant improvement over the examples shown in Figures 3 and 4 is the reduction in the maximum angular correction of the thrust pointing vector in any plan. While the maximum angular correction in the examples shown in Figures 3 and 4 are larger than 20 degrees, in this example the maximum is only as large as approximately 15 degrees. The penalty for this improvement is larger duration corrections

compared to when all the SEP plans were used to correct prior errors. However, these duration corrections are still smaller than when the angles from a minimum of 1 plan were used to update the SEP profile. In this example, when a duty cycle of 79% was applied to prior plans, the last plan of the second SEP segment is changed to $\kappa = 22$.

The designed SEP profile will usually place angular constraints on the updated SEP profiles that are of the order of 10 degrees or less. Therefore, none of the previous three examples would be suitable strategies to correct the SEP profile when applied duty cycles vary by as much as 5% from the designed 85% duty cycle. Figure 6 shows a similar example to that shown in Figure 5, except that now

a 10 degree angular constraint has been applied to the updated SEP profiles. The angular constraints have only been enforced when duty cycles of 81% have been applied to prior duty cycles. By applying these angular constraints there has also been a significant reduction in the durations required to correct the prior errors in the duty cycle. When a duty cycle of 78% was applied to all prior SEP plans, the last plan of the second SEP segment extended to $\kappa = 24$ when no angular corrections were imposed on the updated SEP profiles, but only extended to plan $\kappa = 22$ when the angular constraint was applied to the updated SEP profiles.

These four examples clearly demonstrate that using extreme strategies such as using a minimum of one plan with $k_2 = k_1$, or using all the available plans in the segment with $k_2 = \kappa$, do not provide the most desirable adjustments to the designed SEP profile. Instead, using a minimum of three plans might be considered as a reasonable compromise between correcting any errors as soon as possible, and reducing the angular and duration corrections to the designed SEP profile. Although the angular constraints are constraints imposed by the physical design of the spacecraft, they also appear to improve the efficiency of the adjusted SEP profiles by reducing the duration corrections to the adjusted SEP profiles.

Targeting to only the three dimensional coordinates in position at encounter time changes the velocity and incoming asymptote of the spacecraft at the encounter time, and could prove to be fatal for the spacecraft trajectory to the subsequent encounters. Tests of the autonomous control system have been performed to compare the adjusted SEP profiles that would result from targeting to a six dimensional state (position and velocity, $N = 6$), to those that result from targeting to a three dimensional encounter state (position only, $N = 3$). The corrections to the thrust pointing angles and durations are much smaller when targeting to a three dimensional state and probably better suited to a linear targeting strategy. Also, for the small errors expected in the SEP thrust applied by the IPS, the changes in the velocity of the spacecraft at encounter time caused by targeting to a position only, appear to be small enough to be rectified by a redesign of the SEP profile after each encounter. As such, the DS-1 autonomous control system will be restricted to linear targeting to the desired three dimensional coordinates in position at encounter time, but will maintain the capability to target to a position and velocity at encounter time. Any significant errors in the SEP thrust applied by the IPS which become en-

ergetically disadvantageous for subsequent encounters will require a redesign of the SEP profile by the ground navigation team.

Acknowledgements

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Special thanks to S. Williams for providing a designed SEP profile for the DS-1 trajectory, and B. Kennedy for the software to generate the planetary and asteroid ephemerides.

References

1. Rayman, M. D., and D. H. Lehman, "NASA's First New Millennium Deep-Space Technology Validation Flight", IAA paper L-0502, IAA Conference, Apr. 1996.
2. Riedel, J. E., S. Bhaskaran, S. P. Synnott, W. E. Bollman, and G. W. Null, "An Autonomous Optical Navigation and Control System for Interplanetary Exploration Missions", IAA paper L-0506, IAA Conference, Apr. 1996.
3. Riedel, J. E., S. Bhaskaran, S. P. Synnott, S. D. Desai, W. E. Bollman, P. J. Dumont, C. A. Halsell, D. Han, B. M. Kennedy, G. W. Null, W. M. Owen Jr., R. A. Werner, and B. G. Williams, "Navigation for the New Millennium: Autonomous Navigation for Deep-Space-1", Proceedings 12th International Symposium on Flight Dynamics, Darmstadt, Germany, June 2-6, 1997.
4. Bhaskaran, S., J. E. Riedel, S. P. Synnott, "Autonomous Optical Navigation for Interplanetary Missions, Space Sciencecraft Control and Tracking in the New Millennium", E. Kane Casani, Mark A. Vander Does, Editors, Proc. SPIE 2810, pp32-43, 1996.
5. Coverstone-Carroll, V., and S. N. Williams, "Optimal Low Thrust Trajectories Using Differential Inclusion Concepts", *Journal of the Astronautical Sciences*, Vol. 42, No. 4, October-December 1994, pp 379-393.

Autonomous Optical Navigation
(AutoNav)
DS1 Technology Validation Final Report

Appendix F
DS1 AutoNav Technology Validation
Agreement

New Millennium Program - Deep Space One Project

TECHNOLOGY VALIDATION AGREEMENT

Technology:

AUTONOMOUS OPTICAL NAVIGATION

Technology Provider:

Jet Propulsion Laboratory, Navigation and Flight Mechanics Section (312)

Technology Description:

Autonomous optical navigation (AutoNav) is the primary system to be used for DS1 (Deep Space One) spacecraft navigation. AutoNav is a completely autonomous navigation system which will:

- Provide onboard ephemeris and mass-data services, principally to ACS (Attitude Control System).
- Plan MICAS (Miniature Imaging Camera Spectrometer) picture-taking activity.
- Implement picture-taking activity through interaction with ACS.
- Reduce the resultant images to determine astrometric positions.
- Filter the astrometric data to produce spacecraft state and non-grav information [i.e. Orbit Determination (OD)].
- Compute a correction to a nominal low-thrust mission-burn profile based on encounter targeting parameters, or compute a discrete Reaction Control Subsystem or IPS (Ion Propulsion System) trajectory correction maneuver based on those parameters.
- Provide late ephemeris update information for science targeting to ACS, and start encounter science sequences, based on encounter relative estimated closest approach time.
- Provide all necessary data and file uplink and downlink capability and monitoring telemetry.
- Provide contingency plans and procedures in the event AutoNav is partially or completely disabled.

Other DS1 Technologies Dependent on Given Technology:

IPS (control and calibration - direction and thrust level)
MICAS (encounter, target body ephemeris)

Validation Criteria (Activity Definition/Description):

Pre-Flight

Responsibility: Navigation; Avionics Flight Software and Testbed; and Spacecraft Integration and Test

1. Verify stability and accuracy of main compute elements in long-duration tests in a UNIX-based environment.
2. Provide unit-test verification test runs in "Papabed" and Testbed environments for test of all AutoNav capability.
3. Provide integrated system-level intermediate-duration tests for Testbed and ATLO environments for test of all AutoNav capability in a realistic flight-like configuration.

In-Flight (Expected Flight Observables)

1. Provision of Ephemeris Service to onboard clients
2. Picture Planning, and execution of MICAS/OpNav Sequence
3. OpNav Images (exposed and downlinked)
4. Edited OpNav Images (downlinked OpNav file)
5. Image Processing Results (downlinked OpNav file)
6. OD Results, S/C State and Covariance (downlinked OD file)
7. Estimated Position Spacecraft Ephemeris (downlinked SC-50 file)
8. Estimated Changes to nominal mission burn profile (downlinked Maneuver file)
9. Autonomous operation of IPS during mission burn via execution of Nav Micro-sequences
10. Encounter updates of Spacecraft position
11. Initiation of Encounter Sequences

Success Criteria (Quantifiable/Measurable Goals):

Pre-Flight

Demonstration of ability to meet mission Navigation requirements under simulated flight-like conditions:

- 250 km, 1 m/sec (1 sigma) cruise state {75% of mission success}
- 2.5 km, 0.25 m/sec cross-track, (1 sigma). {Expected downtrack performance is dependent upon flyby altitude, velocity and time of last encounter navigation image, as well as ACS pointing knowledge performance.} {25% of mission success}

In-Flight

Consistent comparison of Radio-Navigation OD results with flight AutoNav results within reasonable 2-Dimensional mutual covariances (2.5 sigma). Demonstration of ability to meet mission Navigation requirements in flight:

- 250 km, 1 m/sec (1 sigma) cruise state
- 2.5 km, 0.25 m/sec cross-track, (1 sigma). {Expected downtrack performance is dependent upon flyby altitude, velocity and time of last encounter navigation image, as well as ACS pointing knowledge performance.}

Validation/Evaluation Documentation Plans:

Complete and publish preliminary technology validation reports approximately 30 days after completion of a defined mission phase (e.g. Initial Checkout, 01 January 1999).

Required Resources from Technology area and/or DS1 Project:

Formal agreements between the DS1 Project and TMOD (Telecommunications and Mission Operations Directorate) have been made and are documented in the appropriate work package agreements and resource cost planner estimates/plans. Any deviations from these agreements will need to be addressed by the appropriate parties.

APPROVALS:

Autonomous Optical Navigation
(AutoNav)
DS1 Technology Validation Final Report

Appendix G
AutoNav—Extended Mission

AutoNav Team: J. E. Riedel, S. Bhaskaran, B.
Kennedy, S. P. Synnott, T. C. Wang, R. A. Werner

Radio Nav Team: B. Kennedy, S. Bhaskaran, J. Thomas

Copyright © AIAA 1997

1.0 EXTENDED MISSION OVERVIEW

Following the successful completion of the Main Mission in the summer of 1999, the Deep Space 1 (DS1) spacecraft began an Extended Mission on September 18, 1999. The main goals of this mission were the flybys of the comet Wilson-Harrington in January 2001 and the comet Borrelly in September 2001. By returning science data from these encounters, DS1 would demonstrate the scientific usefulness of the technologies validated in the Main Mission. It would also further validate the effectiveness of the ion propulsion system (IPS) and the onboard autonomous navigation system (AutoNav). Section 3.2.12 (Post-Braille Cruise Operations) of the AutoNav DS1 Technology Validation Report [1] describes the successful use of the IPS and the AutoNav to drive the spacecraft towards the first of its planned encounters. These two technologies performed their tasks flawlessly during the first two months of the Extended Mission.

Unexpectedly, the spacecraft stellar reference unit (SRU) failed on November 11, 1999. Without this, the flight team was required to leave the spacecraft in a Sun-safehold configuration until a replacement plan could be enacted. While in this state, it became clear that DS1 could not encounter both comets Wilson-Harrington and Borrelly due to the loss of nominal thrusting schedule (or the so-called deterministic mission burns) after the star tracker failed. The DS1 science team met in January 2000 and decided that DS1 should keep the original plan to encounter comet Borrelly in September 2001.

Replacing the SRU and successfully making it to Borrelly required making use of three of the original twelve technologies that were verified in the Main Mission: The MICAS camera, the imaging processing capabilities of the AutoNav and the IPS. This report will describe the roles played by the AutoNav image processing, maneuver

planning and encounter target tracking software in the post-SRU Extended Mission. The roles of the MICAS camera and the IPS will also be described, along with necessary changes made to the comet tracking software in AutoNav.

Following replacement of the SRU, a new, low-thrust trajectory was developed. This trajectory required the use of near-continuous IPS thrust in order maintain spacecraft attitude using IPS thrust vector control (TVC) instead of the reaction control system (RCS) for attitude control. This reduced the usage of hydrazine (the fuel used by the RCS) from tens of grams per day to grams per day. This need to conserve hydrazine was the result of expending large quantities of hydrazine during the extended safehold, and maintaining the spacecraft in an Earth-pointed configuration for high-rate data passes without the use of SRU.

Figure 1 shows the flight configuration of the DS1 spacecraft and the key hardware components used for SRU replacement during the Extended Mission. Table 1 shows the timeline of the Extended Mission.

The replacement of the SRU with the MICAS camera required changes to cruise navigation techniques and the nucleus-tracking software. Since scheduling optical navigation activities would increase the risk of losing celestial inertial reference, the optical orbital determination (OD) capabilities of the AutoNav could not be used during cruise, so radiometric OD would be required. Optical OD would still be used to support the approach phase of the encounter. The encounter with Borrelly required modifications to the tracking software that enabled it to estimate the biases and drifts in the inertial measurement unit (IMU), and provide updated pointing to the Attitude Control System (ACS) during the encounter.




Figure 1. The DS1 Spacecraft

Table 1. Timeline of the Extended Mission

Date	Event
09-18-99	Start of Extended Mission
11-11-99	Loss of SRU
02-99 - 05-00	Redesign, Development, and Testing of SRU Replacement
06-06-00	Upload and Subsequent Restoration of Celestial Inertial Reference
07-00 – 04-01	Deterministic Thrusting
07-16-00 – 07-19-00	Loss of Star Lock and Recovery
11-00	Ka-band Experiments and Solar Conjunction
03-13-00	Upload of Encounter Flight Software and Recovery of Star Lock
05-01 – 09-01	North-South Thrusting
07-15-01 – 07-24-01	Loss of Star Lock and Recovery
08-16-01 – 08-24-01	Loss of Star Lock and Recovery
09-13-01	Loss of Star Lock and Recovery
09-22-01	Borrelly Encounter

2.0 SRU REPLACEMENT

2.1 Need for a Replacement

Without the SRU, the ACS lost the only instrument capable of providing it with inertial attitude quaternions every 0.25 s [2]. This left the ACS with an inertial measurement unit (IMU—the solid-state gyro) and a coarse (0.5 degree) sun sensor assembly (SSA). The IMU was effective at providing spacecraft rate information, which could be integrated to provide attitude, but it was too noisy and unstable to provide a reasonable attitude estimate for more than a few hours. The SSA could be used to keep an accurate fix on the direction to the Sun, but not the spacecraft rotation around that vector. Therefore, measurements from these systems alone would not enable the ACS to sustain a full 3-axis attitude estimate for more than a few hours, far too short to

support lengthy IPS thrust arcs. Please see Reference [3] for an in-depth description of these challenges.

2.2 MICAS as a Star Camera

As the only other optical device onboard DS1, the MICAS camera would become the new *de facto* star camera. AutoNav would be used to process the MICAS images in order to extract the star locations needed by ACS. Due to the small usable field of view (FOV) of the MICAS camera (effectively $0.5^\circ \times 0.75^\circ$ [2], as compared to the $8^\circ \times 8^\circ$ FOV of the original SRU and magnitude limitations (6.0 or brighter), only a single star would be tracked at a given time. Another stellar reference would be needed and was readily available as measurements from the coarse SSA. Since the MICAS camera and the SSA were pointed along orthogonal spacecraft axes, their measurements would provide a strong relative geometry that a new ACS could estimate and control the spacecraft attitude. The ACS would also be able to estimate the current biases and drifts within the IMU that would have to be relied on to maintain correct inertial attitude during turns. With this in mind, a new attitude estimator and a new image-processing manager were written.

By this point the actual image processing software was already in existence. During the beginning of the Extended Mission, work was already underway to develop new software that would be used during the upcoming comet encounter. This software, affectionately called "the Blobber", was designed to search through a specified area of a MICAS image and return a list of any contiguous "blobs". It was expected that these blobs would represent the nucleus of the comet and that additional software could be used to rectify and extract appropriate targeting information for the nucleus tracking software (see Section 5.6 and Appendix

H.) In the context of star identification, it served as a fast means of extracting the pixel and line locations of potential star candidates that needed to be passed along to the ACS.

2.2.1 *New Interfaces for AutoNav Image Processing*—To effectively use the AutoNav image processor for the SRU replacement, three new interfaces were developed between

- Ground and ACS
- ACS and Nav
- ACS and Nav by way of the MICAS camera.

2.2.1.1 *Ground to Nav via ACS*—An ACS storage facility, called a parameter set (PSET) array, was used. It was decided that the old method of configuring AutoNav software by adding new parameter files or expanding old ones would be cumbersome (see Section 2.4.2 of the AutoNav Report [1]) because of the expected high frequency of updates needed in operations. Table 2 shows the entries in the Nav PSET array and their uses.

Table 2. Entries in Image Processing Configuration Array

Nav PSET Array	Uses
pix_start	Column at which the search software started looking for stars. This was typically set to pixel 10. Ignoring data that was too close to the edge of an image was preferred, since the optical distortion was quite prevalent (up to several pixels) near the edge of an image.
pix_end	Column at which the search software stopped looking for stars. This was typically set to line 1013 (out of 1023). See pix_start, above.
lin_start	Row at which the search software would start looking for stars in an image. This was typically set to line 300, which allowed the search software to ignore stray light artifacts that quite literally dominated the images at low sun cone angles (50 to 90 degrees).
lin_end	Row at which the search software stopped looking for stars. This was typically set to line 1013 (out of 1023). See pix_start, above.
ceiling	Maximum pixel signal that would be considered valid star data. This was intended to be used to filter out saturated pixels that might be the result of cosmic ray strikes. This was set to 4000, out of a maximum signal of 4095. In practice, this sometimes resulted in valid signal from particular bright stars being thrown out by the star search software.

Nav PSET Array	Uses
floor	Minimum pixel signal that would be considered for valid star data. This was the key to the performance of the star tracking software. This was set to be 40, which allowed the star search software to ignore the background noise that was prevalent in the images, even after background processing. This allowed valid, potential star signals to be sent to ACS without flooding the ACS Star Identification software with false signals.
ceiling_noisy	This was the maximum value for the ceiling (see ceiling, above) used when background processing was not performed. In practice, this was set to 4000, but it was almost never used in flight.
floor_noisy	This was the minimum value for the floor (see floor, above) used when background processing was not performed. In practice, this was set to 100, but it was almost never used in flight.
blob_boundary_ext	Part of the statistical analysis that was performed to identify a star magnitude relied on a sampling of the background noise. This was used to compute the true signal coming from the star, minus the background interference. The boundary extension was the distance from a star “blob” around which a sample box was circumscribed. The average of pixel values that defined the edges of this box was used as the average background value.
verbosity	Turn on (or off) event reporting during star search processing. This was to allow diagnostic evaluation of the performance of the software when necessary.
acs_filter_width	This defined the maximum width of a star signal in the image, in pixels. In practice, this was set to be 200 pixels. It was intended to be used to filter out large areas that might be stray light artifacts and not true stars. At low cone angles (45- 50 degrees), large stray light artifacts would show up in the middle of the image. This filter was an attempt to prevent them from being mistaken as star signal.
acs_filter_height	This defined the maximum height of a star signal in the image, in pixels. In practice, this was set to be 100 pixels. See acs_filter_width, above.
fg_pic_bias	During picture background processing, a small bias was applied to the foreground image before the background image was subtracted. In flight, this was typically set to 10 DN.

2.2.1.2 ACS to Nav via MICAS—ACS PSET arrays are settable by a single ground command. An additional ground command was used to cause ACS to transfer the array information to the main Nav Task. This ACS-Nav interface required an additional queue interface to be added to the Nav Task. During tracking operations, the ACS task would initiate an image by directly sending image exposure commands to the camera manager, with the request that they be passed to Nav following the exposure. The extended image command interface developed for Nav’s use in the Main Mission was used, since it allowed for user-defined data to be added to the header of the resulting image. This user-defined information provided image handling, routing and processing information to Nav and needed bookkeeping information to ACS. Four image types were handled by Nav:

- Background images
- Solo images
- Parts 1 and 2 of a pair of images.

When Nav received a background image, it was placed in a buffer for later application. ACS routinely requested that background images be taken every 1/2 hour. This was intended to ensure that the background image was replaced often enough to track subtle changes in the stray light signature of the MICAS images (see Section 2.5.1 of the AutoNav Techval report [1]).

Image pairs were shuttered back-to-back and sent to Nav for processing with the intent that persistent star data would show up in each image, but not transient signals from cosmic rays or other interference. This would allow ACS to sort the "wheat from the chaff" and converge on a stable attitude solution. Solo images were requested once ACS had decided that it was receiving a consistent, identifiable star signal. Over 99 % of images taken for star tracking purposes were of the solo type.

Images could also be of a certain class: backgroundA, backgroundB, or no background. The image class type controlled whether background processing would be applied to an image before processing. Although it was intended to use "background free" processing as a means of increasing throughput, in practice this was not necessary. Nearly all images used for tracking underwent background processing to remove the static stray light signatures from the MICAS images.

Table 3 shows the handling definitions and values used during the extended mission.

Table 3. Image Routing and Handling Definitions

Name	Value	Description
Image Type		
IMAGE_BKG	(0x8000)	Indicates that this picture is to be stored in the background image buffer for use in future background processing. This was used as a means of removing most of the noise from stray light artifacts. Images of this type would be of DIFF_CLASS_A or DIFF_CLASS_B (see Image Class, below).
IMAGE_SOLO	(0x8001)	This is the nominal image type. ACS
IMAGE_PART1	(0x8002)	This is the first of two back-to-back images. These images are shuttered within two seconds of each other as way of letting the ACS star identification software discard spurious signatures that might be the result of cosmic ray interference. It also allowed it to identify consistent star signatures, which it needed before declaring that it had locked onto a star.
IMAGE_PART2	(0x8003)	This is the second of two back-to-back images. See PART1, above.
Image Class		
DIFF_NOTHING	(0x8000)	Images of this class did not undergo background processing. In practice, pictures of this class were rarely shuttered.
DIFF_CLASS_A	(0x8001)	Images of this class were to undergo background processing using a background image that was of class "A". If the image in the background was not of type A, ACS would be alerted, and a new background image would be shuttered.
DIFF_CLASS_B	(0x8002)	Undergo background processing with class "B". See CLASS_A, above.

2.3 Star Selection

The key to effective use of the new software was the careful pre-selection of a known reference star, also known as a "lock star." With *a priori* knowledge of where the spacecraft should point the camera for Earth communications or for IPS burn arcs, suitable stars were chosen from a star catalog. These stars were dubbed "Earth stars" and "thrustars," respectively. Over the course of the Extended Mission, it was noted that stars of magnitude 4.0 or brighter were ideal for use as reference stars. Stars of magnitude 5–6 could also be used if they were a "red" spectral type, such as a class-M, since CCD detectors tend to be more sensitive to red. The weak signal from stars less than magnitude 6 could not be relied on for tracking purposes as the tracking software required consistent inputs

to maintain a reliable lock. Due to these magnitude constraints, stars at suboptimal locations occasionally had to be used for inertial attitude reference, with a corresponding loss in thrusting effectiveness for thrusters and a reduced communications bandwidth capability for Earth stars. Once a reference star was selected, its inertial right ascension and declination would be told to the new ACS, which could then use the reported star location within the frame of the image to finely tune its estimate of the attitude.

2.4 Testing Efforts

Extensive testing of the new flight software before it could be approved for upload was necessary and somewhat difficult. Unit testing of the newly rewritten modules was used to evaluate and debug problems in a controlled environment. Performing fully integrated system level tests, required additional changes to the flight software, the DS1 instrument/hardware simulation software, and how they interacted with the flight system testbeds.

2.4.1 Additional Flight Software Changes

Along with the image processing and handling software, AutoNav contained internal image simulation software that could be used to produce an appropriate image of star fields, asteroid/comet bodies and cosmic ray noise (see Section 2.4.2.6 of the original AutoNav Tech Val report). When active, it would intercept an image that was being transferred from the camera to the AutoNav software and perform one or more of the following tasks:

- Remove any previous signal in the image, effectively allowing the simulation software to start from a clean canvas, as it were.
- Based on knowledge of the spacecraft attitude quaternion, it would determine the direction that the camera was pointing and search through an onboard full sky star catalog in order to determine what stars, if any, should be visible within the camera field of view. The original star catalog used during Main Mission only covered an area of the sky within 30 degrees of the ecliptic. This was necessary to conserve space in the onboard file systems. For these tests, a full-sky catalog was needed, and one was developed with a lower (brighter) maximum magnitude in order to stay within the bounds of the file system. Once a set of stars was queried from this new catalog, their locations within the image were computed, and the pixels at these locations were brightened appropriately according to the perceived magnitudes of the stars. The expected signal was also spread across one or more pixels depending on the camera optical characteristics as well as any perceptible motion in the camera due to high spacecraft rates (the spacecraft inertial rate information was available along with spacecraft attitude quaternion).
- Based on the same spacecraft attitude knowledge, this simulation software would determine where in the camera field of view the target comet would appear

based on the relative direction to the comet from the spacecraft at the current time. If this location existed, the comet would be painted at that location and would be the appropriate size in pixels based on the simulated radius (in kilometers) and the distance in kilometers to the comet from the spacecraft.

In order to produce realistic images for the image processing software, the image simulation software would need to be aware of the true spacecraft attitude, not just the spacecraft attitude as estimated by the ACS. This is covered in the next section.

2.4.2 Changes to the Instrument and Hardware Simulation Software

The instrument and hardware simulation software (SIM) needed to be upgraded in order to increase the fidelity of the system integration tests. A new noise model was developed for the Inertial Measurement Unit so that it more accurately modeled the inaccuracies and behaviors observed in the spacecraft IMU. Also, since this simulation software maintained an accurate model of the spacecraft truth attitude (in order to provide inputs to the SRU when it was working), this knowledge could be used by the internal AutoNav simulation software. The old SIM-SRU message interface was overhauled, and a new opmode was created such that the truth quaternion could be sent to ACS, which intercepted it and passed the information along to the micas camera manager. Since the ACS only needed one of the four packets sent from the SIM to ACS, this interface would not prevent the old SRU model from functioning. This was important, since the functioning SRU SIM model was needed to bootstrap the testbed initialization procedure during the early stages of software development and test in early 2000.

Once this quaternion was transferred from the SIM to the ACS and into the Camera Software manager, the manager would embed it into the image data of any nav-bound image. The AutoNav image simulation software would then extract it from the image information before it began the image construction (see Section 2.4.1).

2.4.3 Increased Testing Capabilities

With these additional hooks in place, the flight system testbeds were able to provide an appropriate test platform from which to observe and tune the performance of the tracking software once it achieved a steady state. This increase in fidelity allowed the flight team to assess the expected performance during many flight scenarios, including:

- Steady state attitude control: In a steady state, the new ACS software was required to maintain sufficient attitude knowledge such that the tracked star remained present in the camera field of view. It was also required

to calibrate the current biases and drifts within the spacecraft IMU.

- Attitude transitions: During transitions the spacecraft IMUs would be expected to provide accurate rate measurements to the ACS while the spacecraft was commanded to turn until camera boresight faced a new lock star.
- Post-safing recovery operations: Following a safing event, a reboot, or a loss of celestial inertial reference (see Section 4.4), the spacecraft's knowledge of its inertial attitude would typically be incorrect. Inferring and updating of this attitude knowledge required extensive analysis of spacecraft data and development of processes and sequences that enhanced data acquisition.
- Encounter sequence tested: During an encounter scenario, the ability of the nucleus tracking software to properly update the changing pointing requirements to the ACS could be tested using image data that was completely independent of onboard knowledge.

2.5 Initial Recovery Efforts

The crucial first steps were to determine where the spacecraft was pointing, update its knowledge of its inertial attitude, command it to point towards a known reference star, and activate the tracking software. Due to the fairly volatile nature of the IMU, this was expected to take at least several hours. It was unclear how robust the star tracking capability would be while a star was being tracked. There was considerable concern that in the event of a star tracking failure, the IMU might drive the spacecraft off attitude (and consequently off course if the IPS were thrusting) before the next tracking pass. It was thus expected that ground-directed attitude recovery efforts might become an operational norm.

The maintenance of high-gain antenna (HGA) pointing on the spacecraft in the absence of an SRU is covered in the section entitled "Earth Pointing: The Hard Way" of reference [2]. During attitude recovery operations, this technique would be used to maintain the spacecraft orientation in order to maintain the high-rate communications required for recovery operations.

3.0 TRAJECTORY PROFILE

Before flight testing and thrusting the ion propulsion engine in June of 2000, the DS1 engineers had been designing and planning trajectories for comet Borrelly encounter without star tracker. There were a lot of iterations of the solar electric propulsion (SEP) thrust profile between the Mission Design Team and Navigation Team.

3.1 Wilson-Harrington

Before the loss of the SRU, the original encounter plan for the extended mission had itself been extended to include a

flyby of the comet Wilson-Harrington in January 2001. However, reaching this target would have required thrusting to resume in January 2000. The aforementioned efforts to replace the SRU precluded this from happening. It was therefore decided early in the SRU recovery phase of the mission that a Borrelly-only trajectory would be needed.

3.2 Hydrazine

Hydrazine is the propellant used by the RCS, which is used by the ACS to maintain the spacecraft attitude using the Z-axis- and X-axis-facing thrusters (see Figure 1). However, during the period of time between the loss of the SRU and the restoration of attitude control (over half a year), a large amount of hydrazine was expended maintaining the spacecraft in its safing configuration and maneuvering the spacecraft during HGA communications with the Earth [2]. The remaining mass of hydrazine (approximately 9 kg of the original launch load of 32 kg) needed to be used very sparingly over the next 16 months. Fortunately for the mission, the ACS was able to control the X- and Y-axis attitudes using TVC whenever the IPS was running at a high enough throttle level. This would greatly reduce the duty cycle on the RCS and the usage of hydrazine. TVC is made possible by the thruster being mounted on two gimbals that allow up for ± 5 degrees of slew in the X and Y directions [2]. It was required that the IPS be active for most of that time in order to stay in TVC mode. The limited amount of remaining hydrazine had a large impact on trajectory design and maintenance as DS1 made its way towards Borrelly. To take advantage of TVC as a means of conserving hydrazine, a low-thrust trajectory was needed in which the IPS would be almost continuously active.

3.3 Trajectory Design

With DS1, this initial trajectory was designed to maximize IPS ontime in order to make use of TVC. This trajectory called for 10 months of deterministic thrusting, followed by a 4.5-month ballistic arc before the encounter with Borrelly; this was done to maximize the probability of a Borrelly flyby, allowing time for a possible mission recovery even in the event of an IPS failure. This trajectory plan called for thrusting to resume in early July 2000. The successful operation of the new SRU-replacement software allowed thrusting to resume in late June, one week earlier than expected.

The processes of designing and planning a trajectory to encounter comet Borrelly are described as follows:

1. A computer program named SEPTOP (SEP Trajectory Optimization Program) was used to design the DS1 trajectory at JPL. This program performs a constrained optimization of the propellant (xenon gas) consumption, target encounter time, and the deterministic IPS thrust direction and duration as a function of time. The constraints include adjustments to

the use of hydrazine, forced coasting (no IPS thrusting), and forced thrusting in specific directions (such as stars), cone angle constraints (i.e., restrictions to the thrust direction with respect to the Sun-spacecraft line) so that the radiators and the sensitive instruments will not be pointed close to the Sun, etc.

The results of the SEPTOP outputs were used as starting conditions to the NAVTRAJ (NAVigation TRAjectory) program described in the next step. It is worth pointing out that NAVTRAJ was an integral part of AutoNav during the Main Mission.

2. NAVTRAJ, a numerical integrated trajectory program with high-precision dynamic models, was used to retarget the trajectory based on the results of the optimal SEPTOP trajectory, also called the nominal trajectory. The NAVTRAJ and SEPTOP programs have the same spacecraft power and propulsion models. NAVTRAJ has the ability to make changes in the direction and duration of each thrust segment as defined in the IPS thrust profile (see below for details). It is assumed that the changes in the NAVTRAJ trajectory and the IPS thrust profile are relatively small in comparison with the results from SEPTOP. If there were significant changes, then it would be required to redesign a new SEPTOP trajectory for input to NAVTRAJ. This process is iterated until a converged NAVTRAJ trajectory is obtained. Most of the NAVTRAJ input files are generated by a utility program called SEPPROF (SEP thrust PROFile) that reads the SEPTOP outputs and then generates files for input to NAVTRAJ. The NAVTRAJ input files are described as follows:
 - a) Maneuver File: This file defines the IPS thrust profile. The thrust profile is divided into a sequence of planning cycles containing either IPS thrusting or coasting. In each IPS plan, a duty cycle value is used to specify the ratio of engine “on” vs. “off” time, where “off” time is primarily for telecommunications and autonomous navigation operations. Before the loss of the star tracker, the nominal duration of each planning cycle was 7 days and a duty cycle of 92% was used for the DS1 mission operations. Some planning cycles are shorter due to the operational constraints such as TCMs, close encounter events, etc. The thrust profile may contain several IPS segments (or thrust arcs). Each individual IPS segment is defined as a combination of consecutive IPS plans where the IPS is thrusting continuously except the imposed duty cycle. During the comet Borrelly operations, the design of duty cycle and SEP plans was driven by the DSN (Deep Space Network) tracking schedule.
 - b) OD File: This file includes the starting spacecraft epoch state and covariance for each planning cycle. It can be generated by either SEPPROF or another utility called the ODFILE program. In general, if the epoch state in the OD file is the same as SEPTOP, NAVTRAJ is used to generate flight products (including a trajectory) for upload on the spacecraft. If the OD file contains the current OD solution which is different from SEPTOP, NAVTRAJ is used either to generate a new set of flight products if the deviations from the nominal trajectory are small, or to show that a redesign of a new SEPTOP trajectory is needed if the deviations from the nominal trajectory are significantly large.
 - c) Xenon Mass File: This file contains the estimated available xenon mass as a function of time according to the nominal IPS thrust profile.
 - d) Hydrazine Mass File: This file contains the estimated hydrazine mass as a function of time based on the predicted ACS activities.
 - e) Control File: This file contains the target conditions, gravitation and solar pressure models, spacecraft dry mass, and spacecraft power model. The spacecraft power model is derived directly from the SEPTOP outputs. At a given time, the total spacecraft mass is defined as the sum of spacecraft dry mass, xenon mass, and hydrazine mass.
 - f) Spacecraft Propulsion System File: This file contains a table of the SEP thrust and xenon mass flow rate as a function of power. NAVTRAJ uses this file directly. However, SEPTOP uses the weighted least-squares fits to the table using 4th order polynomials, which produces good approximation for a given power range.
3. A MATLAB utility called THRUSTAR was used to select a set of sufficiently bright stars for use either as the thrust directions for SEP thrusting or Earth-pointed directions for telecommunications. The processes of selecting stars were very complicated and required an iterative procedure to obtain a trajectory (usually not optimal) to arrive at the desired B-plane target conditions. In general, the selection of a thruster is based on the star brightness and color, its angular distance from the Sun, and its location near the optimal thrust directions as derived from SEPTOP. Occasionally, if a thruster could not be obtained near the optimal thrust direction, then the thrust direction was vectorized to select several thrusters to achieve the desired thrust direction. When a desired trajectory was obtained, the locations (right ascensions and declinations) of stars were then implemented in the maneuver file to replace the SEP profile generated by NAVTRAJ. Due to the Sun cone angle constraints, a single thruster was usually locked on by the camera to maintain the spacecraft's attitude for a period of a couple of weeks. Therefore, each individual SEP thrust segment may require several thrusters. As a result, the trajectory is not an optimal one. However, it is the best

available trajectory which is designed to arrive at the comet Borrelly.

4. The initial selection of a set of thrusters was based on the optimal thrust directions derived from SEPTOP. The locations of the thrusters were then implemented in the input maneuver file. A MATLAB utility program IPSTARGET was used to target the trajectory to arrive at the desired B-plane. IPSTARGET first calls a subset of the NAVTRAJ C program to compute nominal B-plane coordinates at encounter, and then perturbs the trajectory to compute the B-plane partial derivatives with respect to duration of thrusting on each star in order to retarget the trajectory at the desired B-plane by adjusting that duration. Similar to NAVTRAJ, IPSTARGET has a capability to make changes in the direction and duration of each thrust segment as defined in the maneuver file. Also note that IPSTARGET uses exactly the same input files as these of NAVTRAJ. The strategy used for IPSTARGET was to change the direction and duration for the first few thrusters (usually one or two) at the beginning of the thrust profile or the thrust segments of interest, and hold the rest of the thrusters as fixed IPS TCMs. After the desired thrust directions were computed, THRUSTAR was used again to select new thrusters as described in the step (3). This process was iterated until the best available trajectory was obtained. If a large deviation from the nominal trajectory occurred as a result of new OD solution, then the processes in steps (3) and (4) were used to redesign a new trajectory instead of going back to SEPTOP. Note that most of the DS1 Borrelly trajectory designs used the THRUSTAR/IPSTARGET interfaces instead of the SEPTOP/NAVTRAJ interfaces.

3.4 Implementation in Operations

The burn profile design methods described above took into account the need for IPS thrusting during Earth passes. These thrustings were constrained to attitudes that allowed the fixed-boresight HGA to point at the Earth during times when a DSN antenna was scheduled to track DS1 and downlink data at a high rate. During these Earth passes (typically eight hours long), the IPS throttle level was set to a low level (approximately 22.4 mN) which still allowed attitude control using TVC. Although this low throttle level minimized the impact on the DS1 trajectory, it still needed to be modeled in order to provide a targeted burn profile.

Following the creation of a nominal thrust profile, the flight sequencing team integrated the orientations and thrust levels into the backbone sequence. Typically, a backbone sequence is a single, absolutely-timed sequence that runs on the spacecraft for several weeks. This sequence controls a majority of the routine spacecraft operations, including (but by no means limited to) telecommunications configuration, operational spacecraft reorientation, star tracking software management, and IPS thrust-level management.

Telecomm configuration is based on the scheduled DSN antenna and the expected off-Earth angle of the HGA boresight. Typically, if the boresight could be pointed to within a degree of the Earth it would enable the maximum supportable data rate. Figure 2 shows a heliocentric view of the Sun, Earth, and spacecraft configuration while in Earth point. During Earth communications, solar panel pointing requirements (“SCARLET Solar Array” [1]) constrained the spacecraft to be either prograde or retrograde thrusting within the plane of the ecliptic. This resulted in a limited set of stars that could be tracked. If the nearest tracking star was suboptimally located, it would result in a decreased supportable data rate.

The sequencing of an attitude transition was fairly straightforward. First, the tracking software would be commanded to stop tracking. Next, the IPS would be turned off, and the spacecraft would be commanded to turn to a new attitude. Since the biases and drifts of the IMU were well calibrated by the previous time spent locked to a reference star, these turns were executed using the IMU as the means of attitude propagation. Without exception, these turns completed with the spacecraft in the desired inertial attitude. Once at this new attitude, the tracking software would be told the magnitude and inertial location (right ascension and declination) of the star that it would expect to see when it started tracking. It would also be told what exposure duration to use for camera commanding. It would then be told to start commanding the camera, at which point it would start receiving star signals from the camera by way of the AutoNav image processor. Shortly afterwards, the IPS was commanded to re-pressurize the plenum, and start




Figure 2. Earth Point Geometry

thrusting. Finally, the ACS would be commanded to start controlling the spacecraft attitude using TVC. The timing of all of these commands was set to allow for a nominal time to pass before progressing to the next command. In other words, the turn was expected to be complete before the star tracking software was enabled, the IPS wasn't turned on until the plenum had re-pressurized, and the ACS wasn't put into TVC mode until the IPS reached a steady state.

3.4.1 Deterministic Thrusting—This is the nominal burn configuration. For 10 months following the replacement of the SRU, the spacecraft was expected to be thrusting deterministically towards an encounter with Borrelly. Table 4 shows a 2-month segment of the burn profile followed by DS1 during of this deterministic thrust period. Typically, odd-numbered arcs in this table represent week-long thrust arcs, while even-numbered arcs represent short 6- to 10-hour burns at which time the spacecraft was aligned to point the HGA at the Earth.

3.4.2 North-South Thrusting—In order to achieve an effectively ballistic trajectory during the last 4 months of the

cruise phase, a burn profile alternating approximately ecliptic north thrust attitudes with approximately south attitudes was used. Adjustments to the nominal north-south burn directions were made to account for thrusting during telecommunications sessions and for deviations from exactly north/south attitudes. Table 5 shows a list of north-south thrust attitudes in the months before the encounter.

3.4.3 Cone Angle Constraints—Due to stray light problems with the camera [2], [3] spacecraft orientations during which the camera boresight was within 45° of the Sun were not allowed, as a flight rule. Theoretically, this constraint prevented certain thrust attitudes from being realized, requiring “vectorization” of a desired thrust arc. In practice, this was not needed during cruise, nor during IPS and RCS TCMs. However, the possibility of having to do so was realized and plans to vectorize TCMs at encounter were developed.

Table 4. Thrust Profile Containing Deterministic Thrust Arcs

Num	Start Day	Start Time	Duration (days)	Right Ascension of Lock Star	Declination of Lock Star	Thrust Level	TBD	TBD
1	01-DEC-2000	02:00:00.0000	4.503	17.148	-10.182	0.022410	4.549	21
2	05-DEC-2000	15:10:00.0000	0.399	339.300	-8.940	0.022410	0.417	22
3	06-DEC-2000	01:10:00.0000	6.524	17.148	-10.182	0.022410	6.590	23
4	12-DEC-2000	15:20:00.0000	0.221	344.410	-6.890	0.022410	0.243	24
5	12-DEC-2000	21:10:00.0000	9.697	12.171	7.585	0.022410	9.795	25
6	22-DEC-2000	16:15:00.0000	0.268	349.500	-4.770	0.022410	0.285	26
7	22-DEC-2000	23:05:00.0000	10.620	12.171	7.585	0.022410	10.674	27
8	02-JAN-2001	15:15:00.0000	0.371	359.640	-0.410	0.022410	0.389	28
9	03-JAN-2001	00:35:00.0000	7.532	22.871	15.346	0.030525	7.608	29
10	10-JAN-2001	15:10:00.0000	0.210	4.720	1.800	0.030525	0.240	30
11	10-JAN-2001	20:55:00.0000	5.696	22.871	15.346	0.031126	5.753	31
12	16-JAN-2001	15:00:00.0000	0.395	9.650	4.400	0.031126	0.451	32
13	17-JAN-2001	01:50:00.0000	7.470	38.969	5.593	0.031126	7.545	33
14	24-JAN-2001	14:55:00.0000	0.213	14.810	6.580	0.031126	0.243	34
15	24-JAN-2001	20:45:00.0000	5.696	38.969	5.593	0.031126	5.753	35
16	30-JAN-2001	14:50:00.0000	0.444	20.040	8.710	0.031126	0.507	36
17	31-JAN-2001	03:00:00.0000	6.597	33.250	8.847	0.031126	6.663	37
18	06-FEB-2001	18:55:00.0000	0.204	25.330	10.770	0.031126	0.233	38
19	07-FEB-2001	00:30:00.0000	9.488	33.250	8.847	0.031727	9.583	39

Table 5. A List of Stars/Magnitudes/Locations Used as Reference Targets to Maintain a Converged "Ballistic" Trajectory During a Segment of North-South Thrusting

Num	Start Day	Start Time	Duration (days)	Right Ascension of Lock Star	Declination of Lock Star	Thrust Level
1*	24-MAY-2001	00:45:00.0000	5.442	276.496	65.563	0.022410
2**	29-MAY-2001	12:00:00.0000	0.510	121.941	21.582	0.022410
3*	30-MAY-2001	02:00:00.0000	6.468	92.812	-65.589	0.022410
4**	05-JUN-2001	14:00:00.0000	0.419	128.177	20.441	0.022410
5*	06-JUN-2001	01:30:00.0000	13.639	276.496	65.563	0.024213
6	19-JUN-2001	18:30:00.0000	6.634	138.808	14.942	0.024213
7	26-JUN-2001	10:30:00.0000	7.566	328.325	-13.552	0.022410
8*	04-JUL-2001	01:00:00.0000	6.364	92.812	-65.589	0.029924
9**	10-JUL-2001	10:30:00.0000	0.346	151.976	9.997	0.022410
10	10-JUL-2001	20:00:00.0000	4.975	263.748	61.875	0.022410
11	15-JUL-2001	20:00:00.0000	8.893	273.475	64.397	0.025114
12	24-JUL-2001	18:30:00.0000	0.401	166.254	7.336	0.022410
13	25-JUL-2001	05:30:00.0000	5.514	92.812	-65.589	0.029924
14	30-JUL-2001	18:30:00.0000	0.365	166.254	7.336	0.022410
15	31-JUL-2001	04:30:00.0000	14.137	92.812	-65.589	0.025415
16	14-AUG-2001	09:30:00.0000	0.528	177.674	1.765	0.022410
17	15-AUG-2001	00:00:00.0000	6.529	276.496	65.563	0.022410

*Arcs 1, 3, 5 and 8 show examples of alternating, self-canceling north-south arcs.

**Arcs 2, 4 and 9 show stars used to allow alignment of the HGA on Earth.

4.0 NAVIGATION OF A LOW-THRUST MISSION WITH RADIO OD

In order to effectively determine DS1's orbit using only radio data, the original methods laid out for navigating the spacecraft under low thrust needed to be modified to match the changed conditions under which the spacecraft would operate. Due to a reduction in the frequency of high-rate and low-rate tracking passes, a decreased availability of range and Doppler data during the Extended Mission was expected. Also, the original methods for modeling the spacecraft IPS and RCS activity needed to be modified to account for data that might no longer be correct. As it turned out, this reduction in tracking data and model fidelity required a change to the filtering strategy.

4.1 Data Types

As with all missions, radiometric data is acquired during tracking passes using the various antennas at the DSN complexes at Goldstone, Canberra, and Madrid. For DS1, conventional Doppler and range data were acquired during tracking passes. Differenced-Differential One-way Range (DDOR) data acquisition was not planned during the cruise phase of the extended mission. Its use in the approach phase of the mission is discussed in Section 4.

4.1.1 Earth Passes—During a high-rate DSN pass, also called an “Earth pass,” the ground communicated with the spacecraft through the spacecraft HGA while the spacecraft was at an Earth-pointing attitude. There were only three Earth passes scheduled per month, on average. This was necessitated primarily by a need to limit attitude transitions. Earth passes typically required a transition before the beginning of a track in order to align the HGA with the Earth and a transition back to a nominal burn attitude following the track. Turning the spacecraft is expensive from a hydrazine standpoint and was considered potentially risky from an attitude knowledge standpoint, given the nature of the tracking software. On the plus side, Earth passes were typically the only time at which ranging measurements to the spacecraft could be taken. Whenever possible, these passes were scheduled so that they spanned the handover between the Goldstone and Canberra complexes. This allowed for near-simultaneous north and south ranging data to be taken. As was discovered during OD validation in the Main Mission, estimating geocentric declination in low thrust trajectories benefits from the strong geometry provided by north and south range data.

As mentioned in Section 2.4, Earth stars were not always optimal with respect to HGA pointing. This often constrained bandwidth, and sacrificed ranging data in favor of downloading the weekly backlog of telemetry. If bandwidth was limited during a north track, operational

efforts were made to obtain range data at the end of the track to provide a stronger geometric correlation with the south range data. As was the case in the earlier phases of the mission, long-range modulation times were needed to prevent out-of-modulo range measurements in the event of missed thrust, or misthrusting. This reduced the amount of range data received.

4.1.2 Midweek Passes—During a low-rate communication session, also called a “mid-week pass,” the ground communicated with the spacecraft through one of the low-gain antennas (LGA) while the spacecraft was at a burn attitude. Due to the use of smaller DSN antennas and the fairly weak LGA, telemetry was rarely available, even at low bit rates. During these passes only a limited amount of Doppler (2–3 hours) was received, but it provided very strong visibility into the burn activity. This was very valuable to the OD and stood in stark contrast to the poor thrusting visibility during the Earth passes. With the absence of telemetry during these tracks, the Doppler signal provided rapid assessment evidence of the health of the spacecraft and its trajectory. With one exception, ranging data was not available during mid-week tracks. Many experiments were attempted with low-modulated ranging to attain data, but these met with mixed results.

4.2 Modeling

The primary spacecraft nongravitational perturbation models needed to navigate DS1 were solar radiation pressure (SRP), IPS thrusting, and RCS activity caused by turns and deadbanding. The SRP model was unchanged from that used in the Main Mission. The original methods for modeling the spacecraft IPS thrust arcs and RCS activity were slightly modified from those used in the Main Mission [4].

4.2.1 RCS Activity—The modeling of RCS activity induced by deadbanding and turns was somewhat simplified in the Extended Mission. Since no OpNav activities were performed, the nongrav file was no longer needed to estimate their effects on the trajectory. It is also worth noting that the occasional loss of attitude lock made the inertial measurements of the RCS activity untrustworthy. Therefore, a modeling scheme that relied on them was not used. However, the nongrav file was still of some use, as it did assist in the placement of impulsive burns that could be used to model the effects of turns by the spacecraft. It was especially useful with respect to modeling the impulse placed on the spacecraft when DS1 was mosaicking. Mosaicking is a set autonomous spacecraft turns, which DS1 underwent whenever it was trying to acquire (or reacquire) its lock star. Since the mosaic turns are so small, the overall effect of the spacecraft is somewhat akin to a mini-RCS TCM (i.e., a delta-V of several cm/s along the spacecraft +Z axis). Also, since many mosaic events occurred outside of a DSN track, a simplified, loose model

had to be used to estimate their impact. While the turn pulses themselves were small enough, they did have a large aggregate effect that needed to be taken into account.

4.2.2 IPS Activity—For IPS activity, a simplified thrusting model made use of the thrust history recorded in telemetry, and assumed that attitude was tied directly to the thruster direction. Due to thrusting uncertainties and approximate location of the star in the camera the true burn attitude was uncertain, so a simplified “use star direction to define burn attitude” strategy was used.

4.3 Filtering

Initially, the nominal pre-SRU loss radio Nav OD strategy was used for post-SRU loss OD. For the first few months using the new models, the solutions were very well behaved. However, subsequently, the OD performance began to degrade, exhibiting slow convergence, large stochastic ranges and multiple-sigma corrections to thrust magnitude and pointing (several mN and several degrees, respectively). It was determined that the filter was trying to extract too much information from the very limited amount of data available, so a simplified filter strategy was used with fewer variables and tighter sigmas (1 mN and 1°). Highly constrained stochastic accelerations were used to help smooth the resulting trajectory and to account for some of the uncertainty induced by the TVC activity and thrust measurements.

4.4 OD Impact During Loss and Recovery of Attitude

Following loss of inertial lock (LOL), inertial reference needed to be quickly restored. If inertial reference is not quickly restored, the bias and drifts of the IMU cause the spacecraft attitude to drift. Since DS1 was thrusting most of the time, this drift caused an ever-increasing divergence

away from the expected trajectory. Following attitude recovery operations, determining the new position and velocity of the spacecraft was of prime importance, since the future thrust profile had to be quickly corrected to keep the spacecraft on course for Borrelly. Once characterized, any velocity errors were accounted for by modifying future burn arcs. If a long time passes before velocity errors can be quantified, an uncomfortably large position error can build up. For example, if the spacecraft is miss-pointed by 20° for 5 days at full thrust, a velocity error of 8 m/s would accrue in a direction normal to the thrust vector. After this time, the position error would be 2000 km and would continue to increase by 5000 km per week. As the spacecraft neared Borrelly, quick evaluation of the LOL effects on DS1’s orbit became important as the planned trajectory was to be modified in a timely fashion. See Table 6 for attitude losses, time ranges, and causes.

4.5 A Case Study: LOL 5

In late August 2001, less than two months from the encounter with Borrelly, solar interference caused the camera to be flooded with false signals. These false signals caused the ACS software to drift away from its planned reference star as it chased the myriad false stars.

The resulting drift lasted two days, after which the spacecraft serendipitously found a real star to track. Recovery efforts began 5 days after the initial LOL at the start of what should have been a routine Earth tracking pass.

At this point in the spacecraft’s orbit, aligning the HGA with the Earth while the spacecraft thrusting was in a prograde direction required pointing the camera little more

Table 6. Attitude Losses, Time Ranges, and Causes

Start	End	Cause
06/12/00	06/12/00	Initial attitude recovery.
07/16/00T20:00	07/19/00T01:00	Solar interference with star observations.
03/13/01T16:00	03/16/01T2000	Planned reboot following FSW upload.
07/15/01T20:00	07/24/01T1800	Unknown, possible lock acquisition failure.
08/16/01T12:00	08/24/01T1100	Solar interference with star observations.
09/13/01T17:00	09/14/01T0100	Inability to acquire initial lock.

than 50° from the Sun. At this attitude, scattered light problems that troubled the camera since the start of the mission [2] were dominating the 3.5-s exposure images that were taken. This made the onboard centroid processing almost unusable, since the high number of false signals overwhelmed any star signatures. (At this phase in the mission, centroid data packets provided picture previews of images taken during recovery activities.) This increased the possibility of downlinking an image that contained an identifiable star field, by only selecting images known to contain stars of sufficiently bright magnitude to make identification likely.

The low Sun cone angle of the camera made attitude recovery operations very difficult, so it was decided to rotate the spacecraft a full 180° from a prograde to a retrograde attitude. This somewhat risky maneuver had two benefits. By flipping, the two and a half days of roughly prograde thrust were mostly canceled out by retrograde thrust. Also, the Sun was no longer able to interfere with camera images, allowing for deeper exposures to be taken. In order to take full advantage of this, the centroid sequences were enhanced to take 10-s exposures and also to run in a continuous loop. Following the flip, one large HGA corrective turn was performed just before the end of the current tracking pass. At the start of the first of two more borrowed passes, the new sequences were uploaded and activated. The new centroid packets contained vivid signatures of dim stars (down to magnitude 8), and provided enough indication of relative motion that a reasonable estimate of IMU drift could be derived. The deep images selected for downlinking proved immediately useful. Less than five hours into the pass, the spacecraft attitude was determined and corrected. The subsequent attempt to turn to and lock onto a suitable reference star was quite successful. Using the second of the two borrowed passes the flight team was able to prepare the spacecraft for its first observation of comet Borrelly, which was scheduled to occur less than twelve hours later.

Modeling all of this activity sufficiently to allow for a useful OD was difficult. Of key importance was identification of the star that the spacecraft had locked onto for the two and a half days before the sequenced turn to Earthpoint. Fortunately, the Nav Team successfully identified this star based on knowledge of its hypothetical location, and the presence of a small “companion” star which showed up periodically in the centroid data (see Figure 3). A simple model, consisting of five days of thrust on the now known star, three days of approximate prograde thrusting, and two days of retrograde thrusting, was developed. This enabled an immediate assessment of the effects on the trajectory. During the recovery period the attitudes of several burn arcs

and turn- ΔV s were estimated. Hypothetical spacecraft rates were approximated by looking at the observed change in locations of stars that appeared in centroid data. Figure 4 shows images from which a drift rate of 0.3° per hour can be determined. After a couple of days, a reasonable OD estimate was produced, and this enabled fine-tuning of the pointing and thrusting for the upcoming North burn arc. The preliminary OD showed that after the end of the recovery efforts, the spacecraft had a position and velocity discrepancy of 5600 km and 20.5 m/s from the nominal trajectory. After three weeks of post-recovery data, an overlap of this fit with an OD comprised entirely of post-recovery modeling showed an agreement of 300 km and 0.7 m/s. The resulting B-plane shift was 18,787 km in B•R, 27,568 km in B•T and 1,158 seconds in time of flight (TOF).




Figure 3. A recreated picture of one of the centroid data packets taken before recovery activities in LOL 5. It shows the 2.5 magnitude reference star that was locked onto. A 4.2 magnitude “companion” star is also visible, along with 11 false star signals caused by solar activity.




Figure 4. Centroid images taken 10 minutes apart. These images show three stars in the camera FOV, with magnitudes of approximately 4.5, 7, and 9. Other signals are stray light artifacts or cosmic rays.

5.0 APPROACH PHASE AND ENCOUNTER USING OPTICAL AND RADIO OD

5.1 Tracking During the Encounter

A closed-loop, onboard tracking system was used to find and maintain lock on the comet nucleus during the flyby. This software was an extension of the original AutoNav software, with an important enhancement: it was able to provide pointing updates to ACS that took IMU drift and bias into account. Since the MICAS camera would be used primarily to observe the comet during the encounter, maintaining attitude using a reference star would not be possible.

5.2 Comet Ephemeris Development

Due to the relatively large non-gravitational forces which act on comets (e.g., outgassing), predicting an accurate ephemeris for even short periods into the future can be quite difficult. Thus, even though ground telescopic observations going back several decades were available for Borrelly, an intensive campaign was undertaken to improve its ephemeris for the DS1 flyby. After its recovery in the sky during its current apparition in May 2001, over 200 observations were obtained from telescopes located at Loomberah, Australia, the United States Naval Observatory in Flagstaff, Arizona, and the Table Mountain and Palomar observatories located in southern California. The observations were processed by members of the Solar System Dynamics (SSD) group at JPL and delivered to the DS1 navigation team. In all, three deliveries were made: the first using just the ground observations and the last two using a combination of spacecraft and ground observations.

More details of the comet ephemeris development effort can be found in Appendix H.

5.3 File Upload Strategy

As during the Main Mission, the comet tracking software used files for configuration and setting initial conditions. Files containing the latest estimates of the spacecraft and comet trajectories were uploaded to the spacecraft before the encounter. This allowed the ephemeris server to provide the ACS with an appropriate *a priori* pointing direction. The parameters that characterized the expected response of the camera to the nucleus, coma, and stray light (background noise) were also uploaded. This was to improve the tracking software's ability to successfully identify the nucleus in the images.

5.4 Radio OD Delivery Accuracy

Even though the OD after LOL 5 looked stable, there was still some concern about unaccounted-for errors. The upcoming observations of Borrelly were expected to resolve some of this uncertainty. The observations taken in early September showed a 1000–1500-km difference between the predicted and observed locations of the comet. Figure 5 shows the results from the observation of Borrelly on September 10. The latest radiometric OD solution was used for the initial prediction of the comet within the camera FOV. At this distance to the comet (22 million km), each 13-microradian pixel spans 282 km. This placed the predicted location of the comet nucleus within 1100 km of where the images showed it to be. Over the first four observations, the position error between observed and predicted comet location was consistent, implying that no significant velocity errors remained from modeling LOL 5.




Figure 5. Left frame: Observed (+) vs. predicted (o) location of Borrelly using co-added images. Middle and right frames: registration performed on two stars seen in co-added images.

5.5 Borrelly Approach Using Radio OD

Determining the heliocentric orbital out-of-plane position errors, as well as establishing the validity of the OD, was accomplished with two DDOR observations taken on September 14 and September 15, one week before the encounter. The resulting OD showed close agreement (20–30 km) to the previous OD. As well as validating the out-of-plane results of the radiometric OD, they also provided a higher certainty on the predicted time of flight (TOF)— ± 3.3 s with DDOR, and ± 14 s without. After one more week of radiometric data, these TOF uncertainties changed to ± 3.5 s with DDOR, and ± 4.7 s without.

5.6 Ephemeris Rectification

Once the DDOR campaign showed that the radio OD was not a major source of error (see Section 4.4 and Section 4.5), efforts shifted to determining why the ground-based comet ephemeris did not agree with the spacecraft observations. Eventually, it was found that if the center-of-brightness computed from the ground observations used the brightest pixel, rather than the standard Gaussian fit to the brightness profile, the results agreed considerably better with the spacecraft. Furthermore, observations taken at Palomar Observatory and processed using the bright pixel method, now were in fairly good agreement with the spacecraft. Nevertheless, discrepancies still existed that were eventually attributed to the lack of an accurate model for outgassing used in the comet orbit estimates. Recently, it was found that an acceleration model that had jets at the assumed comet pole, and varying with the angle between the pole and the sun, resulted in the ability to fit longer data arcs from the ground when combined with spacecraft data. See Appendix H for more information.

5.7 The Borrelly B-plane and the TCM Strategy

The B-plane is a plane passing through the center of the target body and perpendicular to the incoming asymptote, S, of the hyperbolic flyby trajectory. Coordinates in the plane are given in the R and T directions, with T being parallel to

the Earth Mean Ecliptic plane of 2000. The angle θ determines the rotation of the semi-major axis of the error ellipse in the B-plane relative to the T-axis and is measured positive right-handed about S (see Figure 6).

The first of several IPS TCMs occurred on September 11, 2001. This TCM, 1.1, refined the B-plane targeting to place it near an area of the B-plane known affectionately to the Nav Team as the “Magic Control Line.” This line intersected the B•T axis at approximately 2000 km B•R. Its slope was defined as the direction in which the B-plane position was controllable by thrusting while the HGA was aligned with Earth (see Figure 7). Once there, the final targeting of the Borrelly flyby point was controlled solely by Earth-pointed IPS TCMs. This meant that no RCS TCMs were needed for the encounter, and little or no offpointing from Earth. Although there was a reserve of 2 kg of hydrazine for RCS TCMs, not having to use this provided much additional mission assurance, given the severe fuel shortage, especially when the large uncertainty in the remaining hydrazine was considered. Control of the B-plane was exercised in such a way as to arrive at Borrelly with B•T as close to 0 as possible. This was desired, as the encounter sequence was designed assuming Sun-relative geometry. (That geometry allowed the spacecraft to track the target with slews about the spacecraft “Y” axis while keeping the solar panels on the Sun.) Control of the final values of B•R and TOF were not as critical, although accurate knowledge of TOF was still necessary for mission success. It was also desirable to approach 0 B•T from the negative side, as the approach from this side could be controlled by throttling up during Earth telecommunications passes. There was limited ability to throttle down (the IPS has a minimum operable power) to achieve a relative backward motion along the control line, and completely shutting down the engine would have consumed vital hydrazine. If for any reason the spacecraft-comet B-plane shifted into positive B•T, corrective TCMs would have required that the spacecraft be reoriented into a prograde attitude, and this would have been a difficult, fuel-consuming, and dangerous maneuver.




Figure 6. Targeting at JPL is performed in the so-called B-plane coordinate system.




Figure 7. DS1 at Borrelly Encounter B-Plane

The second TCM, 1.3, was scheduled for September 14. Due to the response required by LOL 6, the TCM was cancelled. Originally, the spacecraft was intended to be placed on the magic control line by this TCM, but this was effectively accomplished by reorienting the spacecraft onto a previous Earth star. Following this cancelled TCM, the IPS was shutdown as previously scheduled. This allowed the spacecraft B-plane position to shift day by day, due to unmodeled RCS activity. TCM 2.1 occurred on September 17, at Earth-point orientation. This corrected the targeting to take into account the new updates to the Borrelly ephemeris.




Figure 8. Final DS1 Borrelly Encounter B-Plane

Following TCM 2.1, the spacecraft B-plane target moved closer to the desired aim point (Figure 8 shows the final encounter B-plane). The shifts in the B-plane location from September 18 to September 21 are based on daily OD solutions using optical data and multiple radiometric strategies (long arc, short arc, with and without DDOR, etc.). These shifts were caused by nongravitational impulses from RCS activity. These shifts were expected to occur, and are evident as the B-plane intersection moves “up and to the right, along the magic control line” (see Figure 8). On September 21 and 22 the last two TCMs, 4.1 and 4.2, were designed and executed to line up DS1 for its encounter with Borrelly. Both TCMs occurred at Earth-point orientation. Following their successful execution, it was the task of the nucleus-tracking software to autonomously command the pointing of the spacecraft and the execution of the close-in science sequences. A detailed description of the performance of this software can be found in Appendix H. On September 22, at 22:30:36 ET, DS1 flew past Borrelly at 2171.2 km in B•T, 31.2 km in B•R. This was 6 seconds earlier than predicted. The highest resolution image of the nucleus was obtained approximately two minutes before closest approach and can be seen in Figure 9.



Figure 9. Highest resolution of nucleus, taken during Deep Space 1 encounter with comet Borrelly.

6.0 REFERENCES

- [1] New Millennium Program Technology Validation Reports, <http://nmp-techval-reports.jpl.nasa.gov/>.
- [2] Collins, S., and others, “Deep Space 1 Flight Experience: Adventures of an Ion Drive,” 25th Annual AAS Guidance and Control Conference, Breckenridge, Colorado, February 6-10, 2002.
- [3] Rayman, M.D., and Varghese, P., “The Deep Space 1 Extended Mission,” *Acta Astronautica*, **48**, No. 5-12, pp. 693–7005, 2001.
- [4] McElrath, T. P., and others, “Radio Navigation Of Deep Space 1 During Ion Propulsion Usage,” 15th International Symposium of Space Flight Dynamics, Biarritz, France, June 2000.

7.0 ACKNOWLEDGEMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

The author would like to thank many members of the NASA community, whose efforts and inputs made the DS1 mission a very rewarding and successful experience:

- Tim McElrath, Don Han, and Marc Ryne for developing the DS1 radio navigation strategies during the Main Mission
- Marc Rayman, the DS1 Mission Manager, for having faith in the flight team to make the extended mission a success
- Steve Collins, Tony Vanelli, Sam Sirlin, Sanjay Joshi, and Jim Alexander for their attitude control expertise, the DS1 Flight Team
- The DS1 Flight Controllers, and the personnel of the DSN facilities at Goldstone, Madrid, and Canberra.

Autonomous Optical Navigation
(AutoNav)
DS1 Technology Validation Final Report

Appendix H
Navigation of the Deep Space 1
Spacecraft at Borrelly

Copyright © AIAA 2002

Navigation of the Deep Space 1 Spacecraft at Borrelly

S Bhaskaran, J. E. Riedel, B. Kennedy, T. C. Wang

Navigation and Flight Mechanics Section
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California 91109

Abstract

On September 22, 2001, the Deep Space 1 spacecraft flew by the short period comet Borrelly at a distance of approximately 2200 km. The navigation challenges posed by the flyby were considerable due to the uncertainty in the knowledge of the comet's ephemeris, as well as the difficulty in determining the spacecraft's ephemeris caused by relatively large non-gravitational forces acting on it. The challenges were met by using a combination of radio, optical, and interferometric data types to obtain a final flyby accuracy of less than 10 km. In addition, a closed-loop onboard autonomous tracking system was used to maintain lock on the comet nucleus during the flyby.

Mission Overview

The Deep Space 1 spacecraft was launched on October 24, 1998 as the first mission in the New Millennium Program. The purpose of this program was to fly a series of spacecraft whose goal was to test advanced technologies needed for future missions. Deep Space 1 carried 12 such technologies, including an ion propulsion system, an advanced solar array, and an autonomous optical navigation system. Following the successful completion of its primary mission on July 1999 (the flyby of the asteroid Braille), the spacecraft was approved for an extended science mission to fly by the short period comets Wilson-Harrington and Borrelly. Unfortunately, the onboard star tracker, used as the primary means of maintaining the spacecraft attitude, failed in November, 1999 and for the following 7 months, the spacecraft was placed in an extended safe hold configuration. During this time, new software and techniques were developed to enable the science camera to

function as a replacement for the star tracker. During this period, the thrusting needed to achieve the Wilson-Harrington rendezvous was unable to be performed and was therefore dropped from the mission plan. In June 2000, the software modifications were loaded onto the spacecraft and thrusting resumed to achieve the Borrelly flyby. Finally, in September 2001, the spacecraft flew by Borrelly at a distance of roughly 2200 km, obtaining the highest resolution images of a comet to date.

Due to the unorthodox process of using the narrow angle science camera as a substitute star tracker, the use of ion propulsion as the primary means of propulsion, and the uncertainties in determining precise ephemerides for comets, the challenges in navigating the flyby were substantial. This paper details the navigational techniques and procedures that were used to overcome these obstacles and achieve a successful encounter.

Spacecraft Overview

DS1 was the first interplanetary spacecraft to use solar electric propulsion as its primary means of controlling its trajectory. Its single ion thruster (referred to as the IPS) was capable of producing a maximum of 90 milliNewtons of thrust continuously over many days and weeks. In addition, standard hydrazine thrusters were available for attitude control around all three spacecraft axes, and for some course corrections. Power for the spacecraft was provided by the prototype solar arrays which generate 2.5 kW of power at 1 AU.

The primary science instrument onboard was the Miniature Integrated Camera and Spectrometer (MICAS), which had two visible, one ultraviolet, and one infrared imaging

channels. For navigation purposes, only one of the imaging channels, a standard Charge-Coupled-Device (CCD) chip with a 1024 square pixel array, was used. Each pixel had a field-of-view (FOV) of about 13 μ rad for a total FOV in the CCD of 1.3 mrad, or 0.76 deg. The CCD was coupled to a telescope with a focal length of 685 mm with the boresight fixed to the spacecraft (thus, the entire spacecraft had to be slewed to point at particular region of the sky). Also, the CCD had 12 bit digitization, resulting in data numbers (DN) values for each pixel ranging between 0 (no signal) and 4095 (saturation). This CCD also doubled as the substitute star tracker after the failure of the normal star tracker. In this paper, the horizontal measurement of an object in the frame of the CCD is referred to as its “sample” value, while the vertical is referred to as “line”.

Navigation Overview

Standard navigation data types used on DS1 included Doppler and range, which measure the line-of-sight velocity and position, respectively, of the spacecraft relative to the tracking station. DS1 also used optical data obtained from the MICAS CCD; the images taken of Borrelly during the approach phase were critical in determining the spacecraft’s comet relative position. Finally, DS1 also employed an interferometric data type known as Delta Differential One-Way Range (DDOR). DDOR differences the range signal received simultaneously at two tracking stations to obtain an angular measurement of the spacecraft relative to a line connecting the two tracking stations. The tracking stations used for navigation as well as commanding and telemetry downlink were the three Deep Space Network stations located at Goldstone, California, Madrid, Spain, and Canberra, Australia.

Although DS1’s autonomous navigation system became operational during its primary technology validation mission, the loss of the star tracker precluded its subsequent use for cruise operations since it relied on the MICAS CCD (which was taken over for use as a star tracker). Thus, for the remainder of the cruise to Borrelly, standard radiometric navigation techniques were used to determine its

trajectory. After initial detection using the camera, optical data was added to the orbit determination process.

One important difference between this and other missions was the planning and execution of trajectory correction maneuvers (TCMs). With the IPS, course corrections were burns which could last up to several months long, punctuated at various intervals by periods of ballistic coasting. An additional complication was the fact that the spacecraft’s attitude had to be maintained by locking onto a single bright star using the CCD. Since stars of sufficient brightness were not that common, the attitude used for the thrust profile was often not the optimal one for achieving the desired trajectory. The process of computing a viable thrust profile to keep the spacecraft on course was very complicated, but is out of the scope of this paper and will not be covered in more detail.

The approach phase of the mission began at the first sighting of Borrelly, which occurred roughly 40 days prior to encounter. At this stage, the optical data type became the dominant data type and was relied upon heavily to target the spacecraft to its flyby aimpoint. Due to large uncertainties in the comet’s ephemeris, however, two DDOR data points were taken to help resolve discrepancies between ground and spacecraft based observations of Borrelly. In the end, the spacecraft was guided by referencing its position and target aimpoint to Borrelly itself rather than an inertial location. TCMs during this phase were originally planned to be accomplished using a combination of IPS and hydrazine thruster, but ended up using the IPS alone. The final targeting TCM was performed at Encounter (E) – 12 hours.

Targeting at JPL is performed in the so-called B-plane coordinate system. The B-plane, shown in Figure 1 for the Borrelly flyby, is a plane passing through the center of the target body and perpendicular to the incoming asymptote, S, of the hyperbolic flyby trajectory. Coordinates in the plane are given in the R and T directions, with T being parallel to the Earth Mean Ecliptic plane of 2000; to complete the right-hand coordinate system, T is positive

downwards. The angle theta determines the rotation of the semi-major axis of the error ellipse in the B-plane relative to the T-axis and is measured positive right-handed about S. The horizontal coordinate in the B-plane is referred to as $B \cdot T$ and the vertical is $B \cdot R$.




Figure 1: Borrelly B-plane

The one piece of the autonomous navigation system that remained usable was the closed-loop onboard tracking system. This system enabled the spacecraft to maintain visual lock on Borrelly as it flew by. It was initialized with ground-based ephemeris knowledge about 6 hours prior to encounter. Starting at E-30 minutes, it shuttered images of Borrelly at a rate of about one per minute and used this information to update its estimate of the flyby trajectory. This information was passed to the onboard Attitude Control System (ACS) to point the camera in the correct location to capture Borrelly.

The following sections will describe in more detail the activities and processes used during the approach phase to navigate the Borrelly encounter.

Ground-based Comet Ephemeris Development

Due to the relatively large non-gravitational forces which act on comets (e.g., outgassing),

predicting an accurate ephemeris for even short periods into the future can be quite difficult. Thus, even though ground telescopic observations going back several decades were available for Borrelly, an intensive campaign was undertaken to improve its ephemeris for the DS1 flyby¹. After its recovery in the sky during its current apparition in May 2001, over 200 observations were obtained from telescopes located at Loomberah Australia, the United States Naval Observatory in Flagstaff Arizona, and the Table Mountain and Palomar observatories located in southern California. The observations were processed by members of the Solar System Dynamics (SSD) group at JPL and the ephemeris delivered to the DS1 navigation team. In all, three deliveries were made; the first using just the ground observations and the last two using a combination of spacecraft and ground observations.

Spacecraft Observation Campaign

Starting at roughly E-40 days, an observation campaign was laid out to image Borrelly at various times during the approach. The spacing and timing of these campaigns, referred to as Spacecraft Observations of Borrelly (SOB), had to maintain a balance between obtaining enough images to use for navigation, while not unduly taxing the ground operations teams or placing the spacecraft at risk with unnecessary maneuvers. The final plan for the SOBs is listed in Table 1.

Image Processing

Based on predictions for the brightness of Borrelly's nucleus and coma, and the known sensitivities and noise characteristics of the CCD, it was highly unlikely that Borrelly would be visible in any single frame in the initial observation sets. Thus, the signal-to-noise ratio was increased by co-adding the individual frames together to produce a

Table 1: Spacecraft Observations of Borrelly

SOB #	Date	Range to comet (km)
1	Aug. 25	40,313,000
2	Aug. 29	34,800,000
3	Sep. 7	21,750,000
4	Sep. 10	17,900,000
5	Sep. 13	13,200,000
6	Sep. 15	10,880,000
7	Sep. 16	9,120,000
8	Sep. 18	6,610,000
9	Sep. 20	3,220,000
10	Sep. 21	2,050,000
11	Sep. 22	621,000

composite image. The procedure was started by first determining the inertial pointing direction of the camera boresight. This was done by locating a minimum of two stars in the image and then using a high precision cross-correlation technique to compute their centroids². This technique typically achieved centroiding accuracies of 0.1 to 0.3 pixels. The computed locations of the stars in the FOV, combined with their known right ascension (RA) and declination (DEC) enables a least-squares computation of the boresight pointing direction in inertial space. Then, with the latest best estimate of the spacecraft and comet ephemerides and knowledge of the boresight pointing, the nominal sample/line location of the comet in the camera FOV can be computed. In each frame of the observation set, an $n \times n$ subframe was extracted around the nominal center location and these were added together to form the composite. The subframe size n was chosen such that it encompassed a region larger than the expected errors in the comet's ephemeris errors; the size varied from 20-40 pixels. This co-addition technique was used up

to SOB5, after which the comet was bright enough to centroid in individual frames.

Orbit Determination Strategy

Determining the heliocentric location of the comet was a difficult process requiring careful combination of ground-based and spacecraft observations. However, because the planning of targeting maneuvers was very time critical, waiting for results of this analysis was not a practical way to conduct the encounter. Fortunately, the optical data type offered a means to determine the spacecraft's trajectory independent of the inertial heliocentric orbit of the comet. Since the optical data provided a target relative measurement, it could be used to effectively tie the spacecraft's location to Borrelly; all maneuvers were then computed in this relative coordinate frame. The orbit determination (OD) procedure used was to first obtain the best fit trajectory based on the radio data alone. Then, starting from an initial position and velocity from this estimate, the optical data was used to shift just the spacecraft's position (the velocity was held fixed). Thus, the comet-relative asymptote of the trajectory would not be changed, but its location was translated to where the optical data placed it relative to the comet. Table 2 chronologically lists the various OD solutions, each labeled by the month and day of the last radio data used in the fit, along with the last optical observation used (starting with SOB2 since SOB1 was not accurate enough to use in the fit).

Maneuver Strategy

Maneuver planning and implementation was considerably different on DS1 than on spacecraft with standard chemical propulsion systems. Because the IPS is continuously thrusting over long periods of time, a substantial portion of the trajectory is devoted to performing a maneuver, as compared to chemical maneuvers which occur nearly instantaneously. In addition, IPS thrusting could be separated into two categories – the first is a deterministic “mission burn”, whereby the thrust is needed to impart enough energy to the orbit to achieve a rendezvous, and the second is a statistical trajectory correction

maneuver (TCM), where the course is fine tuned to achieve a specific flyby target. By the time of the approach operations, the former had already been accomplished; only TCMs were needed for hitting the correct flyby aimpoint. For spacecraft safety reasons, the aimpoint distance was chosen to be at 2000 km since it was assumed that at this range, the chance of particle impacts was not severe. For spacecraft geometry reasons, it was to be along the sunline; the combination resulted in the aimpoint chosen to be at a B•R of 2000 km, and a B•T of 0 km.

Two factors were primarily responsible for complicating maneuver planning. The first was the fact that DS1 was continually thrusting at a low level, regardless of the need for mission burns or TCMs, up to a week before encounter. It was found early on in the mission that gimbals on the IPS engine allowed enough thrust vectoring to maintain the spacecraft attitude without the use of the hydrazine-fueled ACS thrusters. Thus, in order to preserve scarce hydrazine for large attitude adjustments, general attitude control was done using the IPS. Although this strategy was critical to mission success, it made the maneuver planning process very difficult. In particular, since maneuvers are planned by first propagating the spacecraft’s trajectory forward to the target conditions, a good prediction of the non-gravitational forces acting on the spacecraft between the current time and time of encounter is necessary. Since it was difficult to predict exactly the future attitude maintenance thrust parameters nor their exact implementation timing with the IPS, the precision of the propagated trajectory was not always very good.

The second complicating factor was caused by the loss of the star tracker. Because bright stars were needed by the camera to lock onto, TCMs could not be performed in completely arbitrary directions. Thus, the IPS thrust vector that would be ideal for reaching the target was not often met. Instead, a suboptimal direction dictated by the nearest bright star was used.

Table 2: Approach OD Solutions

The maneuver strategy that was used during the approach phase turned out to be unusually

complicated, partly due to the above two

OD Solution	Last Used Borrelly Observation
0907	SOB3
0910	SOB4
0912	SOB4
0913	SOB5
0915	SOB6
0916	SOB7
0918	SOB8
0920	SOB9
0921	SOB10
0922	SOB11

factors, but also due to other constraints placed on the spacecraft attitude. In particular, it was desired to place the spacecraft in an orientation such that the high gain antenna was always pointed towards the Earth to maintain a constant communication link. Because of the peculiar geometry of the approach, the line-of-sight direction from Earth to DS1 was almost completely in the B-plane, and at a roughly 45 deg angle. With the spacecraft high gain in this orientation, any IPS thrust would move the spacecraft in the B-plane roughly along this line, pushing the aimpoint negatively in B•R, positive in B•T. Thus, as long as the spacecraft’s trajectory placed the flyby in the bottom left quadrant of the B-plane, it could be corrected by simply throttling up on the IPS without the need to change the attitude. On the other hand, if accumulated OD and IPS execution errors overshoot the aimpoint (above and to the right in the B-plane), the spacecraft would have to be rotated 180 deg. to correct the error, which was highly undesirable from an operations viewpoint. For this reason, the targeting was always performed to bias the aimpoint to the lower left quadrant; as the spacecraft got closer to Borrelly and the OD improved, the aimpoint would be moved closer

to the desired location along the line, but never overshooting it. In all, seven TCMs were planned, but only three were actually executed. Table 3 lists the dates of the maneuvers which were actually executed.

Table 3: Executed IPS Trajectory Correction Maneuvers

TCM ID	Date
2.1	Sep. 17
4.1	Sep. 21
4.2	Sep. 22

Approach Phase

The approach phase of the mission began with the first spacecraft observation set for Borrelly, SOB1. Using the co-addition technique, 8 frames from SOB1 were processed. The result showed a faint signal, barely above the background, which appeared very near the predicted location of Borrelly. The result, though, was not conclusive. Four days later, 12 co-added frames from SOB2 showed a distinct signal about 180 DN's above the background noise. The centroid of this signal (determined relative to the centroids of two co-added stars from the same frames) was roughly 1.8 pixels away from its predicted location, indicating an ephemeris mismatch of about 1500 km, much larger than the predicted value based on ground-based comet observations.

By the time of SOB3, the comet had brightened enough that a composite of 4 frames provided enough signal-to-noise to enable good centroiding. Thus, two sets of composites were produced from the eight usable frames in this set. Due to the closer range to the comet, the observed minus computed location of Borrelly in the FOV had increased to roughly 5 pixels, consistent with the 1500 km error seen in SOB3.

At this point, the cause of the large discrepancy was unknown and could have been due to a

gross error in the estimate of either the comet's or the spacecraft's trajectory. Due to the fact that the ground observations of Borrelly were very consistent and the addition of each day's observations showed only minor changes, the spacecraft was suspected. In order to resolve this, a DDOR campaign was scheduled. It was hoped that the addition of this data type might uncover a subtle error in the Doppler/range based estimates of the spacecraft's trajectory.

In the meantime, the OD and maneuver planning was still implemented using the comet-relative strategy described above. Figure 2 shows the OD results in the B-plane for all the solutions up to September 18 (the ellipses are the formal, 1 sigma uncertainties in the solutions). The shift between solutions 0907 and 0910, which both used observations up to SOB3, was caused by various mismodellings of the attitude control IPS thrusting which occurred in the days between the solutions. This level of B-plane drift is indicative of the general OD accuracy achievable with the difficulty in predicting IPS thrusting events. Changes in IPS thrust on/off times, thrust level knowledge and attitude knowledge inaccuracies are all systematic effects which were difficult to predict and contributed to drifts in the B-plane.

The shift between solution 0910 and 0912, was caused by a larger effect. Originally, TCM 1.1 was planned to move the flyby location to a B•R of 2500 km and B•T of -750 km, which lies roughly along the preferred thrust line direction. Unfortunately, due to the September 11 events, work at JPL was not possible that day and the commands were not sent. Furthermore, the spacecraft lost lock on its star and therefore was unable to maintain attitude, with the result that the spacecraft thrust vector wandered slowly across the sky. This combination shifted the flyby to a location that was coincidentally very near the desired flyby location.

This result was not desired, however, due to the concern that the flyby point might wander above the target, requiring the need for large attitude changes to correct. Fortunately, it was noted that reducing the thrust level at the current orientation would move the flyby point

towards the lower left in the B-plane, where it was originally intended to be. Solution 0913 shows the result of this implementation.

With the addition of SOB6, the SSD group delivered an updated ephemeris which included several apparitions of ground data as well as spacecraft data through SOB6. Solution 0916 was the first to use the new ephemeris, and shows the flyby point to be relatively stable from the 0913 solution. At this time, the OD solutions were accurate enough, and it was getting near enough to the encounter, that a planned TCM, 2.2, was executed to adjust the trajectory to a location nearer to the target. Solution 0918 shows the result after the execution of TCM 2.2.

One disconcerting piece of data was seen in the composite frames of SOB4 and 5. The image of the comet showed several distinct brightness peaks, separated by several hundred km, with the orientation roughly 45 deg away from the sun. The phenomenon was not an effect of the image processing as it appeared in two successive frames with the angular separation of the peaks increasing as would be expected. There was some debate as to whether the secondary peaks was actually the nucleus, and more importantly, whether the comet had fragmented, posing a danger to the spacecraft. Ultimately, it was decided not to change the flyby aimpoint.

On September 14 2001, two DDOR data were taken and folded into the radio solutions. Comparisons of radio based estimates of the spacecraft orbit with and without the DDORs, and trying different combinations of filter assumptions (eg, varying the relative weights of Doppler, range and DDOR, using arcs of differing lengths) showed remarkable consistency. The variation in the B-plane was only on the order of 25-30 km, indicating that the spacecraft's trajectory was probably correct.

With this piece of data, the focus shifted to determining why the ground-based comet ephemeris did not agree with the spacecraft observations. Eventually, it was found that if the center-of-brightness computed from the ground observations used the brightest pixel, rather than the standard Gaussian fit to the

brightness profile, the results agreed considerably better with the spacecraft. Furthermore, observations taken at Palomar Observatory and processed using the bright pixel method, now were in fairly good agreement with the spacecraft. Nevertheless, discrepancies still existed which were eventually attributed to the lack of an accurate model for outgassing used in the comet orbit estimates. Recently, it was found that an acceleration model which had jets at the assumed comet pole, and varying with the angle between the pole and the sun, resulted in the ability to fit longer data arcs from the ground when combined with spacecraft data³.

Figure 3 shows the evolution of the solutions following TCM 2.2 and through the encounter. During this period, the spacecraft had switched to using the RCS thrusters for attitude maintenance. Since these thrusters were not balanced, they imparted a net velocity change to the spacecraft. This is reflected in the roughly 200 km shift between solutions 0918 and 0920. Also, following SOB8, the SSD group delivered the final Borrelly ephemeris to the project. An additional observation set, SOB10, showed the trajectory to be fairly stable in the short span of time between solutions 0920 and 0921.

At this time, less than 24 hours remained until the encounter. TCMs 4.1 and 4.2 were executed to close the remaining gap between the current flyby point and the target, although the roughly 150 km bias in B•R remained. Solution 0922, computed 10 hours before encounter using all the observations, shows a slight drift in the positive B•T direction, but not significant enough to cause concern. In any




Figure 2: B-plane solutions prior to Sept. 18, 2001




Figure 3: B-plane solutions after Sept. 18, 2001

case, no further maneuvers were planned, so remaining targeting errors at this time could not be corrected. Solution 0922 was the final one performed on the ground prior to the flyby; this was uplinked to the spacecraft about 8 hours before the encounter to initialize the RSEN autotracking system.

Encounter Target Tracking

Unlike encounters with planets, the largest error source when targeting a flyby of a small body is the knowledge of the body's ephemeris. Since the gravitational bending of the spacecraft's path is usually negligible, optical images of the target are the only means of precise targeting. However, due to a combination of the high speed of the flyby, light times on the order of tens of minutes, narrow camera fields-of-view (FOV), and the need to load an observation sequence well before the encounter, even the optical data does not provide sufficient accuracy to keep the target in the camera field near closest approach. Therefore, for conventional non-autonomous mission, a sequence is loaded which performs a mosaic; that is, images are taken which cover the navigation uncertainties projected into the camera focal plane. Thus, in order to guarantee an image of the object, multiple frames are returned with empty sky. This is how previous flybys of small bodies, such as Galileo's encounters with Gaspra and Ida, and NEAR's encounter with the asteroid Mathilde, were performed.

With an autonomous closed-loop system onboard DS1, however, the images taken in the tens of minutes prior to encounter can be used to update the spacecraft's target relative position. This system was developed as part of DS1's autonomous navigation system. The target tracking portion was coined RSEN, for Reduced State Encounter Navigation (RSEN). RSEN uses target images to update the spacecraft's position relative to the comet. It performs image processing to locate an approximate center-of-brightness of the target, and, after a number of images have been processed, updates the spacecraft state using a least-squares filter. In order to improve speed, the dynamics are reduced to straight line

motion relative to the target body; since the gravity effects are minimal, this does not lead to loss of accuracy. In addition, since DS1 relied on gyroscopes to maintain inertial attitude during the encounter, the gyroscope drifts and biases also had to be estimated in the filter. A more complete description of the RSEN system can be found elsewhere⁴.

At approximately 8 hours before encounter, the final ground-based navigation solution, 0922, using all available observations, was uploaded to initialize RSEN. Although the formal error ellipse of this solution in the B-plane was only several km, RSEN was initialized with a 20x20 km ellipse to account for systematic errors which may have crept into the solution. At about E-1.5 hours, RSEN snapped its first set of images. These images were processed and the results sent to the ground, but were not actually used. They did, however, provide confirmation that the RSEN system was functional. At about E-30 minutes, RSEN started its encounter imaging sequence, shuttering images about every 30 seconds. As each image was processed, its computed comet center location was stored, but the spacecraft state was not updated at this stage. Finally, at slightly before E-10 minutes, all the accumulated observations were used to estimate a new spacecraft position relative to the comet. The updated ephemeris was provided to the onboard ACS system so that the new information would be used to point the camera. At this point, the solution was updated with every image to keep track of the comet. RSEN was terminated at about E - 2 min., 13 sec. Figure 4 shows the succession of images at several times during the final approach. The comet appears as a bright point of light near the center of the frame (the smeared image across the top of the first two frames is due to stray light). Note the comet drifting slowly out of the FOV as the ephemeris error becomes larger than the FOV; with the state update, the comet returns to near the center of the FOV in the E - 9 minute frame. In the last frames taken, the image of Borrelly is drifting to the top of the frame, even though post-encounter analysis indicated that RSEN was correctly predicting the position to a small fraction of a




Figure 4: Sequence of RSEN images during final approach

FOV. It is believed that this drift was due to a latency of between one and two seconds in the response of the ACS. Although this phenomena had been hypothesized early in the planning of the encounter, there were no means to confirm its actual presence. In the end, none of the planned images were lost due to this effect.

Conclusion

The successful flyby of Borrelly provided the science community with the highest resolution images of a comet nucleus to date, adding considerably to the body of knowledge of these mysterious solar system bodies. Figure 5 shows the final image snapped by the spacecraft about two minutes prior to closest approach, taken at a range of 3514 km, with a surface resolution on the comet of 46 m/pixel. The navigation challenges presented by this encounter were considerable, and was met by the introduction

of several first-of-a-kind technologies. These included the use of an ion propulsion system for course changes, and an autonomous nucleus tracking system. It is hoped that DS1 has paved the way for future missions to use these technologies with confidence, ensuring even greater science returns.

Acknowledgements

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

1. Chesley, S. R., Chodas, P. W., Keesey, M. S., Bhaskaran, S., Owen, W. M., Yeomans, D. K., Garrad, G. J., Monet, A. K. B., and Stone, R. C., "Comet 19P/Borrelly Orbit Determination for the DS1 Flyby", *Bulletin*

of the American Astronomical Society, v. 33, p. 1090 (2001).

2. Bhaskaran, S., Riedel, J. E., Synnott, S. P., “Autonomous Optical Navigation for Interplanetary Missions”, *Space Sciencecraft Control and Tracking in the New Millennium*, E. Kane Casani, Mark A. Vander Does, Editors, Proc. SPIE 2810, pp32-43, 1996.
3. Chesley, S. R., personal communication.
4. Bhaskaran, S., Riedel, J. E., Synnott, S. P., “Autonomous Nucleus Tracking for Comet/Asteroid Encounters: The STARDUST Example”, AAS Paper 97-628, AAS/AIAA Astrodynamics Conference, Sun Valley, ID, August 1997.





Figure 5: Best resolution image of Borrelly



Beacon Monitor Operations Experiment DS1 Technology Validation Report

Dennis DeCoste, Susan G. Finley, Henry B. Hotz, Gabor E. Lanyi,
Alan P. Schlusmeyer, Robert L. Sherwood, Miles K. Sue, John Szijarto,
E. Jay Wyatt
*Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California 91109*



Table of Contents

<u>Section</u>	<u>Page</u>
Table of Contents	ii
Figures	ii
Tables	ii
Extended Abstract	iv
Beacon Monitor Operations Experiment (BMOX)	1
DS1 Technology Validation Report	1
1.0 Introduction	1
2.0 Technology Description	1
2.1 What It Is/What It Is Supposed To Do	1
2.2 Key Technology Validation Objectives at Launch.....	3
2.3 Expected Performance Envelope.....	3
2.4 Detailed Description.....	5
2.5 Technology Interdependencies.....	8
2.6 Test Program	8
3.0 Technology Validation Summary	10
3.1 Tone Experiment Results	10
3.2 Data Summarization Results	10
3.3 Operational Effectiveness Assessment.....	12
3.4 Lessons Learned.....	15
4.0 Technology Application for Future Missions	17
5.0 Acknowledgments	18
6.0 List of References	18
Appendix A. List of Telemetry Channels and Names	19
Appendix B. DS1 Technology Validation Power On/Off Times	21

Figures

<u>Figure</u>	<u>Page</u>
Figure 1. Operational Concept	1
Figure 2. Tone-Signal Structure	2
Figure 3. Monitoring Signal Detector and Message Decoder	5
Figure 4. Prioritized Summary Data Description	7
Figure 5. ELMER Adaptive Alarm Limits.....	7
Figure 6. 1-sec Fourier Spectra of the Input Signals to the Four-Tone Detectors	8
Figure 7. Fourier Spectra of the Output of the Tone-Detectors after Aligning and Summing (and averaging) 10 FFTs of 1-sec Each.....	9
Figure 8. Tracking of Adaptive Alarm Limit to DS1 Solar Array Temperature	13
Figure 9. Battery Temperature Episode Detection	13

Tables

<u>Page</u>	<u>Page</u>
Table 1. Tone Definitions.....	2
Table 2. BMOX Validation Summary	4
Table 3. Summarization Telemetry Packets.....	7
Table 4. BMOX Validation Schedule and Matrix.....	10
Table 5. List of Tone Experiments.....	11
Table 6. Summary of Engineering Data Monitored	12
Table 7. Tracking Cost Per Month (34m BWG, 2 contacts per week).....	15

EXTENDED ABSTRACT

The Beacon Monitor Operations Experiment (BMOX) was one of twelve new technologies that were flight validated on NASA's Deep Space 1 Mission (DS1). The technology enables a spacecraft to routinely indicate the urgency of ground contact using a tone signal rather than telemetry while also summarizing onboard data to be transmitted whenever telemetry contact is required. This technology can be used to lower operational cost, decrease mission risk, and decrease loading on the over-constrained Deep Space Network antennas. The technology is baselined on upcoming NASA missions to Europa, Pluto, and the Sun. Successful flight validation has met a requirement to demonstrate the technology before routine use on the Europa mission.

The end-to-end, Beacon-tone signaling system was developed to provide a low-cost and low-bandwidth method for determining when ground intervention is required. With Beacon monitoring, the spacecraft sets the tone signal and it is transmitted either in a scheduled manner or continuously, depending on spacecraft operability constraints. The tone signal is detected on the ground with smaller aperture antennas than would be required for telemetry on a given mission. Tone detection times are short—on the order of 15 minutes or less for most mission designs. The flight validation experiment checked out the functionality of the tone-detection and message-delivery system, characterized operational performance, obtained parameter limits, and tested selection of tone states by flight software based on the spacecraft's assessment of its own health. The tone system was tested on the DS1 spacecraft in both the X-band and Ka-Band.


Engineering data-summarization flight software creates event-driven and periodic summaries of spacecraft activities since the last contact. Episodes are created by identifying the culprit and causally-related sensors around the time of important events. This data is gathered at a high sample-rate, assigned a priority, and stored for downlink at the next telemetry pass. The gaps are filled in by "snapshots" of all

sensor channels at a much lower sample-rate. The software can use either traditional (static) alarm thresholds or adaptive alarm-limit functions that are determined by a statistical learning network. The adaptive alarm-limit technology, called the Envelope Learning and Monitoring using Error Relaxation (ELMER) is one of two artificial intelligence (AI) components in the current software design. The second AI-based method computes empirical transforms on individual data channels. These pseudo-sensors enhance the value of summaries and serve as an additional input in determining the adaptive limits. The software was originally developed to support Beacon monitor operations, an approach that enables the spacecraft to determine when ground contact is necessary. In this approach, summarization plays a key role in providing operators with the most important data because all of the stored data cannot be downlinked in a single telemetry pass. Efficient summaries also help facilitate quick troubleshooting and thus can reduce the risk of losing the mission. Summarization algorithms can also be applied to nonspace systems to decrease the time required to perform data analysis. The current version of the software runs on VxWorks and has been executed on the PowerPC and RAD6000 target processors.

The experiment also included operational testing of a ground system prototype, called BeaVis (Beacon Visualization), that was designed to facilitate quick interaction with BMOX data. The purpose of this system is to track Beacon-tone states throughout a mission and to display downlinked summary data. For Beacon missions, the user must be able to quickly maneuver through summary data to arrive at an assessment of overall system state and to diagnose any problems that occur. The software enables the user to scroll through a graphical depiction of telemetry downlinks throughout the life of the mission to select the desired data. Summary data is represented graphically with a hypertext style link to the strip charts of the sensor channels contained in each of the four types of summary data packets. A web version of the tool was also implemented.



Beacon Monitor Operations Experiment



What is It?

The Beacon monitor operations technology provides the spacecraft the functionality required to initiate telemetry tracking only when ground intervention is necessary.

Why Is It Exciting Technology?

- Mission operations cost is reduced substantially because there is less contact with the spacecraft
- Reduced loading on ground antennas enables more spacecraft to be operated with existing ground resources
- Beacon uses state-of-the-art techniques for summarizing onboard spacecraft performance data

How Does it Work?

- Instead of routinely sending spacecraft health data, the spacecraft evaluates its own state and transmits one of four Beacon tones that reveal how urgent it is to send high-rate health data
- When telemetry tracking is required, the spacecraft creates and transmits "intelligent" summaries of onboard conditions instead of sending bulk telemetry data to the ground

When Will it be Demonstrated?

- Flight demonstration occurred on the Deep Space 1 mission launched in October 1998
- The technology is being adopted by the DS1 Extended Mission to lower operations cost
- The technology has also been baselined for planned NASA missions to Europa, Pluto, and the Sun

Technology for Low Cost Operations

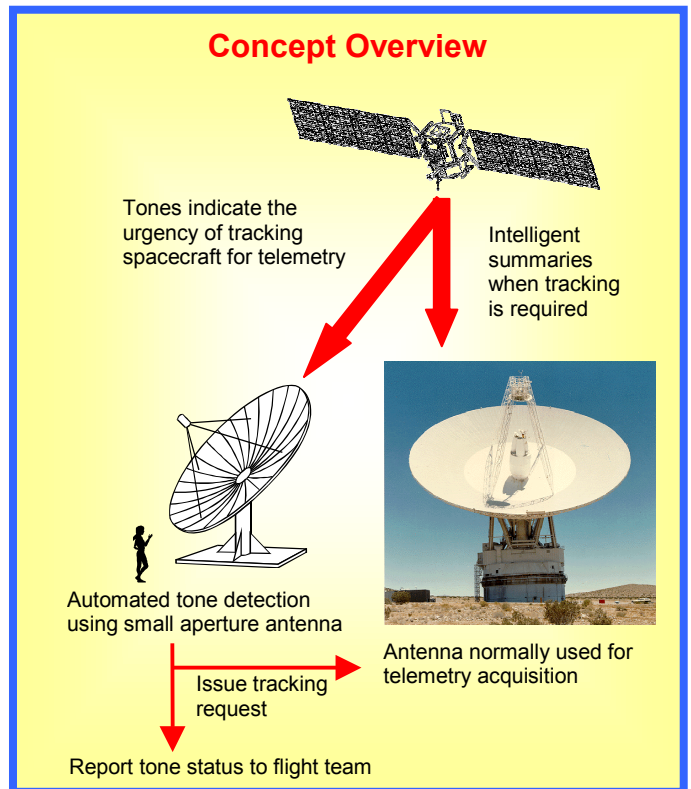
Principal Investigator: E. Jay Wyatt

Email: e.j.wyatt@jpl.nasa.gov

Team: Dennis DeCoste, Sue Finley, Henry Hotz, Gabor Lanyi, Alan Schlutsmeyer, Rob Sherwood, John Szijjarto, Miles Sue

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California 91109

World Wide Web - <http://eazy.jpl.nasa.gov/Beacon>



National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology



Beacon Monitor Operations Experiment (BMOX) DS1 Technology Validation Report

*Dennis DeCoste, Sue Finley, Henry Hotz, Gabor Lanyi, Alan Schlutsmeyer,
Robert L. Sherwood, Miles Sue, John Szijarto, E. J. Wyatt
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California*

1.0 INTRODUCTION

The budget environment that has evolved since the advent of NASA's Faster, Better, Cheaper initiative has caused mission-risk policies and mission designs to change in ways that have been conducive to the inception of new operations concepts and supporting technologies. Such was the case when the Beacon monitor concept was conceived to enable a mission to Pluto to be achieved within the budget constraints passed down from NASA. The technology was accepted into the New Millennium Program and baselined for flight validation on the DS1 mission. As the technology was being developed for DS1, the NASA community has expressed a growing interest and acceptance of adaptive operations and onboard autonomy.

In traditional mission operations, the spacecraft typically receives commands from the ground and, in turn, transmits telemetry in the form of science or engineering data. With Beacon monitoring, the spacecraft assumes responsibility for determining when telemetry will be sent and sends what amounts to a command to the ground to inform the flight operations team how urgent it is to track the spacecraft for telemetry. There are only four such commands. Thinking of Beacon operations in this way creates a paradigm shift over the way operations are traditionally approached. Also, it is very important to not think of the tone message as just a little bit of telemetry. If one does this, it is easy to make the argument that a little more telemetry is better. Our approach is one where telemetry is only transmitted when it is necessary for ground personnel to assist the spacecraft. If the spacecraft goes through long periods (a month or so) without requiring ground assistance. When telemetry tracking is necessary, the intelligent data summaries contain the most relevant information to provide full insights into spacecraft activities since the last contact. The key challenge has been to develop an architecture that enables the spacecraft to adaptively create summary information to make best use of the available bandwidth as the mission progresses such that all pertinent data is received in one four-to-eight-hour telemetry pass.

This work was funded from three NASA funding sources. The NASA Cross Enterprise Technology Development Program (CETDP) Thinking Systems Thrust Area funded flight software development. The Telecommunications and Mission Operations Directorate (TMOD) Mission Services

Technology Program funded development of the tone detection algorithm and also funded development of flight software. Additionally, a small amount of funding from the New Millennium Program was supplied towards the end of the prime mission to help offset the additional costs imposed by DS1 schedule delays.

2.0 TECHNOLOGY DESCRIPTION

2.1 What It Is/What It Is Supposed To Do

Beacon Monitor Operations refers to a spacecraft-initiated operations concept and the supporting technology components. The supporting technology components are the tone subsystem and the onboard engineering data summarization subsystem, both of which were flight validated on DS1. The operational concept shown in Figure 1 depicts a typical end-use scenario where the spacecraft routinely sends one of four X-band tone messages that indicate how urgent it is to track for telemetry. This tone is received at a smaller aperture antenna than would be required for telemetry for that mission. If the tone indicated that telemetry tracking was required, a summary of the important telemetry data stored onboard since the last contact would be downlinked via a normal telemetry link.

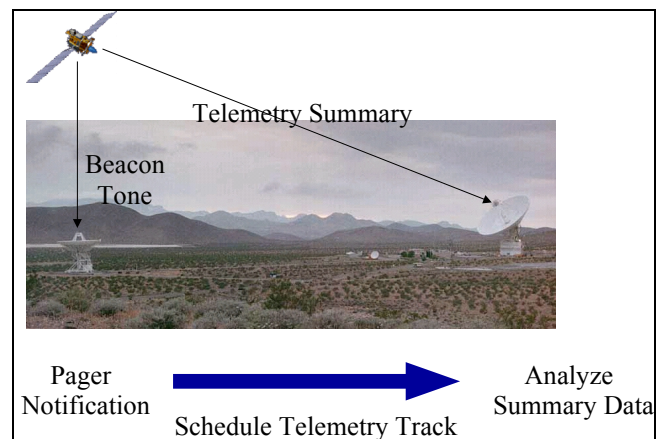


Figure 1. Operational Concept

Advantages of using this technology fall into three categories: reducing mission cost, reducing Deep Space Network (DSN) loading, and reducing mission risk. Operations cost is reduced by reducing the frequency of contact and by reducing the total volume of downlinked data. Savings are realized through staffing reductions


(because fewer people are required to analyze telemetry) and reductions in antenna usage. These reductions help the DSN contend with the oversubscription problem that exists today and that is poised to become worse in the future due to the large number of planned missions. Mission-risk reductions are another major advantage to this technology. At first glance, it may seem that Beacon operations is more risky than traditional operations. However, with today’s faster-better-cheaper missions, scheduled telemetry tracking is being scaled-back due to cost constraints. With Beacon monitoring, the spacecraft can, at low cost, transmit assurances that the spacecraft is behaving as expected in between scheduled telemetry tracks. This reduces the chance of having a catastrophic, time-critical failure and, for ion-propulsion system, affords the additional advantage of verifying that thrusting is ON. If, for example, an ion mission lost thrusting immediately after a scheduled telemetry pass, a week or more may pass before ground personnel become aware of the problem. With Beacon, response time could be cut to just a few days (or less). Loss of thrusting for a week or more could cause the mission to not reach the target body.

2.2.1 Beacon Tone Monitoring System—As mentioned before, the tone system is used to routinely monitor the health of the mission. There are four tone signals; each signal uniquely represents one of the four urgency-based Beacon messages. The DS1 tone definitions are summarized in Table 1. These tones are generated as the spacecraft software reacts to real-time events.

Table 1. Tone Definitions

Tone	Definition
Nominal	Spacecraft is nominal. All functions are performing as expected. No need to downlink engineering telemetry.
Interesting	An interesting and non-urgent event has occurred on the spacecraft. Establish communication with the ground when convenient. <u>Examples</u> : device reset to clear error caused by Single Event Upset (SEU), other transient events.
Important	Communication with the ground needs to be achieved within a certain time or the spacecraft state could deteriorate and/or critical data could be lost. <u>Examples</u> : memory near full, non-critical hardware failure.
Urgent	Spacecraft emergency. A critical component of the spacecraft has failed. The spacecraft cannot autonomously recover and ground intervention is required immediately. <u>Examples</u> : PDU failure, SRU failure, IPS gimbal stuck.
No Tone	Beacon mode is not operating. Spacecraft telecom is not Earth-pointed or spacecraft anomaly prohibited tone from being sent.

It is important to communicate the urgency of ground response using a telecommunications method that has a low detection threshold and short detection times. Ease of detection translates to lower cost operations. The signal structure is shown in Figure 2. Each message is represented by a pair of tones centered about the carrier frequency. Tones are generated by phase-modulating the RF carrier by a square-wave subcarrier using a 90-degree modulation angle. The carrier frequency (F_c) is completely suppressed. The resulting downlink spectrum consists of tones at odd multiples of the subcarrier frequency above and below the carrier. Four pairs of tones are needed to represent the four possible messages.



B=Frequency uncertainty F_c =Carrier frequency
 f_i =Subcarrier frequency for the i message

Figure 2. Tone-Signal Structure

2.1.2 Onboard Summarization System—If the Beacon tone indicates that tracking is required, the onboard summarization system provides concise summaries of all pertinent spacecraft data since the previous contact. This subsystem gathers high-level spacecraft information—such as the number of alarm crossings, spacecraft mode and state histories, and other pertinent statistics—since the last ground contact. It also gathers episode data for the culprit and causally related sensor channels whenever a sensor violates an alarm threshold and stores the data at a high sample rate. It collects snapshot telemetry at a much lower sample rate for all sensors and transform channels. Snapshot data serves only for rough correlation and to fill in the gaps between episodes. The last component of the downlinked summary—performance data—is similar to episode data but captures maneuvers or other events known in advance to be of interest to people on the ground. All of the summary algorithms are implemented in C for the VxWorks operating system.

The summary algorithms incorporate AI-based methods to enhance anomaly-detection and episode-identification capability. The Envelope Learning and Monitoring using Error Relaxation (ELMER) technology replaces traditional redlines with time-varying alarm thresholds to provide faster detection with fewer false alarms. The system uses a statistical network to learn these functions; training can be performed onboard or on the ground (ground-based for DS1). ELMER is particularly powerful because it requires very little domain knowledge and trains the statistical network with nominal sensor data. Another artificial

intelligence (AI) method produces empirical transforms that have a heritage in previous AI research at JPL in selective monitoring. Once computed onboard, these act as virtual sensors. The current transforms for DS1 compute high, low, and average values, and first and second derivatives. Alarm limits can be placed on these transforms and also serve as an input to the ELMER adaptive-alarm limits. Additional transforms, if desired, can easily be defined and uplinked to the spacecraft as the mission progresses.

2.2 Key Technology Validation Objectives at Launch

The primary validation objective was to verify that the two subsystems (tone and summarization) were fully deployed and operating as expected. This was accomplished through a series of experiments to test the basic functionality of the deployed system. An additional validation objective was to evaluate the operational effectiveness of using the technology on future missions and on DS1 in the extended mission phase.

Validation objectives were captured in a signed Technology Validation Agreement between the BMOX Team and the DS1 project.

2.2.1 Objectives Prior to Experiment Turn-on—

1. Test summarization algorithms and ground visualization environment using representative spacecraft data (Topography Experiment (TOPEX/Poseidon)) prior to DS1 testbed data availability
2. Provide unit-test verification test runs in “Papabed” and Testbed environments for test of all BMOX flight software capability
3. Verify that the tone detector can automatically detect weak signals using schedule and predicts information

2.2.2 Expected In-flight Observables—

1. Tones detected at DSS 13 during experiment activities, conducted periodically throughout the prime mission
2. Tone message delivery to JPL
3. Engineering data summaries downlinked during scheduled DS1 project telemetry passes
4. Characterization of tone system behavior with mission distance
5. Demonstration of the ability to detect spacecraft anomalies, map to Beacon tones, and detect the tones on the ground in a timely manner
6. Produce summary data that provides value-added information if Beacon monitoring were to be used as the primary mode of operations

7. Characterization of DS1 staffing level for routine operations and a comparison of that staffing level to the expected level of support required in performing Beacon operations
8. Detailed analysis of antenna tracking time with and without Beacon operations
9. Assessment of the number of mission anomalies or events requiring ground intervention

Success Criteria (Quantifiable/Measurable Goals):

2.2.3 Prior to Experiment Turn-on—

1. Tones detectable at DSS 13 throughout the primary mission phase
2. Adaptive summaries of spacecraft health information that result in downlink bandwidth savings over traditional downlink approaches
3. Telecom system capable of generating X-band tones per Small Deep Space Transponder specifications

2.2.4 In-Flight—

1. Determination of the size of engineering data summaries required to adequately analyze spacecraft conditions when the tone indicates that ground intervention is required
2. Tone detection probability of 95% or greater
3. Onboard tone selection accuracy of 95% or better for urgent conditions
4. Message delivery latency less than 1 hour
5. Major (urgent) event capture in summary data 90% or better using traditional alarm limits, 70% or better using adaptive alarm limits (after initial checkout period)
6. Summary data sufficient for determining corrective actions at least 75% of the time
7. Ability to display summary data within 2 hours of downlink data available to DS1 project
8. Determine, through operational experiments, that Beacon operations will reduce routine operations cost on DS1 by at least 25%
9. Determine, through operational experiments, the exact level of expected savings in operations-staffing cost and antenna-tracking cost on future JPL missions.

2.3 Expected Performance Envelope

Table 2 illustrates the full set of validation objectives and the weighting of each in computing the percent validated at any point during the mission and includes brief descriptions of the experiments that were conducted and the associated success criteria.

Table 2. BMOX Validation Summary

Experiments	Goal	Success Criteria	Validation %	Antenna	When	How many tone passes	Pass Duration, hr
1. Engineering Summary Data Generation & Visualization, and Tone Selection			50%				
1.1 Data Generation and Visualization – Functional checkout	Demo end-to-end functionality of on-board data summarization system.	Summarization algorithms work as expected during DS1 mission operations.	25%	HGA	Starting late Feb., 99	(Regular DS1 Telem.)	Depends on bandwidth
1.2 Data Generation and Visualization – Detailed performance verification	Performed detailed analysis of all features of the software.	Summarization data successfully determines spacecraft anomalies with enough detail for spacecraft engineers to respond appropriately.	15%	HGA	Jul. – Dec., 99	(Regular DS1 Telem.)	Depends on bandwidth
1.3 Tone Selection	Demo FSW functionality to set and reset the tones and meaningful mapping from spacecraft health to urgency-based request.	Tones are set as a result of a spacecraft data out-of-limits condition. Parameter file can be easily updated and uploaded. Tones selector is reset.	10%	HGA or LGA	Apr. – Dec., 99	Some telemetry, some mid-week	1
1.4 Final analysis & report generation	Analyze and document results, lessons learned, and as-flown design in a final report.	The software system provides a viable means for conducting spacecraft-initiated operations on future space missions.	Not included in validation				
2. Tone Trans. & Detection			40%				
2.1 SDST functionality checkout	Verify that the SDST can correctly generate Beacon tones.	SDST generates and transmits the 4 Beacon tones, as instructed via uploaded commands.	20%	HGA	Jan., 99	1	2.5
2.2 Tone Calibration - X	Calibrate Beacon frequency & tone detector parameters, and verify predicts. Establish the lowest threshold and the longest integration time possible.	Successfully detect Beacon tones and obtain frequency uncertainty estimates.	10%	HGA or LGA	Feb. - Mar., 99	4	1
2.3 Tone Detection - LGA	Demonstrate weak-signal detection.	Detect signal with power level 5-10 dB Hz.	5%	LGA or HGA	Mar., 99	1	1
2.4 Tone Detection - Ka	Obtain Ka-band Beacon signal characteristics.	Successfully detect and record Ka-Beacon signal.	5%	HGA	Mar. – Apr., 99	1	1
2.5 Detailed analysis & report generation	Analyze and document tone-transmission and detection system results in a final report.	Beacon signaling system provides a viable means for conducting spacecraft-initiated operations on future space missions.	Not included in validation				
3. Multi-mission Ground Support			10%				
3.1 Functional demo of tone notification process	Demonstrate a low-cost and reliable process to detect and deliver Beacon messages in a realistic environment.	The tone detector detects and delivers Beacon messages within 1/2 hr after the Beacon tone pass.	10%		Feb. - Mar., 99	Use passes from 2.2 above	
3.2 DSN Track Automation	Demonstrate viable demand-based DSN antenna scheduling schemes and methods for automating the tone detection process.	Beacon- triggered DSN passes can be successfully scheduled using a real DSN station schedule.	Optional for extended mission				
4. Ops Concept Assessments			N/A				
4.1 Effectiveness Assessment	Produce a final report documenting results of cost benefit analysis.	Quantify future mission-tracking cost and personnel cost savings for Beacon operations.	Not included in validation				
4.2 Perform Beacon operations during DS1 prime mission operations	Evaluate effectiveness through Beacon ops for DS1 ops benefit.	Beacon ops is mature enough to support DS1 extended mission.	Optional post-validation activity				
4.3 Perform Beacon operations during DS1 extended mission	Provide updates to flight software and continue performance assessment.	Demonstrated ops-cost savings during DS1 extended mission.	Optional for extended mission				

2.4 Detailed Description

2.4.1 Tone Experiment Detailed Description—The tone monitoring technology consists of generation, transmission, and detection of the tone signals. The primary requirement was to transmit tones in X-band; however, Ka-band was tested to help pave the way for future missions that may use a Ka-band transponder. The experiments were also constructed so that detection of weak signals, such as from a mission to Pluto, could be validated. Finally, tone-message handling and reporting and overall low-cost operation of the tone system was assessed.

There are four tone signals. Each tone uniquely represents one of the four urgency-based Beacon messages. For a description of the tone meanings, refer to Table 1.

BMOX was designed so that the urgent Beacon tones are sent when the spacecraft fault protection puts the spacecraft in standby mode. This condition occurs when the fault protection encounters a fault that it cannot correct. Standby mode halts the current command sequence, including IPS thrusting. The software to control this condition was onboard the spacecraft but never enabled.

During the DS1 tone experiment, the Beacon tone was sent at prescheduled times for about 30 minutes. The Beacon tone was not operated continuously because DS1 requires as much power as possible for IPS thrusting and the tone transmission reduces the power available for thrusting.

The tone is sent using the DS1 Small Deep Space Transponder (SDST). The signal structure is shown in Figure 2. A pair of tones centered about the carrier represents each message. These tones are generated by phase-modulating the RF carrier by a square-wave subcarrier using a 90-degree modulation angle. The frequency carrier (F_c) is completely suppressed. The resulting downlink spectrum consists of tones at odd multiples of the subcarrier frequency above and below the carrier. For the DS1 experiment, the four-subcarrier frequencies (f_1 , f_2 , f_3 , and f_4) are 20, 25, 30, and 35 kHz, respectively. Different frequency allocations can be assigned to different missions. The monitoring system is designed to achieve a low-detection threshold. The goal is to reliably detect the monitoring messages with 0 dB-Hz total-received-signal-to-noise-spectral-density ratio (Pt/No) using 1000 seconds observation time.

The Beacon message is first received and decoded by the Goldstone site and subsequently transmitted to a signal detector at JPL. Next, the Beacon message is forwarded to DS1 Mission Operations and other end users, including the Demand Access Scheduler, using e-mail or pagers.

The signal detector contains four tone detectors, one for each message. To ensure proper signal detection, the band-

width of each tone detector must be sufficiently large to accommodate the frequency uncertainty and frequency drift of the downlink frequency: i.e., the Beacon tones for a given message will not drift outside of the passband of the detector for that message. The FFT (Fast Fourier Transform) is employed to compute the energy of all spectral pairs having spacing corresponding to the four Beacon signals. Because of oscillator instability, Fourier transforms cannot be produced over long time intervals. The total observation time is divided into short intervals. FFTs are first performed over these short intervals and then incoherently combined after the frequency drift has been removed. The maximum of the outputs of the four tone detectors is then selected and compared against a pre-determined threshold to determine which message has been received. A block diagram for the signal detector and the message decoder is shown in Figure 3.

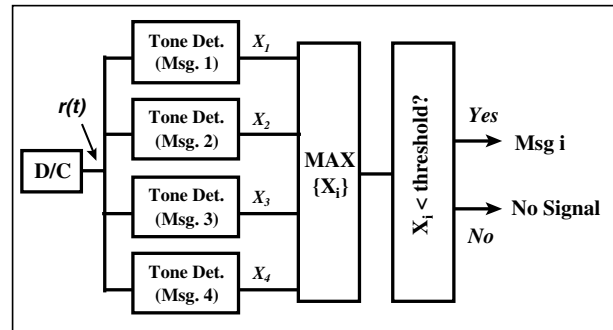


Figure 3. Monitoring Signal Detector and Message Decoder

2.4.1.1 Tone Transmission and Detection Experiment—The four Beacon messages are represented by four pairs of tones; these tones will be generated by modulating the downlink carrier with an appropriate subcarrier using a 90-degree modulation angle. The four subcarriers selected to represent the four Beacon messages are:

<u>Beacon Message</u>	<u>Subcarrier Frequency, KHz</u>
NORMAL	20
INTERESTING	25
IMPORTANT	30
URGENT	35

The DS1 spacecraft is equipped with two transmitters: X-band and Ka-band. When Beacon tones are being transmitted via one of the two links, no telemetry can be sent over the same link. However, DS1 can transmit Beacon signals using one link (e.g., X-band) and simultaneously downlink telemetry using the other link (e.g., Ka-band).

The first tone pass was used to verify the functionality of the Small Deep Space Transponder (SDST). Four commands were sent directly to the SDST software manager, each representing a different tone. The Beacon flight software

was not used during this test. The tones were detected on the ground beginning in January, 1999.

The next part of the experiment used four Beacon passes to calibrate the signal and compare against prediction. A set of Beacon tone states was loaded into the command sequence on the spacecraft. Beacon tones were then generated onboard, transmitted to the ground, and detected by DSS 13 at Goldstone. The detector and tone frequencies were calibrated, predicts were verified, and detector parameters were determined. In the first three tests, the Beacon tone states were pre-selected, but unknown to the tone detection personnel. In the last test, a tone was generated by the onboard Beacon flight software. These tone passes occurred between February and April, 1999. This set of four tone passes was the minimum required to calibrate the detection system and validate its performance.

All Beacon passes require dedicated use of either the LGA or HGA during a Goldstone pass. Telemetry and Beacon signals cannot be transmitted simultaneously over the same communication link (of the same frequency, X- or Ka-band); therefore, Beacon passes were scheduled to accommodate the DSN telemetry passes. In addition to the above calibration-tone experiments, two additional experiments were scheduled to test the performance of the Beacon-tone detector using the Ka-band frequency and using the X-band frequency in a weak-signal regime. The Ka-band experiment was identical to the X-band experiment except for the frequency. The purpose of the weak-signal X-band experiment was to determine the threshold at which the signal can no longer be detected. These two experiments were scheduled to occur during March–April, 1999.

2.4.1.2 Multi-mission Ground Support Experiment—The objective of the Multi-mission Ground Support Experiment was to demonstrate a low-cost, reliable process to deliver Beacon messages to the flight project within a reasonable amount of time. For the DS1 Beacon experiment, this time was defined to be less than 30 minutes. The Beacon tone passes from the tone transmission experiments were used in this experiment. During these passes, Beacon messages were generated, transmitted, and subsequently detected by the ground station (DSS 13). The detected messages were delivered to the BMOX team at JPL via e-mail or pager. Post-Beacon pass telemetry was used to verify the correct transmission times.

2.4.2 Data Summarization Detailed Description—If the Beacon tone indicates that tracking is required, the onboard summarization system provides concise summaries of all pertinent spacecraft data since the previous contact. The summarization system performs three functions: data collection and processing, mission activity determination, and episode identification. The data collection sub-routine receives data from the engineering telemetry system via a

function call and applies summary techniques to this data, producing summary measures for downlink to the ground. The mission activity sub-routine determines the overall spacecraft mode of operation. This determination is used to choose the appropriate data and limits monitored by the episode sub-routine. The mission activity is intended to be exclusive. When a new mission activity starts, the previous mission activity is assumed to have ended. The episode sub-routine combines summary and engineering data received internally from the data-collection sub-routine with the mission activity received from the activity sub-routine and compares the data with mission-activity-specific alarm limits. It is necessary to use the mission activities to determine which data to use for episode identification and to identify the limits of these data. If the limit is exceeded, the sub-routine spawns a new episode and collects past relevant data from the data collection sub-routine. The past data collected will be one-minute summaries that go back in time as far as the user has defined. (Therefore, a five-minute episode would contain summaries starting five minutes before the episode to five minutes after the episode.) At the end of the episode, the sub-routine outputs data to the telemetry subsystem for downlink.

Three different types of summarized data are produced onboard: overall performance summary, user-defined performance summary, and anomaly summary. Six different telemetry packets have been defined to contain this information (see Table 3). Taken as a whole, the telemetry packets produce summary downlinks that are used to enable fast determination of spacecraft state by ground personnel. The summary data is prioritized in the downlink so that the most important data is sent first (Figure 4). The first telemetry sent is a summary of events since the previous downlink. Next, the episodic data, the nominal data, and, finally, the user performance are sent.

The performance summaries are generated at regular intervals and stored in memory until the next telemetry-round contact. They are computed by applying standard functions, such as minimum, maximum, mean, first derivative, and second derivative, to the data. User-defined summary data can provide detailed information on a particular subsystem and are created at the user's discretion. Anomaly summary data (episodes) are created when the raw and summarized data violate high or low limits. These limits are determined by the subsystem specialist and stored in a table onboard the spacecraft. The limit tables are based on the current mission activity.

The software also has the capability to use AI-based envelope functions instead of traditional alarm limits. This system, called Envelope Learning and Monitoring using Error Relaxation (ELMER), provides a new form of event detection will be evaluated in addition to using the project-specified traditional alarm limits. Envelope functions are

essentially adaptive alarm limits learned by training a statistical network with nominal engineering data (see Figure 5). The network can be onboard or on the ground. For DS1, envelope functions are trained on the ground and

then uploaded to the spacecraft. DS1 spacecraft fault protection will only be based on project-specified static-alarm limits; however, the summary data can be generated based on the adaptive limits.

Table 3. Summarization Telemetry Packets

Telemetry Name	Description	Output Frequency
Activity	Current value of mission activity	Output on change
Data Sample	Records a snapshot of every raw and summarized data channel	Regular interval: i.e., 15 min.
Episode Summary	Records general data about an out-of-limits data condition called an “episode”	One per episode
Episode Channel	Records specific data about a single data channel’s behavior during an episode	One or more per episode
Tone Change	Current state of the Beacon tone	Output on tone change
Channel Summary	Summary data about a single data channel’s behavior since the last downlink	One for each channel out of limits
User Summary	A user-specified packet containing raw and/or summarized data	Duration user-specified

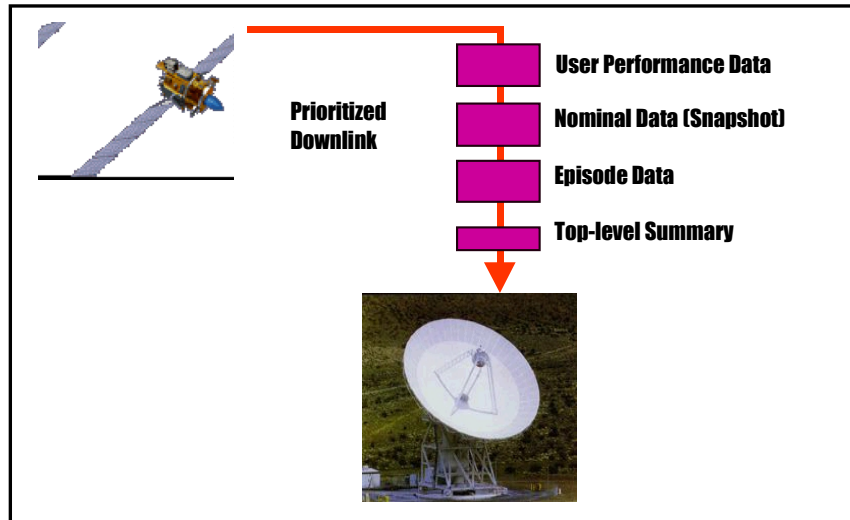


Figure 4. Prioritized Summary Data Description




Figure 5. ELMER Adaptive Alarm Limits

The sampler module and its related data-gathering module currently consist of 3038 lines of source code and 222 KB of memory on the Power PC series processors. Activity determination is a rare event and processing time is negligible. The once-per-wake-up processing time for DS1 averages 30 ms.

2.5 Technology Interdependencies

DS1 BMOX was designed to have minimal impact on the operation of the baseline DS1 mission. There are, however, some important interdependencies to note for future missions that may be interested in deploying the technology. These are summarized as follows:

- The transponder should be capable of transmitting beacon tone signals. The Small Deep Space Transponder (SDST) has this capability, as does the Space Transponding Modem (STM).
- The algorithms used for anomaly detection within the Summarization System should be the same as those used for fault detection within the fault-protection subsystem. Otherwise, summary data may not capture the relevant data.
- Bandwidth-constrained missions will likely have more of a use for tone monitoring.
- Operationally-constrained spacecraft designs make unattended operations difficult, adding cost and decreasing the utility of Beacon operations.

2.6 Test Program

2.6.1 Ground Test—A number of system-level tests/demonstrations were conducted throughout the development process to validate the design concept and hardware/software interfaces. These tests/demonstrations were also conducted to satisfy project-related requirements.

2.6.1.1 SDST/Tone Detector Compatibility Tests—The first major test was to validate the compatibility between the tone detector and the SDST. Beacon signals were generated by the SDST (engineering model) in the radio laboratory in Building 161. The signals were transmitted to a test facility in Woodbury, where the signals were down-converted to 300 MHz IF and recorded by the Full Spectrum Recorder (FSR). The recorded signals were processed by the tone-detection algorithm installed in the FSR.

An example of the detection results is shown in Figure 6 and Figure 7 using 20 KHz as a signal frequency. Figure 6 gives the Fourier spectrum of a 1-sec snapshot of the monitoring signal before being processed by the detector: i.e., the spectra of the input signals to the four tone detectors. Figure 7 gives the Fourier spectra of the outputs of the four tone detectors after aligning, summing and averaging over 10 FFTs, each of 1-sec duration. The horizontal line is the detection threshold corresponding to a given false-alarm probability. As shown in the figure, the aligning and summing process significantly reduces the noise fluctuation and enhances signal detection.




Figure 6. 1-sec Fourier Spectra of the Input Signals to the Four-Tone Detectors

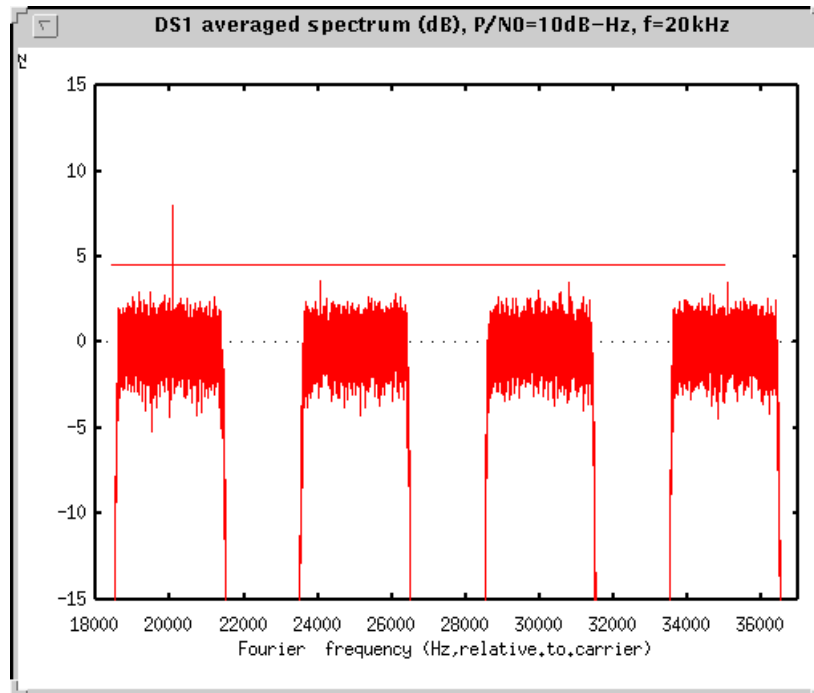


Figure 7. Fourier Spectra of the Output of the Tone-Detectors after Aligning and Summing (and averaging) 10 FFTs of 1-sec Each

The recorded data was subsequently and successfully used in a concept demo, which is one of the requirements imposed on technologies by DS1. During the conceptual demo, segments of the previously recorded SDST Beacon signal data were selected for the tone detector to perform real-time detection. The detector, located in Building 111, was remotely operated from the SMOCC room in Building 301, where the concept demo was given. Detection results were sent to the SMOCC room via a network connection and displayed on a projection screen in real-time. Segments of the recorded data were selected for the demo and the tone detector successfully detected the signals and displayed the detection results.

A second compatibility test was performed with the flight transponder, during which the spacecraft was in the thermal vacuum chamber and the tone detector was transported to the Telecom Development Laboratory (TDL). The SDST was commanded to send Beacon tones one at a time to the TDL using a fiber-optic link. The signal was demodulated and down-converted to IF at the TDL. The received signal was displayed on a spectrum analyzer. The observed spectra confirmed that the SDST had correctly generated and transmitted all monitoring signals as commanded. In addition, the received monitoring signals were fed to the tone detector, where they were digitized, recorded, and subsequently detected. These tests revealed that there are no interface or compatibility issues between the SDST and the tone detector and ensured that they would work smoothly as a tone system.

2.6.1.2 Tone Detection System Test—In addition to being able to detect very weak signals, it is envisioned that an operational tone system would be capable of schedule-driven, predicts-driven, fully-automated tone detection and message delivery. This would lower the operations cost, which is critical if this technology were to be employed as an operational capability. The original DS1 experiment plan was to leverage on the DST technology to demonstrate in-flight such a capability. A series of system tests was designed and conducted in the TDL to demonstrate (1) predicts generation capability, (2) DST/Tone detector interface and file transfer, and (3) automated detection using frequency predicts. Frequency predicts were generated by the DST controller using a SPK file obtained from the DS1 Project database. The predict file along with a trigger file were then sent to the tone detector and were subsequently used to detect the TDL-simulated Beacon signals. Two automated Beacon detection demonstrations were conducted by using simulated spacecraft tones at TDL. DS-T-generated frequency predicts and a trigger file were used to initiate the detection of a scheduled pass. The detector detected Beacon signals at the 7 dB-Hz power signal-to-noise level using 10-s integration time with a probability of false detection of 0.01. BMOX team members, Section 331 engineers, and DS1 management attended this demo. It fulfilled the pre-launch readiness requirement. This test also paved the way for a subsequent in-flight demo.

2.6.2 Flight Test—The test program consisted of executing the experiments described in Section 2.3. Testing began in

January, 1999 and continued through the end of the prime mission in September, 1999. Table 4 depicts the flight-validation schedule.

3.0 TECHNOLOGY VALIDATION SUMMARY

The technology was declared fully validated in July, 1999, after both the summarization and tone systems were fully deployed and tested as described in Section 2. The overall system performed as expected and was considered a success.

3.1 Tone Experiment Results

A series of experiments were run to test the end-to-end tone delivery system. These experiments were designed to incrementally test additional capability for the Beacon-tone system. Prior to launch, the ability of the SDST to generate Beacon tones was tested by the telecom engineers. A similar test was performed on the spacecraft several times after launch. This test was called “X-tone” because it tested the capability to send the Beacon tones using X-band transmission. The X-tone test, expanded to use a series of tones to test the ground detection system, was repeated several times throughout March and April, 1999. The dates of these and other tests are listed in Table 5.

The ability of the software to select tones and transmit them in DS1 telemetry was tested on February 26, 1999. This test, called b-tone, consisted of ground commands that set the Beacon tone during a downlink pass. The tone was verified in regular DS1 telemetry but was not transmitted to the tone detector. Each tone was verified during the b-tone test. In addition, the tone-reset command was tested.

The next test to run onboard DS1 was the b-transmit test. This test involved setting the Beacon tone using information from the software on board, then transmitting the tone using the SDST. The tone was received at the DSS 13 antenna and

forwarded to the tone detector at JPL. No advance knowledge of the commanded tone was given to the ground detection engineer. After the tone was detected, it was delivered to other members of the Beacon team in an e-mail message. The b-transmit test was run three times in April, 1999.

The last tone test to be run was the Ka-tone test. This test was identical to the X-tone test except that it used the Ka-band transmitter to send the Beacon tone. This test was run in April, 1999.

3.2 Data Summarization Results

The data summarization was first turned on February 19, 1999. The Beacon team determined the limits applied to the engineering data for testing the summarization capability. The limits were set just outside of the minimum and maximum value seen for the data since launch. Shortly after the first turn-on, several of the data channels went into episode (out-of-limits) condition. Upon further inspection, it was determined that many limits were based on engineering units (EU), but much of the data was being stored using data numbers (DN) in EH&A. The data summarization was turned off after several hours, and the initialization file (also called sampler init file, or SIF) was updated with DN-based limits.

On March 8, 1999, the data summarization was turned back for several hours. A few channels went into alarm; however, the number was reduced from the previous test. Inspection of the data revealed negative values for some eight-bit sensors. This was impossible because all eight-bit sensors should range from 0 to 255. After careful debugging in the DS1 test bed, an error was found in the DS1 flight software. It was discovered that when data are passed from the originator to EH&A, EH&A converts the data to its own internal double-precision format as though it were 8 bits and signed. This results in the values from 0 to 127 being

Table 4. BMOX Validation Schedule and Matrix

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep
SDST Checkout		△							
Tone Calibration		△	—————	△					
Tone Notification		△	—————	△					
Data Summarization - functional Checkout			△	—————	△				
Weak Signal Detection				△					
Ka-Band Detection					△				
Software Update & Testing				△	—————	△			
Data Summarization performance verification						△	—————	△	
Extended Mission Planning								△	—————

Table 5. List of Tone Experiments

Date	Experiment Type	Results
Jan 6	X-tone, 20, 30, 25, & 35 kHz	Tones found in this order after accounting for 20-second offset in spacecraft internal time. Detection time = 5 min. Frequency offset (FRO) = -4.25kHz, (high gain antenna)
Feb 4	X-tone, 35 & 20 kHz	Noisy and stable sub-carriers used with low modulation indexes from low gain antenna. All successfully detected. FRO = -1.98kHz
Feb 26	B-tone & X-tone	Software tone test. All four tones were commanded and transmitted through regular telemetry.
Mar 3	X-tone, 35 & 20 kHz	Antenna computers down and wind speeds halted antenna several times and early, but several detections were successful at very low levels. FRO=1.25 kHz: 20.0001 kHz, DN=3, Pd/N0=8.8, 10 sec, 35.0013 kHz, DN=2, Pd/N0=4.2, 15 sec.
Mar 18	X-tone, 30, 20, 25, & 35 kHz	X-tone successful. After 4.4 kHz carrier offset was found and applied. Spacecraft time found to be 10 seconds later than predicted. IPS was on.
Mar 24	X-tone	X-tone semi-successful. X-tones found but wrong frequencies because carrier predicts were off by 4.5 kHz and not entered in FSR.
Apr 7	X-tone, 20, 25, 30, & 30 kHz	X-tone successful. Station needs 45 minutes pre-cal vs. 30. FRO=5.0 kHz.
Apr 13	B-transmit & X-tone, 20, 25, 30, & 35 kHz	B-transmit successful, 25 kHz tone, needed visibility of carrier before carrier suppression to get correct FRO of 5.5 kHz. X-tone was also successful.
Apr 19	Ka-tone	The FSR at DSS 13 tracked the Ka carrier but the Ka-tone sequence did not get transmitted to the S/C as the auto-nav processing took longer than expected. FRO=0.0 (3-Way).
Apr 20	B-transmit	B-transmit successful, detection code found 25 kHz tone, needed visibility of carrier to find correct FRO of 6.0 kHz.
Apr 26	Ka-tone, 20, 25, 30, & 35kHz	Ka-tone was successful for the sequence that was activated. Detection of 20 kHz tone at DN=1 was 4.5 Pd/N0 for 15 sec. FRO=9.9 kHz (wrong up-link freq. in predicts).
Apr 27	B-transmit	Detection code found 25 kHz tone, FRO of 6.9 kHz was used to center the signal.

represented correctly, and the values from 128 to 255 being represented as -128 to -1, respectively. EH&A apparently does not have a data-type code for unsigned 8-bit integers. The effect of this problem was that limits were harder (and sometimes impossible) to specify. With a new set of rules, it was possible to create a SIF that would work around this problem for some of the data. If both high and low limits were 128 or greater, they had to be converted by subtracting 256. However, if the low limit was 127 or less and the high limit is 128 or greater, the limits won't work. Sensor values with both limits less than 127 could remain unchanged. With these rules, another SIF was created and uploaded to DS1. Data summarization was restarted on March 22, 1999. Everything appeared to operate correctly in data summarization. A few data channels went into episode condition. It was determined that temperature sensors were drifting colder due to DS1 moving away from the sun. The limits were updated and a new SIF was uplinked.

Data summarization ran smoothly on and off during the month of April and May, with minor modifications to the SIF due to noisy channels. During this period, a new version of the Beacon FSW was developed and tested. This version

included a work-around for the limitation of EH&A data described above. In addition, the following new features were added:

- The criteria for determining mission activity was parameterized in the SIF
- Episodes will now end if a new SIF is loaded
- Additional protection for divide-by-zero conditions
- SIFs can now be loaded from EEPROM or RAM
- User-data packets can now have start and stop times associated with them

The new version was started up on June 15, 1999. A new SIF was included with limits determined by the DS1 spacecraft engineers. Since that time, data summarization has needed a few updates due to false alarms. There are several reasons for these false alarms. The Beacon FSW is able to sample the data once per second. This is a much higher rate than the data sent to the ground for analysis. Because of the higher rate, the FSW is able to see events that are normally missed on the ground. These events have been confirmed by correlating with fault-protection monitors that capture maximum excursions on the same sensors.

Another reason for false alarms has been activities such as optical navigation (OPNAVs) that move power and thermal sensors outside their normal ranges. The subsystem engineers respond, “Yes, these events take the sensors outside their normal ranges, and yes, this is expected behavior.” So where does the Beacon team set the limits? Since the Beacon data summarization is context sensitive, a new “mission activity” for OPNAVs could be created with its own set of limits. An OPNAV activity consists of several spacecraft turns, with picture taking occurring at each target. This is similar to a maneuver. With this in mind, the mission-activity determination criteria for maneuvers has been changed to include optical-navigation activities. This will also make the maneuver activity determination more robust. Prior to this change, switching to maneuver activity when DS1 was actually firing thrusters was only used to change the velocity. Maneuvers involve turning to a thrusting attitude and turning back after the thrusting. Now, the maneuver activity includes these turns and their respective settling times as well. This makes sense because it is during this entire period that power and thermal sensors may deviate from their nominal cruise values. This change was uplinked in early September, 1999. The current list of engineering data being monitored is listed in Appendix A. A summary of this list is contained in Table 6.

Table 6. Summary of Engineering Data Monitored

Subsystem	Number of Channels
Attitude Control	8
Fault Protection	1
Navigation	1
Other	2
Power	22
Propulsion	1
Telecommunications	6
Temperature (all subsystems)	35

Beacon data summarization has been an evolving process requiring several limit refinements from the spacecraft team. This should be expected in the development of any data summarization system. This process is very similar when any new mission launches. For the first several months, ground alarms are updated as the flight team learns about how the spacecraft really operates. The ground-testing activities give a good first cut at setting alarm levels; however, the spacecraft never operates exactly as it did in test. Implementing context-sensitive limits is a similar process. Engineering data limits are no longer set based on the worst case. Now the worst case can be viewed based on the spacecraft activities. This should ensure more accurate discovery of anomalies.

One activity that produced important results involves analyzing summary-system performance on DS1 anomalies to date. Although capabilities were limited due to onboard

memory restrictions, preliminary results when running ELMER on historical data are showing that adaptive alarm thresholds can track gradual trending of sensor data much tighter than the current DS1 static alarm limits. This is seen in monitoring the gradual drift in eight solar-array-temperature sensors, one of which is shown in Figure 8. Comparing traditional limits with ELMER limits during the 81 days of operations, ELMER limits track actual spacecraft performance much more precisely than static limits, which would be off the scale of this chart.

Another validation exercise has confirmed that summarization can capture subtle, yet important spacecraft episodes. In ground tests, ELMER detected an unexpected heater turn-on that occurred when the solar panels went off-axis during a spacecraft maneuver. Since ELMER trains across multiple parameters using nominal data, the summarization system detected this event without explicit a priori knowledge of the scenario. This data is shown in Figure 9.

ELMER has been running onboard with only 10 sensors, all temperature. This limitation is primarily due to limited onboard memory. There have only been three ELMER limit violations (episodes) during the primary mission. Two have occurred during OPNAV events and can be explained by the temperature excursions associated with spacecraft turns. These are basically “false alarms.” The third episode has not yet been explained. The ELMER limit functions were developed after training on data from the first four months of the mission. It is hoped that additional training on spacecraft data since February will correct these false alarms in an extended mission. There will be additional ELMER limit functions added in an extended mission as well.

3.3 Operational Effectiveness Assessment

The experiment afforded insights into the operational cost savings that a future mission might realize. Computing cost savings for DS1, however, was not possible in the prime mission because Beacon technology was not used operationally by the mission. Although not specified in any plans, the best measure of the effectiveness of the technology turned out to be the interest expressed by the DS1 team in using it for the extended mission phase. In August, 1999, work began with the DS1 team to help infuse the technology into the planned two-year extended mission to two additional target bodies. The technology was seen as a way to contend with the severe cost constraints that extended missions face. Luckily, one of the BMOX design objectives was to deploy the technology experiment in a manner that would allow the mission to use it once validated.

There were many important results on how to design, implement, and operate Beacon-monitor operations systems

on future missions. The entire end-to-end experience of working with a flight project team to field this experiment resulted in uncovering important design considerations and lessons learned that will be useful to future missions that plan to use the technology. These are described in the remainder of this section.

3.3.1 Data Processing Issues—Beacon summary data was delivered to the Beacon team through an automated batch

script that queried the data each night. The data was placed in a public directory and then processed by the Beacon team the next morning. The processing was a simple task, but was not automated because data summarization was frequently turned off for days to weeks at a time. During DS1's extended mission, data summarization should be on continuously and, therefore, the data processing should be automated.




Figure 8. Tracking of Adaptive Alarm Limit to DS1 Solar Array Temperature




Figure 9. Battery Temperature Episode Detection

The database used to store Beacon summary data was created specifically for the Beacon task. Because summary data is not easily formatted for commercial databases, it was decided to develop a DS1 database. In hindsight, this was the wrong decision. It has been very difficult to maintain a custom database. The users do not have good visibility into the database if the tools are not working correctly. Changes to the database take a programmer to change the code instead of running a tool that would be provided with a commercial database. In addition, commercial databases have built-in query features that are easy to set-up and use. There were instances in which data was requested, but it could not be provided in a timely fashion. Also, custom requests such as one for all episodes involving a specific channel could not be provided. The limitations of using a custom database hindered the operational effectiveness of Beacon.

3.3.2 Data Summarization Software Enhancements—The data summarization software was not relied upon for determining spacecraft state. Although the algorithms and returned summary data seemed adequate, there were several suggestions made by the Beacon personnel and flight team for further enhancements. Some of these suggestions will be incorporated into the M7 version of the flight software to be uploaded during DS1 extended mission operations.

The episode data was lacking depth because it only provided ten samples, each separated by two minutes. The long time between samples was set to ensure that Beacon summary data would not overflow the telemetry buffer in the event of repeated episodes on a single channel. For the M7 version of the software, the number of samples is being changed to 20 and the user will be allowed to set the number of times a channel can go into episode before it stops producing episode packets. With these changes, the sample interval can be set much shorter. In fact, a six-second-sample interval will be used. This will give the episodes more visibility while not overloading the telemetry buffer with false alarms. Making a change and adding all data on change-to episodes was considered; however, the DS1 project only wanted very minor software changes in M7.

During the course of operations, the initialization file with the episode limits was changed and uplinked many times. Many times the changes only involved one or two limits in the file. Because the file is on the order of 15 kilobytes, there were periods of low communications bandwidth when it would take several minutes to uplink the file using the low-gain antenna. Operationally, it is much easier to have a capability to update limits without sending out the entire initialization file.

The flight team made a few suggestions for improving the usefulness of the summary data. The derivative summary functions, but one of the subsystems suggested that integrals

be added to the summary functions. Several other flight-team members suggested adding different persistence for each episode limit check. Currently, there is a global persistence parameter that applies to all episodes. This change will be implemented in our M7 software release. Another suggestion was to add a sample rate to user-performance packets.

Two capabilities that fault-protection monitors have that should be present in Beacon are conditional monitors and maximum excursion tracking. Conditional monitors enable the user to check multiple sensors based on the values of the sensors. The DS1 fault protection software also has the capability to track and save the minimum and maximum values for sensors. The summarization software will only track these values if the sensor goes into an episode condition. This may be important data for future missions relying on summary data even though the sensors are not outside their limits. As mentioned in the Lessons-Learned section, there should be tighter integration between the Beacon software and the fault-protection software.

3.3.3 Reporting Results to the Flight Team—A set of tools for examining the summary data was developed. These tools were only located on the Beacon team workstation. Since launch, some web-based tools were developed to access the summary data. These tools have made it easier to report the results to the flight team, but are very limited in their capabilities. These tools will be improved during extended mission. The goal is to make the data easily accessible to the flight-team users. Easy access to the Beacon data is very important for making the technology operationally effective; unfortunately, access was not available during the DS1 primary mission.

3.3.4 Automation of Tone Detection—Tone-detection automation is proceeding as an activity in support of DS1 Extended Mission and was not an objective of the as-launched system. Tone-detection automation was an objective prior to the TMOD redirection wherein BMOX antenna support was changed from DSS 26 (which supported automated demand-access antenna operations) to DSS 13. Full automation involves automatic-predicts generation, automatically running scripts to perform tone acquisition, detection, and automatic tone-message reporting. Tone-message reporting can, in fact, be quite elaborate, where the autonomous-reporting system expects confirmation from users that tones were received. If not, a fully automatic reporting system would have a roster of the team members and would keep contacting people until the tone message was acknowledged. The lessons learned from conducting tone-detection operations during the mission is that tone acquisition is highly amenable to automation and would substantially lower the cost of performing Beacon operations. Automatic-predicts generation would also serve

other users of DSS 13 and would support broader DSS 13 automation objectives.

3.3.5 Cost Savings from Using Beacon—Part of future work in Beacon technology involves infusing the Beacon technology into DS1 mission operations as an end-to-end system. Technology infusion is not an easy task and traditionally has not been done well. DS1 will benefit from this work by reducing the amount of tracking time used.

In extended mission, DS1 will have two tracking passes per week, an 8-hour, high-gain pass on Mondays, and a 4-hour mid-week pass to check spacecraft status. Utilizing Beacon, the DS1 project will not have to use a 4-hour mid-week DSN pass to check spacecraft status. It can use a 30-minute (or less) Beacon pass that actually provides them with additional information over a carrier-only pass. In addition, the frequency of eight-hour telemetry passes can be reduced and 30-minute Beacon passes substituted. The number of 8-hour telemetry passes that can be eliminated has not been determined, but DS1 expects it could be as many as every other pass. In this case, there would only be two eight-hour telemetry passes each month and four 30-minute Beacon passes each month. The overall savings for this case are summarized in Table 7. This results in savings of 30 hours of DSN tracking time or \$18,248 per four-week period. This does not include the substantial savings of mission-engineering-labor costs of performing routine telemetry analysis.

The benefits of infusing a regular Beacon operation technology on DS1 are apparent in the cost savings of reduced-DSN utilization. In addition, the four-hour mid-week passes are replaced with 30-minute Beacon passes that contain additional status information. Future missions will benefit from the experience of a flight mission using a regular Beacon tone for an extended period of time. This includes the experience of scheduling the DSN for Beacon operations as well as the success of the Beacon tone system in relaying the spacecraft status to the ground. New missions that could benefit from this technology include ST-4, Pluto Express, Europa Orbiter, and MDS. Each of these missions is planning on using either part or all of the Beacon operations technology. The continuation of work on the Beacon technology by revising the operations concept will add value to these mission customers. In addition, the

operations procedures for using the Beacon technology can be fully developed.

Demand-access scheduling of DSN antennas is another important feature of an operational Beacon system. Scheduling antennas based on demand rather than a pre-negotiated agreement is important to the success of this technology within the DSN. During the DS1 extended mission, there is no funding to demonstrate automated scheduling of antenna resources. If a Beacon tone is received that requires contacting the DS1 spacecraft, it will be necessary to manually request a station pass. Until the DSN changes their scheduling paradigm, it will be difficult to implement demand-access scheduling.

3.4 Lessons Learned

3.4.1 Ion Propulsion Missions—The utilization of the ion propulsion system (IPS) (also called solar-electric propulsion) on DS1 offers an additional advantage in using Beacon monitoring. The IPS provides continuous thrust for much of the cruise phase. The operational margin for IPS thrusting represents the duration for which IPS could be off and still allow the spacecraft to reach the target asteroid. Due to the low thrust associated with IPS and because actual thrusting did not start until several weeks after launch, the operational margin is only a few weeks. Telemetry-downlink passes are becoming less frequent as the DS1 mission progresses. Eventually, there will only be one telemetry pass per week. If the spacecraft experiences a problem that requires the standby mode, the IPS engine will be shut down. It could be up to one week before the flight team has visibility to that standby mode. Using the Beacon-tone system during the periods between scheduled-telemetry downlinks can be a cost-effective way to decrease mission risk because it reduces the likelihood of losing thrusting time and not making the intended target. Other future IPS missions have taken note of this fact and requested Beacon-tone services to lower their mission risk.

3.4.2 Software Testing—It was decided to redesign the DS1 flight software about 18 months before launch. This decision greatly compacted an already full schedule to complete the software. As a result, the testing of all non-essential software functions was delayed until after launch. The Beacon experiment was considered a non-essential piece of software and, therefore, was only tested pre-launch

Table 7. Tracking Cost Per Month (34m BWG, 2 contacts per week)

	Monthly cost: DS1 Operations without Beacon	Monthly Cost: DS1 Operations with Beacon	Monthly Savings
8-hour telemetry passes	\$19,465	\$9,733	\$18,248
4-hour carrier only passes	\$9,733	not applicable	
Beacon tone passes	not applicable	\$1,217	
Total	\$29,198	\$10,950	

* assuming reduction of two 8-hour telemetry passes per month

for non-interference with the other flight software. In post-launch testing, a few problems were discovered that prevented the Beacon software from starting until a new version could be uploaded. These problems related to differences between the flight-hardware based testbed and a simulated-hardware testbed. This is the age-old lesson learned by performing system testing on the software prior to use. But even beyond that, it is important to run tests on the actual hardware-based testbed. Unfortunately, the DS1 schedule would not allow this until post launch.

3.4.2 Fault Protection Integration—Before the software redesign, the Beacon software was tightly integrated with the DS1 fault-protection software. The decision was made after the redesign to de-couple the two pieces of software. Previously, the fault-protection monitors triggered the Beacon tones. After the redesign, the mapping of faults to tones was performed using two different methods. All spacecraft standby modes are now mapped to the urgent Beacon tone. The interesting and important Beacon tones are mapped using Beacon software-determined limits. Decoupling the fault protection software from the Beacon software gives this organization maximum flexibility to determine what sensors to monitor. Unfortunately, our algorithms for determining faults are not nearly as sophisticated as the fault-protection monitors. These monitors can look at many different values based on conditional logic before determining what fault has occurred. Future spacecraft designed to use Beacon operations should plan on completely integrating the Beacon tone software with the fault-protection software.

3.4.4 Beacon Signal Frequency Stability—The signals used for Beacon monitor are characterized by three things: (1) the signal strength can be extremely low, (2) the initial tone frequencies, which are derived from an onboard auxiliary oscillator, are not known exactly, and (3) the tone frequencies are constantly drifting. The tone detector is designed to detect these types of signals with a high level of confidence. The maximum-frequency uncertainty and the maximum-frequency drift rate for the tone detector were established using a Galileo spare transponder. An operational issue was encountered with the DS1 Beacon experiment: How and to what extent can the auxiliary oscillator's temperature be stabilized before the start of a Beacon pass? Stabilizing the temperature will reduce the frequency uncertainty and frequency drift, making it easier for the tone detector to detect the Beacon signal. Based on data provided by the DS1 telecom personnel, the auxiliary oscillator temperature can undergo a wide range of changes after an OPNAV maneuver. This results in a very large frequency uncertainty and a very high rate of change (>6 Hz/sec), both of which would exceed the limits of the tone detector (when the signal level is low).

One solution to the OPNAV-related problem is to wait for the transponder temperature to stabilize. Studies by the DS1 telecom personnel indicated that about four hours are needed for the transponder temperature to stabilize after running the OPNAV activity. This operational constraint would not have much impact on the spacecraft and is believed to be the simplest, lowest-cost solution to this problem. This procedure is recommended to improve weak-signal detection for DS1 and future missions using Beacon Monitor.

During the DS1 tone experiments, the initial frequency uncertainty was much larger than expected. A bias was manually introduced to keep the received signal in the recorded band. Without the bias, the frequency might be outside the recorded band. In an automated detection mode, it is necessary to record at least 3 times the current bandwidth, unless a better way to predict the frequency can be found. One possibility is to make use of the auxiliary-oscillator frequency vs. temperature-calibration table to improve frequency prediction.

3.4.5 Downlink Carrier Phase Noise—Post analyses of the received-signal frequency indicated that the phase noise of the downlink carrier was fairly significant. This would result in detection loss. Analyses should be performed to estimate the impact of this phase noise on detector performance and to factor this into future detection experiments.

3.4.6 Spacecraft Clock Accuracy—During one of the experiments, it was observed that the actual tone switching times did not seem to agree exactly with the predicted switching times. This led to the discovery by the DS1 team that there was an error of 18 to 19 seconds in the SCLK/SCET conversion.

3.4.7 DSN Equipment Issues—A couple of tone passes were not successful due to DSS 13 weather and equipment. In one experiment, the spacecraft started transmitting tones before it rose above the horizon of DSS 13. In another case, a scheduled pass was cancelled due to spacecraft activities. While the overall tone experiments have been very successful, future experiment plans should allow for this kind of contingency.

3.4.8 Beacon Operations Paradigm—The Beacon software makes determinations of spacecraft anomalies. The data summarization component of Beacon attempts to summarize related data from these anomalies. These determinations are based upon high and low limits on sensor data. It is important to involve the spacecraft subsystem engineers in the determination of which data to monitor and the setting of the limits on these data. They are the personnel most familiar with the operational characteristics of each subsystem and, therefore, should be determining interesting and fault conditions for their subsystem. Also, by involving them in the data summarization definition, they will become better

acquainted with the Beacon software and will be more inclined to use it during crisis situations.

Ground-alarm limits on telemetry are generally set using the worse possible state of each data channel. This practice can hide problems with the spacecraft if the alarm limits are set at wide boundaries. Beacon data summarization offers context-sensitive limits. In the case of DS1, limits can be set for cruise, downlink, IPS thrusting, maneuver, and standby modes. Spacecraft operations personnel are not accustomed to working with summarized-engineering telemetry or context-sensitive limits. When data limits were requested, generally one set of limits was received with instructions to apply them to all mission activities. Setting limits like this does not utilize the capabilities of the Beacon data summarization. For future implementations of Beacon, it will be important to educate the flight team about Beacon's capabilities early in mission design. Beacon data summarization should also be used during spacecraft testing to familiarize operators with the technology. This will help ensure reliance on Beacon data during the mission.

3.4.9 Systems Engineering—As previously mentioned, there were problems with false-episode alarms due to mission activities such as Optical Navigations, camera calibrations, etc. It is important to carefully define each of the mission activities and how they are related to engineering data. In the DS1 case, the maneuver activity was defined to only occur when the thrusters were firing. Since maneuvers also involved turning the spacecraft, it was important to include all events that turned the spacecraft in our maneuver-mission-activity criteria. Once mission activities are carefully defined, then episode limits for those activities can be developed.

4.0 TECHNOLOGY APPLICATION FOR FUTURE MISSIONS

There are essentially three paths to future work in this area. One is continuing to follow the technology-development roadmap for AI-based onboard-summarization methods. In the coming year, this involves also investigating the notion of summarizing spacecraft data in order to create a comprehensive onboard archive in addition to downlinking summary telemetry. Missions to Europa and Pluto only plan to downlink about 5% of the total volume of engineering data. The summarization algorithms developed for DS1 form a good foundation for investigating how to intelligently capture the most important data in order to maintain an adaptive long-duration onboard archive. This archive may serve as an input to other onboard-autonomy software or it may just be available for downlink if ground personnel require

additional insights into past-spacecraft activity. In addition to pursuing this archiving concept, there are many, many new automated data-analysis methods to investigate for use in onboard summarization systems. This will also be researched in the coming year.

The second thrust has to do with future mission deployments. After the DS1 Extended Mission, the next mission customer is the Europa Mission. Europa is the first mission funded by the JPL Outer Planets/Solar Probe program and currently has a planned launch in 2003. New versions of flight software for summarization and tone selection will be developed in the coming year and will be compatible with the JPL Mission Data System architecture. This architecture is currently baselined for the Europa mission. MDS-compliant software prototypes that build on lessons learned from the DS1 experiment will be delivered to the Europa mission in November, 2000. More generally, the technology is useful to a broad range of deep-space missions. In this era of faster, better, cheaper, there are many advantages to using this type of operations approach instead of more traditional operations. Earth-orbiter missions have different requirements, but can benefit from having Beacon-based adaptive operations. The Beacon-monitor team has long standing ties to Stanford University, Santa Clara University, and the University of Colorado, all of which are developing Beacon-based operations concepts and systems for Earth-orbiting missions.

There is another proposed Beacon concept for an Earth-trailing spacecraft (SIRTF) that involves using one tone. SIRTF plans to track every 12 hours, but would like to have Beacon tracking every 2 hours. The idea is that the spacecraft would only send a Beacon tone if it had a problem. The possible Beacon detections are 1) help tone or 2) no detection. Normally, the spacecraft would be busy doing observations; however, if it had a problem it would turn to Earth point and start transmitting a carrier signal. This Beacon signal could shorten the anomaly response time from 12 hours to a maximum of 2 hours. This requires no modification to the already-designed spacecraft since there is no need to distinguish fine levels of urgency. SIRTF management considers this important because their design does not include a transponder that supports Beacon tones. There is one drawback with this operation. When the tone detector fails to detect a Beacon signal, one can not tell whether (1) the spacecraft is fine and no Beacon has been transmitted or (2) the spacecraft has an anomaly and fails to transmit.

The third thrust involves development of the ground-system infrastructure for conducting Beacon operations. The NASA Space Operations Management Office (SOMO) and the JPL Telecommunications and Mission Operations Directorate have high-level objectives to support Beacon monitoring on future missions. The exact scope and implementation of this multi-mission support has not yet been worked. In the meantime, tone detection for the DS1 Extended Mission is being

supported through special arrangements with the experimental DSS 13 ground station. A more generic tone detection system needs to be implemented if the DSN antennas will support Beacon-monitor missions. In addition, the full benefit of adaptive operations requires demand-based scheduling of DSN antennas. This is also a high-level objective for the DSN.

5.0 ACKNOWLEDGMENTS

Acknowledgments fall into three categories. This experiment would not have been possible without the dedicated support of the BMOX Team Members:

E. Jay Wyatt – Principal Investigator
Robert L. Sherwood – Flight Operations Lead
Miles K. Sue
Henry B. Hotz
Susan G. Finley
Alan P. Schlutsmeyer
Tim Fogarty
John Szijarto

In addition to the core team, several members of the DS1 team had a substantial impact on the quality of the experiment:

Dave Lehman
Marc Rayman
Phil Varghese
Pam Chadbourne
David Morabito

The following program management and project staff also provided critical support throughout development and operations:

John Carraway – Outer Planets/Solar Probe Mission Operations Lead
Richard Doyle – Division 36 Manager
Guy Man – New Millennium Program Autonomy IPDT Manager
Peter Shames – TMOD Mission Services Technology Program Manager
Larry Teitlebaum – DSS 13 Manager

This work was funded by the Autonomy Program and TMOD Technology Program at the Jet Propulsion Laboratory. The research described in this paper was carried out at the Jet Propulsion Laboratory, California

Institute of Technology, under a contract with the National Aeronautics and Space Administration.

6.0 LIST OF REFERENCES

- Doyle, R. et al., “Autonomy and Software Technology on NASA’s Deep Space One,” *IEEE Intelligent Systems Magazine*, May/June 1999.
- Wyatt, E. J., et Al., 1999, “Flight Validation of On-Demand Operations: The Deep Space One Beacon Monitor Operations Experiment,” Fifth International Symposium on Artificial Intelligence, Robotics, and Automation in Space, Noordwijk, The Netherlands.
- Sherwood, R., et al., 1999, “Lessons Learned During Implementation and Early Operations of the DS1 Beacon Monitor Experiment,” *Third International Symposium on Reducing the Cost of Ground Systems and Spacecraft Operations*, Tainan, Taiwan.
- Wyatt, E. J., et Al., 1998, “Beacon Monitor Operations on the Deep Space One Mission,” Fifth International Symposium on Space Mission Operations and Ground Data Systems, Tokyo, Japan.
- DeCoste, D., 1997, “Automated Learning and Monitoring of Limit Functions,” International Symposium on Artificial Intelligence, Robotics, and Automation in Space, Tokyo, Japan.
- Wyatt, E. J., et al., 1997, “An Overview of the Beacon Monitor Operations Technology,” International Symposium on Artificial Intelligence, Robotics, and Automation in Space, Tokyo, Japan.
- Rayman, M. and D. Lehman, 1997, “Deep Space One: NASA’s First Deep-Space Technology Validation Mission,” 48th International Astronautical Congress, Turin, Italy.
- Sherwood, R., et al., 1997, “Flight Software Implementation of the Beacon Monitor Experiment On the NASA New Millennium Deep Space 1 (DS1) Mission,” Second International Symposium on Reducing the Cost of Spacecraft Ground Systems and Operations, Oxfordshire, UK.
- Chien, S., et al., 1997, “Resource Scheduling for a Network of Communications Antennas,” Proceedings of the IEEE Aerospace Conference, Aspen, CO.
- Staehele, R. L., et al., 1996, “Pluto Express: Advanced Technologies Enable Lower Cost Missions to the Outer Solar System and Beyond,” International Low Cost Planetary Missions Conference, Laurel, MD.
- Wyatt, E. J. and J. B. Carraway, 1995, “Beacon Monitoring Approach to Spacecraft Mission Operations,” First International Symposium on Reducing the Cost of Spacecraft Ground Systems and Operations, Oxfordshire, UK.


Appendix A. List of Telemetry Channels and Names

Channel #	Description	
A-0259	ACS_TELEM_ALLOCATED_ENTRY_59	X
A-0534	ACS_TELEM_ALLOCATED_ENTRY_363	X
A-0563	ACS_TELEM_ALLOCATED_ENTRY_117	X
A-0762	ACS_TELEM_ALLOCATED_ENTRY_149	X
A-1619	ACS_TELEM_ALLOCATED_ENTRY_182	X
A-1621	ACS_TELEM_ALLOCATED_ENTRY_184	X
A-1622	ACS_TELEM_ALLOCATED_ENTRY_188	X
B-2014	FSC_IPCU_VME_N15_SUP_VOLT_MEAS	X
B-2040	FSC_BTf_SOFTWARE_VERSION_MEAS	X
B-4001	FSC_RAD6000_TEMP_MEAS	X
B-4004	FSC_UDL_OSC_TEMP_MEAS	X
D-0900	DWN_PRYOR_STATE_0	
F-0692	FPR_SYMPTOM_SUMMARY1	
F-1098	MON_ACS_INFO_EHA_MDC_STATE	X
G-4001	FSC_PEPE_TEMP1_MEAS	X
G-4002	FSC_PEPE_TEMP2_MEAS	X
G-4003	FSC_PEPE_CALORIMETER_TEMP_MEAS	X
I-4002	FSC_MICAS_OPT_BENCH_NXNZ_TEMP_MEAS	X
I-4003	FSC_MICAS_OPT_BENCH_PYNZ_TEMP_MEAS	X
I-4004	FSC_MICAS_M1_MIRROR_TEMP_MEAS	X
I-4006	FSC_MICAS_OPT_BENCH_CUBE_TEMP_MEAS	X
I-4007	FSC_MICAS_IR_DET_TEMP_MEAS	X
I-4008	FSC_MICAS_UV_DET_TEMP_MEAS	X
I-4010	FSC_MICAS_COVER_MECH_TEMP1_MEAS	X
N-0141	NAV_EHA_WHICH_MACHINE_RUNNING	X
O-4001	FSC_UPPER_BUS_TEMP1_MEAS	X
O-4002	FSC_UPPER_BUS_TEMP2_MEAS	X
P-0020	FSC_BATTERY_1_SOC	X
P-0022	FSC_BATTERY_2_SOC	X
P-2002	FSC_BATTERY_VT_MODE_MEAS	X
P-2010	FSC_BATTERY_MID_VOLT_1_MEAS	X
P-2011	FSC_BATTERY1_CURRENT_MEAS	X
P-2020	FSC_BATTERY_MID_VOLT_2_MEAS	X
P-2021	FSC_BATTERY2_CURRENT_MEAS	X
P-2030	FSC_SCARLET_VOLT_MEAS	X
P-2031	FSC_SCARLET_VAL_MOD_CUR_1_MEAS	X
P-2032	FSC_SCARLET_VAL_MOD_VOLT_1_MEAS	X
P-2040	FSC_SCARLET_WING1_CUR_MEAS	X
P-2050	FSC_SCARLET_WING2_CUR_MEAS	X
P-2060	FSC_PDU_ESS_BUS_CUR_MEAS	X
P-2061	FSC_PDU_ESS_BUS_VOL_MEAS	X
P-2062	FSC_PDU_NEb1_CUR_MEAS	X
P-2063	FSC_PDU_NEb1S_CUR_MEAS	X
P-2064	FSC_PDU_NEb2_CUR_MEAS	X
P-2065	FSC_PDU_NEb3_CUR_MEAS	X
P-2070	FSC_PDU_RELAY_FET_STATUS_WORD0_MEAS	X
P-2071	FSC_PDU_RELAY_FET_STATUS_WORD1_MEAS	X
P-2072	FSC_PDU_RELAY_FET_STATUS_WORD2_MEAS	X
P-4011	FSC_BATTERY_TEMP1_MEAS	X
P-4021	FSC_BATTERY_TEMP2_MEAS	X

Channel #	Description	
P-4022	FSC BATTERY CHARGE TEMP MEAS	x
P-4041	FSC SCARLET WING1 VAL TEMP1 MEAS	x
P-4042	FSC SCARLET WING1 VAL TEMP2 MEAS	x
P-4043	FSC SCARLET WING1 VAL TEMP3 MEAS	x
P-4044	FSC SCARLET WING1 VAL TEMP4 MEAS	x
P-4045	FSC SCARLET WING1 VAL TEMP5 MEAS	x
P-4046	FSC SCARLET WING1 VAL TEMP6 MEAS	x
P-4047	FSC SCARLET WING1 VAL TEMP7 MEAS	x
P-4048	FSC SCARLET WING1 VAL TEMP8 MEAS	x
P-4051	FSC SCARLET WING2 VAL TEMP1 MEAS	x
P-4052	FSC SCARLET WING2 VAL TEMP5 MEAS	x
T-0001	FSC SDST_XPDR STATE MEAS	x
T-0014	FSC SDST_X_PWR MEAS	x
T-0024	FSC SDST_EXCITER_SPE MEAS	x
T-2015	FSC PDU_SDST_CUR MEAS	x
T-2016	FSC PDU_KASSPA_CUR MEAS	x
T-2017	FSC PDU_XSSPA_CUR MEAS	x
T-4002	FSC XSSPA_TEMP MEAS	x
V-2005	ACS N2H4 TANK PRSS MEAS	x
V-4001	FSC PROP_MOD_TEMP1 MEAS	x
V-4002	FSC PROP_MOD_TEMP2 MEAS	x
V-4003	ACS N2H4 TANK_TEMP1 MEAS	x
V-4011	ACS RCS_CLUSTER1_TEMP MEAS	x
V-4012	ACS RCS_CLUSTER1_CAT_TEMP MEAS	x
V-4021	ACS RCS_CLUSTER2_TEMP MEAS	x
V-4022	ACS RCS_CLUSTER2_CAT_TEMP MEAS	x

Appendix B. DS1 Technology Validation Power On/Off Times

Date	Experiment Type
Jan. 6, 1999	X-tone, 20, 30, 25, & 35 kHz
Feb. 4, 1999	X-tone, 35 & 20 kHz
Feb. 19, 1999	Data Summarization turned on
Feb. 26, 1999	B-tone & X-tone
Mar. 3, 1999	X-tone, 35 & 20 kHz
Mar. 8, 1999	Data Summarization turned on
Mar. 18, 1999	X-tone, 30, 20, 25, & 35 kHz
Mar. 22, 1999	Data Summarization turned on
Mar. 24, 1999	X-tone
Apr. 7, 1999	X-tone, 20, 25, 30, & 30 kHz
Apr. 13, 1999	B-transmit & X-tone, 20, 25, 30, & 35 kHz
Apr. 19, 1999	Ka-tone
Apr. 20, 1999	B-transmit
Apr. 26, 1999	Ka-tone, 20, 25, 30, & 35kHz
Apr. 27, 1999	B-transmit
June 1999 - May 2000	During this period, beacon tone passes were done just about every week and data summarization was left on continuously.



Ion Propulsion System (NSTAR) DS1 Technology Validation Report

John R. Brophy, Roy Y. Kakuda, James E. Polk, John R. Anderson,
Michael G. Marcucci, David Brinza, Michael D. Henry, Kenneth K. Fujii,
Kamesh R. Mantha, John F. Stocky
*Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California 91109*

James Sovey, Mike Patterson, Vince Rawlin, John Hamley
*Glenn Research Center
Cleveland, Ohio*

Tom Bond, Jon Christensen, Hap Cardwell, Gerald Benson, Joe Gallagher
*Hughes Electron Dynamics Division
Torrance, California*

Mike Matranga
*Spectrum Astro, Inc.
Gilbert, Arizona*

Duff Bushway
*Moog, Inc.
Aurora, New York 13026*



Table of Contents

<u>Section</u>	<u>Page</u>
Extended Abstract	iv
Fact Sheet	vii
1.0 Introduction	1
2.0 Technical Description	1
2.1 The NSTAR Ion-Propulsion System	1
2.2 Key Technology-Validation Objectives	6
2.3 Expected Performance Envelope.....	6
2.4 Detailed Description.....	6
2.5 Technology Interdependencies.....	8
2.6 Test Program	16
3.0 Technology Validation Summary	38
4.0 Future Applications	38
5.0 Acknowledgment	39
6.0 List of References	39

Figures

<u>Figure</u>	<u>Page</u>
Figure 1. Functional Block Diagram of the NSTAR Ion Propulsion System.....	2
Figure 2. Diagram of the NSTAR Ion Engine (with the plasma screen removed).....	3
Figure 3. PPU–Block Diagram	4
Figure 4. NSTAR Power-Throttling Strategy	7
Figure 5. NSTAR Thruster Thermal Environment on DS1.....	9
Figure 6a. Gimbal 1 Slew at Start Up and Recycle @ Mission Level 6	12
Figure 6b. Gimbal 2 Slew at Start Up and Recycle @ Mission Level 6	12
Figure 7. IPS Acceptance Test 2, X-Band Signal to Noise	13
Figure 8. Location of Diagnostics Hardware on the DS1 Propulsion Module	14
Figure 9. QCM Deposition Rate vs. Time and Mission-Throttle Level.....	15
Figure 10. QCM Deposition Rate Increases with Mission-Throttle Level.....	15
Figure 11. Comparison of Measured BOL Thrust with Calculated Thrust at BOL and EOL	17
Figure 12. Engine Efficiency as a Function of Time and Power Level during the 8,000-hr Test of EMT2	18
Figure 13. Margin Between the NSTAR Neutralizer Flow Rates and the Transition from Spot to Plume Mode.....	18
Figure 14. Beam Current Density Distribution Measured at BOL for Six Throttle Levels (EMT2).....	19
Figure 15. Beam Potential Measurements at BOL for Six Throttle Levels (EMT2).....	19
Figure 16. Variation in Electron-Backstreaming Voltage at Full Power over the Course of the 8,000-hr Test (EMT2).....	20
Figure 17. Accelerator Grid Aperture Diameters Measured after the 8,000-hr Test Indicate Significant Enlargement from Their Original 1.14-mm Diameter Values.....	21
Figure 18. Electron-Backstreaming Voltage Margin at the End of the 8,000-hr Test (EMT2).....	21
Figure 19. FT2 Efficiency Versus Power during ELT	22
Figure 20. Comparison of Electron-Backstreaming Limits for FT2 and EMT2	22
Figure 21. Comparison of the Long-Term Behavior of the Purveyance Margin for FT2 and EMT2	22
Figure 22. Comparison of the Long-Term Behavior of the Discharge Voltage for FT2 and EMT2.....	23
Figure 23. NSTAR Ion-Thruster Discharge-Propellant Utilization Efficiency	26
Figure 24. The Difference Between Measured and Calculated Thrust over the NSTAR Throttle Range.....	27
Figure 25. Difference Between a Given Power Level and the Beginning-of-Life Power	27
Figure 26. In-Flight Measurements of PPU Efficiency Compared to Ground Test Data	28
Figure 27. Difference Between a Given Input Power to the Flight PPU on DS1 and the Corresponding Throttle Table BOL Value	30
Figure 28. Difference Between Measured and Calculated Thrust in Flight Compared to Ground Measurements	30
Figure 29. Thrust Measured in Flight as a Function of PPU Input Power Compared to the Throttle Table Values	31

Figure 30. In-Flight Gimbal Positions as a Function of Mission Throttle Level.....	32
Figure 31. Example of Flow Rate Throttling	32
Figure 32. Time History of Peak Cathode and Neutralizer Heater Voltages in Flight.....	33
Figure 33. Time History of Cathode and Neutralizer Ignition Delays in Flight.....	34
Figure 34. Example of In-Flight Throttle-up and Throttle-down Sequences	34
Figure 35. Discharge Loss Measured in Flight Compared to the Throttle Table Values	34
Figure 36. Discharge Voltage Measured in Flight Compared to the Throttle Table Values and Ground Test Measurements.....	35
Figure 37. Discharge Current Measured in Flight Compared to the Throttle Table Values.....	35
Figure 38. Accelerator Grid Impingement Current Measured in Space Compared to Ground Test Measurements	36
Figure 39. In-Space Ratio of Accelerator Grid Impingement Current to Beam Current.....	36
Figure 40. Neutralizer Keeper Voltage Measured in Space and in Ground Tests.....	37
Figure 41. Neutralizer Common Voltage Measured with-Respect-to Spacecraft Ground in Space and with-Respect-to Facility Ground in Ground Tests.....	37

Tables

<u>Page</u>	<u>Page</u>
Table 1. NSTAR IPS Component Masses.....	2
Table 2. PPU Power Supply Requirements.....	5
Table 3. Derived Objectives from the QFD Process	7
Table 4. Table of Expected End-of-Life Performance	8
Table 5. Approximate Hydrazine Consumption Per Activity	10
Table 6. MICAS Image-Noise Comparison with the IPS On and Off	11
Table 7. NSTAR Project Tests (NPT).....	16
Table 8. NSTAR Development Tests.....	16
Table 9. NSTAR Characterization Tests.....	24
Table 10. NSTAR Engineering Development Tests	25
Table 11. Flight Throttle Table of Parameters Controlled by the DCIU.....	25
Table 12. Flight Engine Performance Measured in Space	33
Table A1. IGOR Data Channels.....	41
Table B1. Beam On/Off Time.....	41

Appendices

<u>Appendix</u>	<u>Page</u>
Appendix A. List of Telemetry Channels and Names.....	41
Appendix B. Date of Turn-on/off and Frequency of Data Capture.....	44

EXTENDED ABSTRACT

1.0 TECHNOLOGY VALIDATED

The *Deep Space 1* (DS1) spacecraft uses a single-engine, xenon ion propulsion system, provided by the NASA Solar electric propulsion Technology Applications Readiness (NSTAR) project, for primary on-board propulsion.

Technology-validation requirements for the NSTAR Project were developed early in the project life cycle. A quality functional deployment (QFD) exercise conducted in 1993 resulted in a documented set of user, customer, stakeholder, and sponsor needs that the NSTAR Project needed to satisfy in order to be declared successful. All items from that complete list are shown in this report along with the benchmark data that was demonstrated in flight. One of the prime objectives of the project was to satisfy future users that this technology was flight-proven; therefore, retiring the perceived risk issues was a significant part of the validation effort. The details of these efforts are described in the full report. Some of these important issues were retired through an extensive ground test program while the others were retired through the flight test on DS1.

2.0 RISKS ASSOCIATED WITH THIS TECHNOLOGY

The following key risks were addressed by the NSTAR project as part of ground testing and during the flight of the ion propulsion system on DS1:

1. Adequate engine life—Prior to the NSTAR project, no ion engine intended for primary propulsion had ever been successfully operated for its full design life.
2. Guidance, navigation and control (GN&C) of a solar-electric propulsion (SEP) spacecraft—The low-thrust nature of SEP, together with large solar arrays, makes GN&C sufficiently different from conventional deep-space spacecraft that this is a significant risk area.
3. Mission operation costs—SEP systems require the propulsion system to operate continuously for long periods of time, leading some observers to project that a standing army of propulsion and power engineers would be required to operate the spacecraft, resulting in high-mission operations costs.
4. Spacecraft contamination by the SEP system—Slow erosion of the engine results in a non-propellant efflux from the thruster that could contaminate sensitive spacecraft surfaces.
5. SEP impacts on science instruments—The charge-exchange plasma generated by the operation of the SEP system is easily detected by on-board plasma instruments.
6. SEP impacts on communication—The charge-exchange plasma generated by the operation of the SEP system,

as well as the primary beam plasma, could affect the transmission or reception of electromagnetic waves.

7. Electromagnetic compatibility (EMC) of the SEP system with the spacecraft—The high-power nature of SEP and the use of strong permanent magnets in the ion engines could make it difficult for the SEP system to be electromagnetically compatible with the spacecraft.

How these risks were successfully retired is discussed in the full report.

3.0 VALIDATION OBJECTIVES AND APPROACH

The NSTAR project was designed to overcome the barriers preventing the use of SEP on deep-space missions and enable ion propulsion to enter the mainstream of deep-space propulsion options. To accomplish this, the project had to achieve two major objectives:

1. Demonstrate that the NASA 30-cm diameter ion engine had sufficient life and total impulse capability to perform missions of near-term interest.
2. Demonstrate through a flight test that the ion propulsion system hardware and software could be flight qualified and successfully operated in space and that control and navigation of an SEP-based spacecraft could be achieved.

To demonstrate sufficient engine life, the ground test program was designed to first demonstrate 100% of the engine design life and, subsequently, to demonstrate 150% of the engine life. The flight of the NSTAR system on DS1 addressed the integration, compatibility, and operations issues associated with the use of SEP on a deep space mission.

4.0 TEST PROGRAM

The NSTAR test program employed an extensive ground test activity together with the flight test on DS1 to validate the ion propulsion technology.

The NSTAR ground test program was planned around the use of engineering model thrusters (EMTs) built by NASA Glenn Research Center (GRC) and eventually flight model thrusters fabricated by Hughes Electron Dynamics (HED). A total of four EMTs and two sets of flight hardware—consisting of thrusters, power processor units (PPUs), and digital interface & control units (DCIUs)—were fabricated and tested. In addition, the NSTAR project designed and fabricated an engineering model xenon feed system. The flight xenon control assembly (XCA) was fabricated by Moog. The four EMTs enabled a series of more than 40 engineering tests that addressed wear mechanisms, thermal behavior, mechanical fidelity, low-power performance, and, finally, lifetime in order to instill confidence in the thruster design. An 8000-hour life test demonstrated—for the first

time in history—that an ion engine for primary propulsion could be successfully operated for its full design life.

The two sets of flight units were subjected to acceptance and qualification testing, after which selected flight units were delivered to the spacecraft for the DS1 test program and, ultimately, for flight. The spare flight set is, as of this writing, being used in an extended life test to demonstrate 150% of the engine design life.

5.0 TEST RESULTS

Ground Tests

Early tests of the GRC-built engineering model thrusters validated an initial set of design features and enabled measurement of engine-component wear under a variety of thruster operating conditions. A 2000-hour test of EMT1 led to design improvements that were successfully verified in a subsequent 1000-hour test of this thruster. These tests resulted in a final design that was incorporated into the second engineering model thruster, EMT2. This thruster was used in the Life Demonstration Test (LDT), which was designed to operate the thruster for 8000 hours at full power.

The LDT was the most successful endurance test of a high-power ion engine ever performed. A total of 8,192 hours of operation were achieved at an input power of 2.3 kW with a specific impulse of 3200 s before it was voluntarily terminated. A total of 88 kg of xenon propellant was processed, demonstrating a total impulse of 2.73×10^6 N-s. Risks associated with neutralizer lifetime, thrust performance degradation, engine efficiency degradation, material deposition, thrust vector drift, electrode wear, long-term thermal characteristics, and initial start-up conditions were successfully retired by this test.

The last major test in the NSTAR project plan is the Extended Life Test (ELT), which is designed to demonstrate 150% of the engine design life using the DS1 flight spare engine (FT2). The engine design life is most easily expressed in terms of the total amount of xenon propellant that the thruster can process. For the NSTAR project, the engine design life is 82 kg of xenon, which corresponds to about 8,000 hours of operation at full power. To demonstrate 150% of the engine life, therefore, requires a test in which approximately 125 kg of xenon is processed by the engine. A secondary objective of this test is to demonstrate extended operation at throttled conditions since the previous project-level life tests had all been performed at the full-power point. It is believed that the full-power point is the most stressing to the engine; however, the ELT is designed to obtain the data necessary to support this assertion.

As of this symposium (February 2000), the ELT has operated FT2 for more than 8,000 hours covering three

different throttle levels and has processed more than 75 kg of xenon. The test is scheduled to demonstrate the 125-kg throughput by the end of the year. The Deep Space Exploration Technology program is considering extending this test to determine the actual thruster end-of-life. This would significantly benefit the potential future users listed in Section 6.0 below.

Flight Test

Aside from an initial hiccup, the operation of the NSTAR ion propulsion system (IPS) on DS1 has been flawless.

The initial hiccup occurred 4.5 minutes after the engine was first started in space when continuous high-voltage recycling caused the thruster to shutdown. Subsequent troubleshooting efforts identified that the fault was most likely due to a piece of conductive debris lodged between the grids. To dislodge this debris, the spacecraft was turned several times to move the ion engine in and out of the Sun. This results in thermally cycling of the engine's ion accelerator system causing the electrodes to move relative to one another. Subsequently, another start attempt was made at thirty-one days after launch. The engine started normally and has operated perfectly since this time.

As expected, operation of the ion engine, PPU, and xenon feed system in space produced performance that closely matched that measured on the ground. In addition, the flight on DS1 enabled the following resolution of the key risk areas listed earlier:

1. Guidance, navigation and control—The operation of the SEP system on DS1 demonstrated that GN&C is not more difficult with an SEP spacecraft, just different.
2. Mission operation costs—The electrical nature of SEP lends itself well to autonomous operation, resulting in essentially no significant increase in mission operations cost for SEP vehicles.
3. Spacecraft contamination—Data from DS1 indicates that this efflux travels largely in line-of-sight from the engine and does not pose a significant health risk to a properly designed spacecraft.
4. SEP impact on science instruments—DS1 showed that the low-energy, charge-exchange plasma generated by the operation of the ion engine does not interfere with measurements of the much more energetic solar wind plasma
5. SEP impacts on communication—No impact of the SEP system operation on communications with DS1 could be detected.
6. Electromagnetic compatibility (EMC) of the SEP system with the spacecraft—DS1 showed that while this issue requires careful engineering, it is an easily tractable problem.

6.0 APPLICABILITY AND POTENTIAL FUTURE BENEFITS

Many missions have been identified by JPL's advanced mission planning activity as being either enabled or strongly enhanced by the use of solar electric propulsion. These were based on NSTAR or derivatives of the NSTAR ion propulsion technology, including: Comet Nucleus Sample Return, Mercury Orbiter, Neptune Orbiter, Titan Explorer, Saturn Ring Observer, Europa Lander, and Venus Sample Return.

To illustrate the benefits enabled by the use of an NSTAR-derivative SEP system for a mission to a comet, the performance of a SEP-based spacecraft to the comet 46P/Wirtanen is compared to ESA's chemical propulsion-based Rosetta mission to the same target. The Rosetta spacecraft has an initial wet mass of 2,900 kg and is launched on an Ariane 5. This spacecraft takes more than 9 years to reach the comet, arrives with a net spacecraft

mass of 1300 kg, and does not return a sample from the comet. An SEP-based spacecraft, on the other hand, with an initial wet mass of 1830 kg, could be launched on a Delta IV medium launch vehicle. The SEP system would take only 2.6 years to deliver a 1300-kg spacecraft to the comet. The same SEP system could then return the spacecraft and a comet sample to Earth in an additional 4.5 years. Thus, the SEP-based spacecraft could travel to the comet and return to Earth in less time than it takes for a chemical-propulsion-based spacecraft to fly to the comet!

7.0 CONCLUSIONS

The success of the NSTAR SEP system on the DS1 spacecraft, as well as the success of the NSTAR engine life test program, has resulted in SEP now becoming a legitimate propulsion option for deep space missions. The project's successful validation effort now enables exciting new missions to benefit from the substantial performance capabilities of ion propulsion.



NSTAR Fact Sheet



The NSTAR project and DS1 successfully validated ion propulsion enabling exciting new missions to benefit from the substantial performance capabilities of this technology

Flight Engine Performance Measured in Space

NSTAR Throttle Level	Mission Throttle Level	PPU Input Power (kW)	Engine Input Power (kW)	Measured Thrust (mN)	Main Flow Rate (sccm)	Cathode Flow Rate (sccm)	Neutralizer Flow Rate (sccm)	Specific Impulse (s)	Total Efficiency
12	85	1.99	1.86	75.34	19.99	2.91	2.82	3035	0.602
11	83	1.94	1.82	72.55	18.63	2.75	2.67	3125	0.610
11	83	1.96	1.83	72.63	18.62	2.75	2.67	3131	0.609
10	77	1.84	1.72	69.54	18.59	2.75	2.67	3000	0.594
10	76	1.82	1.70	67.21	17.31	2.58	2.51	3109	0.602
10	75	1.79	1.68	66.81	17.33	2.58	2.51	3087	0.601
10	74	1.77	1.66	66.11	17.33	2.59	2.51	3054	0.595
10	73	1.75	1.65	65.64	17.31	2.59	2.51	3035	0.594
10	72	1.73	1.63	65.15	17.31	2.59	2.51	3012	0.592
9	69	1.67	1.57	62.27	16.08	2.50	2.43	3070	0.597
6	48	1.29	1.22	47.43	11.42	2.50	2.42	3006	0.573
6	48	1.29	1.22	47.39	11.44	2.49	2.42	3004	0.571
3	27	0.89	0.84	31.70	6.93	2.50	2.43	2770	0.511
0	6	0.50	0.48	20.77	6.05	2.50	2.43	1961	0.418

Thruster

Thrust 19.0 to 92.7 mN
Input Power 423 to 2288 W
Isp 1814 to 3127 s
Mass 8.33 kg
Xe Throughput Demonstrated 88 kg
Xe Throughput Planned 125 kg by the end of Y2K
Manufacturer Hughes Electron Dynamics

Power Processing Unit (PPU)


Input Power 474 to 2522 W
Efficiency 0.92 to 0.94
Input Voltage 80 to 160-VDC
Mass 13.3 kg
Manufacturer Hughes Electron Dynamics

Digital Interface & Control Unit (DCIU)

S/C Interface 1553
PPU Interface RS-422
Mass 2.47 kg
Manufacturer Spectrum Astro, Inc.

Xenon Feed System

Flow Rate Accuracy +/-3%
Cathode, Neutralizer Flow 2.39 to 3.7 sccm
Main Flow 5.98 to 23.43 sccm
Xenon Control Assembly (XCA) 7.78 kg
XCA Manufacturer Moog, Inc.
Plenum Tanks (two) 1.5 kg (each)
Tank Assembly Mass 7.94 kg
Tank Volume 49.2 liters
Mass of Xenon Stored 81.5 kg
Main Tank Manufacturer Lincoln Composites



JPL 2/00

Ion Propulsion System (NSTAR) DS1 Technology Validation Report

*John R. Brophy, Roy Y. Kakuda, James E. Polk, John R. Anderson, Michael G. Marcucci, David E. Brinza,
Michael D. Henry, Kenneth K. Fujii, Kamesh R. Mantha, John F. Stocky
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California*

*James Sovey, Mike Patterson, Vince Rawlin, John Hamley
Glenn Research Center, Cleveland, Ohio*

*Tom Bond, Jon Christensen, Hap Cardwell, Gerald Benson, Joe Gallagher
Hughes Electron Dynamics Division, Torrance, California*

*Mike Matranga
Spectrum Astro, Inc., Gilbert, Arizona*

*Duff Bushway
Moog, Inc., Aurora, New York*

1.0 INTRODUCTION

The first use of solar-electric propulsion (SEP) on a deep-space mission began with the launch of the Deep Space 1 (DS1) spacecraft on October 28, 1998. This marks a milestone in the development of advanced propulsion for deep-space missions. The DS1 spacecraft uses a single xenon-ion engine, provided by the NASA Solar electric propulsion Technology Applications Readiness (NSTAR) project, as the primary onboard propulsion system. This propulsion system is designed to deliver a total ΔV of 4.5 km/s to DS1 while using only 81 kg of xenon.

The NSTAR project was designed to overcome the barriers preventing the use of SEP on deep-space missions and enable ion propulsion to enter the mainstream of deep-space propulsion options. To accomplish this, the project had to achieve two major objectives:

1. Demonstrate that the NASA 30-cm diameter ion engine has sufficient life and total-impulse capability to perform missions of near-term interest.
2. Demonstrate through a flight test that the ion-propulsion system hardware and software could be flight qualified and successfully operated in space and that control and navigation of an SEP-based spacecraft could be achieved.

By all measures, these objectives have been met with unqualified success. Aside from an initial hiccup, the operation of the NSTAR ion propulsion system (IPS) on DS1 has been flawless: the IPS successfully provided the ΔV required for the July 29, 1999 flyby of the asteroid Braille. Consequently, ion propulsion is now a credible propulsion option for future deep-space missions. Details of how the NSTAR ion-propulsion technology was validated for deep-space missions are given in the sections that follow. This report is a summary version of the full NSTAR Flight Validation Report given in Reference [1].

2.0 TECHNICAL DESCRIPTION

As is rigorously explained in Reference [30], the NSTAR IPS was one of 12 breakthrough technologies to be validated on the DS1 spacecraft. Each was to be validated in different ways depending on the technology usage and would require different periods of time. Through joint planning, the DS1 operators and NSTAR personnel produced a validation plan that fit into the DS1 overall-mission plan. How DS1 was conceived and how the individual validation results were perceived from an overall-spacecraft perspective are also explained in Reference [30]. This paper, therefore, concentrates on the validation results from the technology's standpoint and illustrates some risk-reduction issues that could be applied to future programs.

The NSTAR project developed and delivered an ion propulsion system to DS1 that was based on the NASA 30-cm diameter xenon ion engine. This section provides a description of the NSTAR IPS, the key technology objectives, and a summary of the ground- and flight-test results.

2.1 The NSTAR Ion-Propulsion System

A block diagram of the four major components of the NSTAR IPS is given in Figure 1. The ion thruster uses xenon propellant delivered by the xenon feed system (XFS) and is powered by the power processing unit (PPU), which converts power from the solar array to the currents and voltages required by the engine. The XFS and PPU are controlled by the digital control and interface unit (DCIU), which accepts and executes high-level commands from the spacecraft computer and provides propulsion subsystem telemetry to the spacecraft-data system. To accommodate variations in the solar array output power with distance from the Sun, the NSTAR IPS was designed to operate over an engine-power range of 500 W to 2,300 W. Discrete levels within this range are often referred to as "throttle levels." The mass of the NSTAR IPS is given in Table 1.

Table 1. NSTAR IPS Component Masses

Component	Mass (kg)
Ion Engine	8.33
Power Processing Unit (PPU)	15.03
XFS minus Xenon Propellant Tank	12.81
Xenon Propellant Tank	7.66
Digital Control and Interface Unit (DCIU)	2.47
PPU to Ion Engine Cable	1.70
Total	48.00

2.1.1 Ion Engine—The NSTAR—ion engine produces thrust by ionizing a low-pressure xenon gas (of order ~ 0.1 Pa) and electrostatically accelerating the resulting positive ions. Ion acceleration is accomplished through the use of two closely spaced, multi-aperture electrodes positioned at one end of the engine across which an accelerating voltage of 1.28 kV is applied. The velocity of the ion exhaust is determined by the magnitude of the applied-net-accelerating voltage and the charge-to-mass ratio of the ions. A magnetic field created by rings of permanent magnets is used to improve the efficiency with which the engine ionizes the propellant.

Electrons stripped from the propellant atoms in the ionization process are collected and injected into the positive-ion beam by the neutralizer cathode in order to space-charge neutralize the ion beam and to prevent the spacecraft from accumulating a large negative charge.

The electrostatic-acceleration process is extremely efficient. In practice, the NSTAR ion-accelerator system has an efficiency of converting electrical-potential energy to kinetic energy of $\geq 99.6\%$. This nearly perfect ion acceleration efficiency enables the ion engine to produce a specific impulse of more than 3,000 seconds while maintaining low-engine-component temperatures. It also results in the ion engine being the most efficient type of electric thruster at specific impulses greater than approximately 2,500 seconds. The combination of high efficiency and high specific impulse makes ion engines attractive for a wide variety of mission applications, including north-south station keeping (NSSK) of satellites in geosynchronous orbit, Earth-orbit transfer, orbit repositioning of Earth-viewing spacecraft, and robotic solar-system exploration.

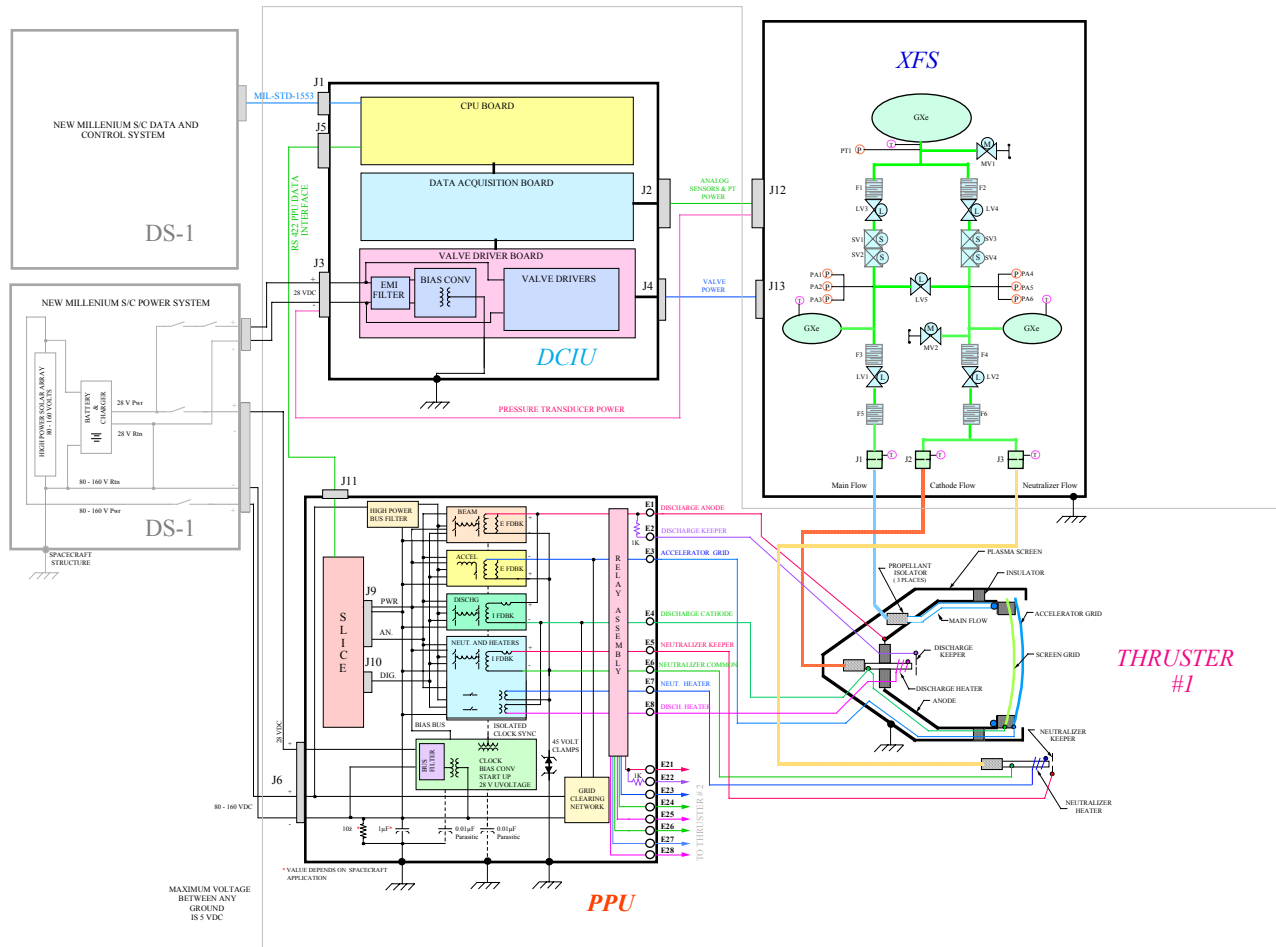


Figure 1. Functional Block Diagram of the NSTAR Ion Propulsion System

A schematic diagram of the NSTAR 30-cm diameter ion engine fabricated by Hughes Electron Dynamics (HED) is shown in Figure 2. The engine is based on technologies developed by NASA [2] and is designed to produce a thrust of 20 mN to 92 mN with a specific impulse of 1950 seconds to 3100 seconds over the input-power range of 500 W to 2,300 W. The engine-design life is 8,000 hours at the full-power-operating point. This is equivalent to a total propellant throughput capability of 83 kg and a total impulse of 2.65×10^6 N-s. The engine is designed to provide this throughput for any throttling profile.

On DS1, in order to maintain the thrust centerline through the spacecraft center of gravity (CG), the thruster is mounted on a 2-axis gimbal ring whose orientation is controlled onboard.

2.1.2 Xenon Feed System (XFS)—The NSTAR xenon-feed system, shown schematically in Figure 1, is designed to store up to 81.5 kg of xenon propellant and provide three separate flow rates to the engine: main flow, cathode flow, and the neutralizer flow. The XFS controls these flow rates to within $\pm 3\%$ over a range of 6 to 24 sccm for the main flow, and 2.4 to 3.7 sccm for the cathode and neutralizer flows. The flow-rate control and accuracy are achieved by controlling the pressure in the two plenum tanks upstream of the three porous-metal-plug flow-control devices (FCDs) labeled J1, J2 and J3 in Figure 1. The pressures in the plena are measured with multiple redundant pressure transducers and controlled with two bang-bang solenoid-valve regulators. The main flow is fed from one plenum, while the cathode and neutralizer-flow lines are manifolded into the other. The FCDs for the cathode and neutralizer are closely matched, so these flows are approximately equal over the

entire throttling range of the engine. The flow rate through each FCD is a function of the upstream pressure and temperature; therefore, each plenum pressure is controlled by commands from the DCIU, which compensates for changes in FCD temperature to achieve the desired-flow rate. Upstream-latch valves serve to isolate the main tank from the rest of the system during launch, while the downstream-latch valves start and stop the flow to the engine during operations.

All of the XFS components except the tanks were assembled into a xenon control assembly (XCA) and mounted on a single plate by Moog, Inc. The FCD assemblies were manufactured by Mott, Inc., and the plenum tanks were manufactured by Structural Composites, Inc. (SCI). The propellant feed lines exit the XCA, cross the gimbal mechanism and attach to the engine with resistoflex fittings. The mass of the XFS given in Table 1 includes the flow-control components, the tubing, the wiring, and the XCA plate.

The xenon is stored in a super-critical state to minimize the storage volume. To maintain a single-phase state throughout the entire mission, it is necessary to maintain a minimum propellant-tank temperature of 20° C. Depending on the propellant load, if the temperature goes below this minimum, the xenon could go into a liquid state that may result in tank slosh or the injection of liquid into the feed system resulting in xenon-flow spikes. To keep the composite xenon-propellant tank from over pressurizing, the maximum temperature limit is set to 50° C. The XFS propellant tank has a volume of 49.2 liters and was manufactured by Lincoln Composites.




Figure 2. Diagram of the NSTAR Ion Engine (with the plasma screen removed)

2.1.3 *Power Processing Unit (PPU)*—The PPU is designed to take an 80 V to 160 V input directly from the solar array and supply the appropriate currents and voltages to start and operate the engine. This large input-voltage range was designed to accommodate the expected variation in solar-array-output voltage resulting from a large variation in spacecraft-Sun distance. The PPU is packaged in an enclosure separate from the DCIU and is designed to be bolted onto the spacecraft bus in an area where its excess heat output can be thermally radiated to space. In addition to the high-voltage input, the PPU requires a 28-VDC input for housekeeping power. Both input-power buses have electromagnetic-interference filters to meet the conducted emission requirements of MIL-STD-461. Enclosed within the PPU is a digital “slice” board that operates an RS422 serial-command and telemetry interface with the DCIU, digitizes the PPU telemetry, and controls the PPU-power supplies based on commands from the DCIU.

During normal-engine operation, the PPU provides four steady-state outputs. The beam voltage, the accelerator-grid voltage, the discharge current, and the neutralizer-keeper current are provided by four power supplies as shown in Figure 3. They are the beam supply, the accelerator supply, the discharge supply, and the neutralizer supply, respectively. In addition, during engine startup the PPU provides heater power to the cathode and neutralizer heaters

and an ignition voltage of 650 V to the cathode and neutralizer-keeper electrodes. The PPU output requirements are summarized in Table 2. The high-voltage input to the PPU is distributed to three inverters operating at 20 kHz that drive these power supplies. The power-supply outputs are routed to internal relays that allow them to be switched to one of two terminal blocks, so that a single PPU could be used to run either of two engines. External power-output cables attached to these terminal blocks route power to the field joint on the DS1 spacecraft.

The PPU contains internal protection for input over- and under-voltage conditions. In addition, each power supply is short-circuited protected. When a short-circuit is detected on the beam or accelerator power supplies, internal logic initiates a recycle event to clear this short, based on the assumption that this short is the result of an arc discharge between the electrodes of the ion-accelerator system. The recycle sequence includes turning both supplies off, ramping the discharge current to 4.0 A, enabling both supplies again, and then ramping the discharge current back to the original setpoint. The PPU also contains a “grid-clearing-circuit,” which can be used to attempt to clear an electrical short-circuit between the accelerator-system electrodes that cannot be cleared by the recycle sequence. This circuit includes relays that place the discharge-power supply across the accelerator-system electrodes. The




Figure 3. PPU-Block Diagram

Table 2. PPU Power Supply Requirements

Power Supply	Parameter
Beam Power Supply Output Voltage Output Current Regulation Mode Ripple	650 to 1100 VDC 0.5 to 1.8 ADC Constant Voltage < 5% of Setpoint, Regulated Parameter
Accelerator Power Supply Output Voltage Output Current Regulation Mode Ripple	–150 to –180 VDC 0 to 0.02 ADC, 0.2 A surge for 100 ms Constant Voltage < 5% of Setpoint, Regulated Parameter
Discharge Power Supply Output Voltage Output Current Regulation Mode Ripple	15 to 35 VDC 4 to 14 ADC Constant Current < 5% of Setpoint, Regulated Parameter
Neutralizer Power Supply Output Voltage Output Current Regulation Mode Ripple	8 to 32 VDC 1 to 2 ADC Constant Current < 5% of Setpoint, Regulated Parameter
Discharge Cathode Pulse Igniter Pulse Amplitude Pulse Duration Rate of Rise Repetition Rate	650 V peak 10 μ s 150 V/ μ s 10 Hz minimum
Discharge Cathode Pulse Igniter Pulse Amplitude Pulse Duration Rate of Rise Repetition Rate	650 V peak 10 μ s 150 V/ μ s 10 Hz minimum
Discharge Cathode Heater Supply Output Voltage Output Current Regulation Mode Ripple	2 to 12 VDC 3.5 to 8.5 ADC Constant Current < 5% of Setpoint, Regulated Parameter
Neutralizer Cathode Heater Supply Output Voltage Output Current Regulation Mode Ripple	2 to 12 VDC 3.5 to 8.5 ADC Constant Current < 5% of Setpoint, Regulated Parameter

discharge-power supply is then commanded to a current of 4.0 A, which is sufficient to vaporize small flakes of conductive material that may be shorting the accelerator system. The flight PPU mass listed in Table 1 includes 1.7 kg for micrometeoroid shielding.

2.1.4 Digital Control and Interface Unit (DCIU)—The DCIU, built by Spectrum Astro, Inc., serves as the data acquisition, control, and communications unit in the IPS and is packaged in a box designed to bolt onto the exterior of the spacecraft. The functions of the DCIU include: acquisition, storage, and processing of the signals from the sensors on

the XFS and telemetry from the PPU slice; control of the valves in the XCA; control of the power supplies in the PPU (through the slice), and communication with the spacecraft data-and-control system. The DCIU executes stored sequences that control IPS-operating modes in response to high-level commands generated on the ground or autonomously by the spacecraft. The DCIU is powered by the 28-VDC spacecraft auxiliary-power bus and contains three half-width VME boards that perform the data acquisition, communications and processing, and valve-drive functions. The communications with the PPU slice occur over an RS422 interface; telemetry commands are

transmitted to the spacecraft on a MIL-STD-1553 interface. The mass of the DCIU shown in Table 1 does not include the weight of the thermal-control hardware provided by the DS1 spacecraft.

2.2 Key Technology-Validation Objectives

There are two key objectives of the NSTAR project:

1. Provide the information necessary to allow a project manager to baseline solar-powered ion propulsion technology on a spacecraft.
2. Stimulate commercial sources of, and uses for, ion-propulsion technology.

The NSTAR Project was started in 1992 to meet these objectives. Ion-propulsion technology had been under development in the laboratory for several decades, yet had never been included in a planetary or Earth-orbital mission application. While there are several different forms of electric propulsion thrusters, the NSTAR electrostatic ion engine design originated in 1960 when Harold Kaufman designed and tested the first broad-beam, electron-bombardment ion engine at NASA's Lewis Research Center (now NASA's Glenn Research Center). Early models of ion thrusters used Cesium or Mercury as propellant; demonstration models were flown in 1964 and 1970 on SERT I and II, among others [3]. While these flights showed that such thrusters could operate in space, they did not show that the thruster system could be built and tested with the reliability standards necessary for a flight mission or that the thruster could demonstrate the lifetime necessary for typical mission applications. Therefore, the NSTAR Project was initiated to validate this technology using a two-pronged approach: a ground-test program that was aimed at validating the full lifetime of the ion engine for future missions and a flight-test program that had the objective of demonstrating the delivery, integration, launch, and operations of flight-quality hardware and software. The overall objective of the entire effort was to produce the test-and-operational data that would allow a future spacecraft project manager to baseline this electric-propulsion system.

From these principal objectives, the NSTAR project developed and prioritized a list of derived objectives using a Quality Functional Deployment (QFD) technique. A QFD report was published May 2, 1995 [4]. This report described in detail the NSTAR QFD process. Many project stakeholders, including sponsors, scientists, and spacecraft managers, must have confidence in ion propulsion for it to be used. The NSTAR Project used QFD to merge the needs of a diverse set of stakeholders into a detailed list of technical requirements. Specifically, QFD allowed NSTAR to focus on the most important tasks as viewed by the future users of SEP. A summary of the prioritized QFD-derived objectives for the NSTAR project is given in Table 3 (with a high rating corresponding to a higher priority).

2.3 Expected Performance Envelope

The expected end-of-life (EOL) performance for the NSTAR IPS is specified at the 16 discrete-throttle levels shown in Table 4. These EOL values were developed based on the 8,000-hr life test of an engineering-model NSTAR engine [5,6].

Power throttling over the 16 NSTAR throttle-level settings is accomplished by varying the beam current at constant-beam voltage for throttle levels from 2 through 15. For NSTAR throttle levels 0 and 1, both the beam current and beam voltage are reduced. This throttling strategy maximizes the engine-specific impulse and efficiency at each power level. The engine-throttling envelope capability (with lines of constant-beam power) is shown in Figure 4. The upper boundary of this envelope represents the allowable maximum-beam voltage; the right-hand boundary represent the maximum allowable beam current. The lower boundary is determined by the ion-extraction capabilities of the ion-accelerator system and represents the minimum beam voltage that the engine can be operated at for a given beam current.

The minimum beam-voltage limit for a given beam current is called the "purveyance limit." The left-hand boundary represents the minimum beam current and is determined primarily by the minimum allowable discharge current. The minimum discharge current is a function of the cathode thermal characteristics. For the NSTAR engine, the minimum discharge current is 4.0 A, resulting in a minimum beam current of 0.5 A. The NSTAR throttle table was designed to run along the top of the engine throttling envelope to maximize the specific impulse and maximize the voltage margin between the beam voltage set point and the purveyance limit. This has the effect of minimizing the thrust at each power level. Other throttling strategies are possible; however, the potential benefits of alternate throttling strategies are highly mission specific.

The second column in Table 4 indicates the "Mission Throttle Level." There are 111 mission-throttle levels even though there are only 16 NSTAR throttle levels. These "extra" throttle levels result from specifying 6 new throttle settings between each NSTAR throttle level. These new "finer" throttle settings are used to take better advantage of the available onboard power and are achieved by reducing the beam voltage in 6 steps of 20 V each at constant beam current between each of the NSTAR macro-throttle levels.

2.4 Detailed Description

More detailed descriptions of the NSTAR hardware may be found in References [2, 5 to 18].

Table 3. Derived Objectives from the QFD Process

Customer Attributes	Rating
Low Life-Cycle Cost	9
Enhance US Industrial Competitiveness	9
Minimum SEP impact on Science Instruments	7.4
Short Interplanetary Cruise	7.4
Low Risk of Ion Propulsion Failure	7
Demonstrated Compatibility of SEP with Spacecraft	7
Compatibility With Small Spacecraft	7
Benefit to Successive Missions	7
Demonstrated Integration and Test of Ion Propulsion	6.4
Maximize Spacecraft Resources for Payload	6.4
Acceptable Development Cost Profile	6
Short Development Cycle	6
Low SEP Recurring Cost	5.6
System Reliability Quantified	5.6
Minimize Tracking Requirements	5.6
Minimal Development Risk	5.4
Simple/Proven Spacecraft Operation	5.4
Multiple Launch Opportunities	5.4
Minimal Cost Uncertainty	5
Minimal Development Schedule Uncertainty	5
Good In-Flight Recovery Options	4.6
Minimize Long Duration Ground Tests	4.4
Capture of Large Mission Set	4
Low-Cost Launch Vehicle	3
Minimize MOS Resources	2.6
Low SEP Non-Recurring Cost	2.4
Flight Heritage of SEP Hardware	2.4
Use Off-the-Shelf Components	1

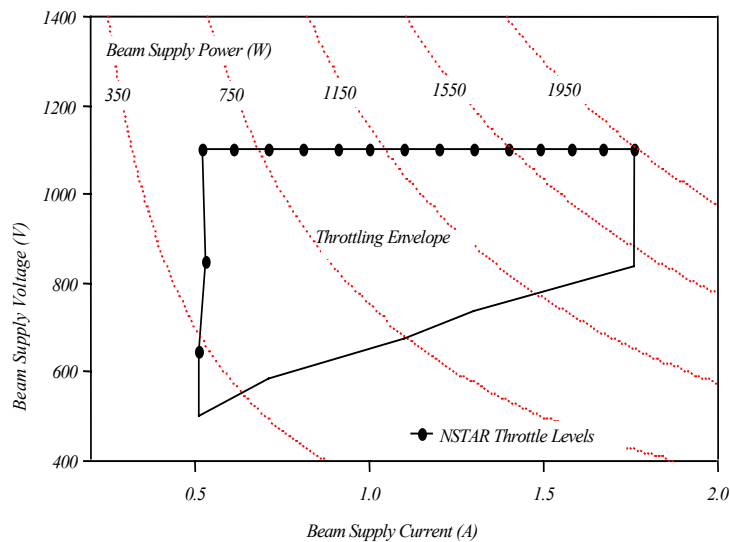


Figure 4. NSTAR Power-Throttling Strategy

Table 4. Table of Expected End-of-Life Performance

NSTAR Throttle Level	Mission Throttle Level	PPU Input Power (kW)	Engine Input Power (kW)	Calculated Thrust (mN)	Main Flow Rate (sccm)	Cathode Flow Rate (sccm)	Neutralizer Flow Rate (sccm)	Specific Impulse (s)	Total Thruster Efficiency
15	111	2.567	2.325	92.67	23.43	3.70	3.59	3127	0.618
14	104	2.416	2.200	87.87	22.19	3.35	3.25	3164	0.624
13	97	2.272	2.077	83.08	20.95	3.06	2.97	3192	0.630
12	90	2.137	1.960	78.39	19.86	2.89	2.80	3181	0.628
11	83	2.006	1.845	73.60	18.51	2.72	2.64	3196	0.631
10	76	1.842	1.717	68.37	17.22	2.56	2.48	3184	0.626
9	69	1.712	1.579	63.17	15.98	2.47	2.39	3142	0.618
8	62	1.579	1.456	57.90	14.41	2.47	2.39	3115	0.611
7	55	1.458	1.344	52.67	12.90	2.47	2.39	3074	0.596
6	48	1.345	1.238	47.87	11.33	2.47	2.39	3065	0.590
5	41	1.222	1.123	42.61	9.82	2.47	2.39	3009	0.574
4	34	1.111	1.018	37.35	8.30	2.47	2.39	2942	0.554
3	27	0.994	0.908	32.12	6.85	2.47	2.39	2843	0.527
2	20	0.825	0.749	27.47	5.77	2.47	2.39	2678	0.487
1	13	0.729	0.659	24.55	5.82	2.47	2.39	2382	0.472
0	6	0.577	0.518	20.69	5.98	2.47	2.39	1979	0.420

2.5 Technology Interdependencies

The ion propulsion system effects the design and performance of many other spacecraft subsystems as well as the mission operations. These subsystems include the solar array, the spacecraft power subsystem, thermal control, attitude control, communications, science instruments, command & control, and navigation. Part of the validation effort was to investigate and measure, if possible, the IPS direct effects on each of these systems.

2.5.1 Power System—The operation of the ion propulsion system is intimately coupled to the spacecraft power system. The IPS is by far the largest load on the power system. The power subsystem is designed to allow the battery to support occasional spacecraft loads during IPS thrusting. This enables IPS operation under transient and short-term negative power-margin conditions that maximizes power utilization. The spacecraft-power system is composed of:

1. A 2500-Watt (@1 AU solar range) concentrator solar array (SCARLET) power source.
2. Two 12-amp-hour (@ ~32 V) batteries to supply energy during power short falls.
3. An high-voltage power conditioning unit (HPCU) that supplies low-voltage power, controls the battery charge and discharge, and adjusts for changes in peak-power voltage.
4. A power distribution unit (PDU) to distribute and switch power.

The solar-array output and the high-voltage bus are tied together and have a voltage range from 80 V to 120 VDC.

To provide maximum power to the IPS during the thrusting phase, the spacecraft has to operate near the peak power

point (PPP) of the array. The spacecraft requires a predetermined minimum level for each mission phase. Based on a projected PPP voltage, an uplink command is sent to the HPCU to have the array’s operating-voltage set point selected slightly greater than the expected PPP. The set-point selection is updated every week during spacecraft tracking.

The IPS is commanded to a throttle level that corresponds to the maximum projected power from the array minus the expected spacecraft power consumption. If the battery is projected to discharge too deeply (defined as reaching 65% State of Charge (SOC) in about 30 minutes), an onboard software algorithm sends an autonomous command to IPS to throttle back one step.

The DS1 flight has shown that although the PPU with a thruster load generated some noise on the high-voltage bus, the high-voltage power-converter unit performed in a stable manner. The design of the HPCU on DS1 allows both the spacecraft avionics and ion propulsion to operate in a stable manner near the PPP of the solar array. This approach relies on a fairly well-defined solar-array model to determine the projected PPP. DS1 demonstrated that collapsing the solar-array voltage (pulling a larger load than was sustainable, resulting in an under-voltage condition) did not damage either the HPCU or the PPU. Onboard flight tests indicate that the HPCU can operate at a set-point voltage greater than the voltage corresponding to the PPP without collapsing the array voltage as long as the battery is capable of handling the needed power. The noise observed on DS1’s high-voltage bus during normal operation is a function of the grounding configuration. A single-point ground approach was used for power-return lines on the spacecraft

with the star ground near the power source. The observed noise could be minimized on future spacecraft through improved routing of ground lines and shields.

2.5.2 Thermal—During IPS operation, the PPU can dissipate up to 200 Watts at 80° C. The top plate of the spacecraft (+Z axis) was used as the PPU and spacecraft’s thermal radiator. The plate could radiate 235 Watts at 80° C and 85 Watts at 0° C. The PPU was designed to operate with baseplate temperatures between –5° C and 50° C with survival-temperature limits of –25° C and 55° C. The PPU was temperature controlled using a combination of 70- and 100-W heaters when not operating. During thrusting, the internal dissipation of the PPU maintains the PPU temperature, with the heaters being required only for operation at the lowest throttle levels. To minimize the power needed to heat the PPU at low throttle levels, the PPU temperature is kept near the lower limit allowed for normal operation.

The DCIU temperature is heater controlled and presents a constrained thermal load to the thermal system. The changing solar aspect angle is the chief driver to a change in thermal operation. The DCIU is designed to operate from –15° C to 50° C with survival limits of –25° C to 55° C.

The XFS temperature is also heater controlled. To minimize the power needed to heat the XFS, the XFS temperatures are kept near the lower limit of normal operation. The flow-control devices are kept above 20° C to maintain their calibration. The Xe propellant tank is kept between 20° C

and 50° C to maintain the super-critical gas state while not over pressurizing the tank.

The thruster is placed inside the conical launch-vehicle adapter within the gimbal rings as shown in Figure 5. During normal IPS operation, the thruster is self-radiating and no additional thermal control is required. The waste heat from the thruster is isolated from the spacecraft and blocked by the gimbal rings and adapter. Consequently, the only significant thermal emission is in the -Z axis (thruster plume direction). The thruster is buried in the launch vehicle adapter such that only the neutralizer is in sunlight when the Sun is perpendicular to the -Z axis. This minimizes the solar load on the thruster. When the Sun is in the -Z axis hemisphere, the solar load increases significantly. To keep the solar load from over heating the thruster magnets, the Sun was not permitted to go closer than 30 degrees to the -Z axis when the thruster was operating at a high power and 1 AU from the Sun.

2.5.3 Attitude Control—The initial and continuous control of the IPS thrust vector was an important IPS validation activity because of its potential to impact the spacecraft’s attitude-control subsystem. When the IPS was not thrusting, 3-axis control of the DS1 spacecraft was accomplished using a blow-down hydrazine system. Each of the three-axis dead bands was controlled to various levels depending upon the mode of operation and hardware constraints. The dead bands were tightened when imaging and loosened to save propellant when in an IPS thrusting or non-thrusting cruise mode.




Figure 5. NSTAR Thruster Thermal Environment on DS1

When starting the IPS, the 3-axis dead bands are set to ± 1 degree. This is done to ensure that attitude control is maintained when stabilizing the control loop during thruster start. After the engine is started, the ACS gimbal is slewed ± 1 to 2 degrees to measure the IPS control torque. Gimbal slews during the initial IPS start up are given in Figure 6a and Figure 6b. This slew procedure is also performed during IPS recycles (indicated by the solid circles) and can be seen following the solid circle just before 1 AM. The slew algorithm during recycles was suppressed later in the mission since the recycles are very short and do not change the gain of the control loop. Note that the thrust level at mission level 6 (NSTAR Throttle Level 0) is only 20 mN and requires the smallest control authority. The attitude-control loop operation was validated for IPS thrust from 20 mN to 78 mN during the initial acceptance test.

The gimbal controller is used to center the thrust through the spacecraft center of mass and maintain the spacecraft attitude along the spacecraft X- and Y-axis. The Z (roll) axis is maintained by the hydrazine thrusters. The X- and Y-axis thruster do not fire once the control loop is stabilized.

Periodically the spacecraft orientation is changed as the gimbal angle deviated from zero degrees. This is done to compensate for a shift in the spacecraft center of mass. The thruster, spacecraft, hydrazine tank, and xenon propellant tank, however, were centered extremely well, eliminating the need for this compensation. Further, the gimbal potentiometer became very noisy as the mission progressed causing an erroneous pointing of the thrust vector. The potentiometer was eliminated from the control loop later in the mission.

Two stepper-motor drives are used to control the gimbal position and can slew the gimbals +6 degrees before running into the mechanical stops. The data from DS1 suggests that the gimbal travel could have been limited to +2 degrees.

The spacecraft attitude-control system (ACS) consumes about 7 grams of hydrazine per day when the IPS is on. In this mode the spacecraft ACS uses the:

- IPS and gimbals to obtain 2-axis control.
- Reaction control subsystem (RCS) to control:
 - The third axis.
 - All major turns.

When IPS is off, spacecraft consumes about 10 grams of hydrazine per day, and the ACS uses the RCS to control:

- The three axes.
- All major turns.

The approximate propellant consumption required for various operations is given in Table 5. Note that the effect of solar distance is ignored.

Table 5. Approximate Hydrazine Consumption Per Activity

RCS Activity	Average Propellant Consumption (gram/day)
IPS thrust on with no OpNav	7
IPS thrust on with 1 OpNav per week	15
IPS off with no OpNav	9.7
OpNav	52
Spacecraft turn to vector	40

2.5.4 Science Instruments—No interference has been observed by the remote sensing instruments when the IPS is thrusting. This was validated by the miniature integrated camera and spectrometer (MICAS) instrument when 3 CCD, 3 APS, 3 IR, and 3 UV exposures were taken with IPS off followed by a second 3 CCD, 3 APS, 3 IR, and 3 UV exposures taken with the IPS on. The IPS was operating at 1 kW with the MICAS pointing well away from the Sun to minimize solar reflection. The results, shown in Table 6, indicate that there is no impact of IPS operation.

The particle and field measurement sensors were mildly affected by the IPS. With the IPS off, the magnetometer from the IPS Diagnostics System (IDS) was used to measure the thruster’s magnetic field. The thruster magnetic field was observed to vary as the gimbal/thruster was rotated. With the IPS on, the IDS magnetometer was able to see the variation in the thruster-produced magnetic field due to the motion of the gimbal/thruster, a change in the thruster power, and variations in the thruster’s magnet temperature. Future magnetometers can correct for the IPS’ magnetic field by incorporating a conventional boom and inboard and outboard magnetometers.

The plasma experiment for planetary exploration (PEPE) instrument was able to measure residual xenon using a mass spectrometer. Future sensors using high-voltage accelerator/detectors may find it necessary to filter the xenon line in their spectra. However, operating the IPS did not interfere with PEPE’s solar-wind measurement.

2.5.5 Communications—The radiative- and conductive-electromagnetic interference of IPS upon the spacecraft and instruments appears to be extremely small. The only interference noted was an increase in telemetry-system noise, mostly due to a spacecraft’s ground loop. X-band transmission through the IPS plume was performed at various angles and IPS power levels. No significant effect

Table 6. MICAS Image-Noise Comparison with the IPS On and Off

IPS State	Star Field	Micas Sensor	Exposure (S)	Pixels or Elements	Minimum Level (dn)	Maximum level (dn)	Mean Signal (dn)	Standard Deviation (dn)
off	1	CCD	0.218	1064960	107	576	134	5.18
on	2	CCD	0.218	1064960	107	283	135	5.16
off	1	CCD	1.750	1064960	106	1025	136	7.27
on	2	CCD	1.750	1064960	104	1002	136	7.20
off	1	CCD	9.830	1064960	106	3284	149	45.80
on	2	CCD	9.830	1064960	105	3248	150	45.30
off	1	APS	0.874	65534	97	158	122	2.93
on	2	APS	0.874	65534	98	153	122	2.69
off	1	APS	1.750	65534	89	155	123	2.71
on	2	APS	1.750	65534	89	159	122	2.86
off	1	APS	4.920	65534	97	151	121	2.83
on	2	APS	4.920	65534	97	154	122	2.91
off	1	IR	0.874	139392	0	3515	287	168.00
on	2	IR	0.874	139392	0	3525	288	168.00
off	1	IR	3.500	139392	0	3512	291	171.00
on	2	IR	3.500	139392	0	3529	291	172.00
off	1	IR	9.830	139392	0	3519	297	176.00
on	2	IR	9.830	139392	0	3526	297	177.00
off	1	UV	4.920	20020	723	815	785	24.60
on	2	UV	4.920	20020	713	805	776	24.10
off	1	UV	14.000	20020	726	950	893	58.40
on	2	UV	14.000	20020	716	946	884	57.90
off	1	UV	28.000	20020	735	1168	1070	95.20
on	2	UV	28.000	20020	723	1165	1060	97.70

was noted during any of the tests. Figure 7 shows that there is no discernible difference in signal to noise when the IPS was throttled at NSTAR level 0, 1, 2, 3, and 4 (500 to 1000 Watts) and from when the IPS was not on. During this test, the low-gain antenna was used for two-way Doppler through the IPS plume, which was pointing at the Earth, and the DSS 55 Block V receiver was used in right-hand circular polarization operating at 8.42 GHz. The vertical scale on Figure 7 is from 12 to 21.3 dB, while the horizontal scale covers from DOY 148 14:37 to 23:15. Note that the IPS was operating at NSTAR throttle level 0 for more than 7 hours before this test to ensure that the thruster was operating in steady state xenon-flow conditions.

2.5.6 Command & Control—The IPS command, control, and telemetry were made very simple to ensure that the integration of IPS to the spacecraft was uncomplicated and the operability by the MOS team straightforward. The basic commands used during normal thrusting were Safe, Standby, Thrust On, Thrust Off, and Throttle Level. A few other commands were used to: initially start up the DCIU, power on the DCIU, perform special diagnostic tests, initially prepare the IPS after launch, and prepare IPS for startup. The control of the IPS was automated so that no monitoring was needed.

The IPS telemetry stream from the DCIU to the spacecraft was composed of a data packet containing all measured IPS parameters sampled every second. This maximum quantity was often filtered by the spacecraft’s telemetry manager to packets each 2 seconds, each 5 seconds, each 5 minutes, etc., in length for insertion into the downlink because of data management issues on board and the robustness of the telemetry link with the ground. IPS data volume was high during critical operating times, such as engine start, and was lower during cruise operations.

2.5.7 Mission Design and Navigation—The DS1 mission design and navigation teams demonstrated that IPS can be reliably flown to multiple planetary targets. Further, the teams have demonstrated that autonomous operation is possible. Since DS1 was the first low-thrust mission, a number of processes had to be modified, tested, and integrated. The first category of process was comparable to conventional mission-design and navigation software:

- Preliminary trajectory-design.
- Intermediate trajectory-design.
- Ground-navigation.




Figure 6a. Gimbal 1 Slew at Start Up and Recycle @ Mission Level 6




Figure 6b. Gimbal 2 Slew at Start Up and Recycle @ Mission Level 6




Figure 7. IPS Acceptance Test 2, X-Band Signal to Noise

To implement autonomous navigation, a number of other processes had to be developed, tested, and integrated. These can be put into three autonomous software categories:

- Orbital determination.
- Trajectory design.
- Command and control.

The low-thrust trajectory program that was used to develop the preliminary heliocentric trajectory neglects the Earth’s mass. To refine the trajectory it was necessary to propagate the launch trajectory using a trajectory program that includes the Earth’s gravity. By propagating the trajectory out of the Earth’s gravitational sphere of influence and determining the spacecraft’s state at that point, a starting point was used to begin the low-thrust trajectory program.

The available IPS power over the mission is required for trajectory optimization. This requires that the solar array and spacecraft’s power be defined as a function of solar distance, aging, and radiation-dose. The solar-array power changes as a function of solar-array temperature, aging, and the spacecraft’s load characteristics. The spacecraft’s power changes as a function of solar distance, and aging, which changes the amount of heater power required to maintain subsystem temperatures.

During the flight of DS1, the trajectory was re-optimized to take into account changes in thrusting profiles. Whenever the original IPS thrust profile was not followed, the trajectory was re-optimized and re-planned with very little performance penalty. In addition, DS1 demonstrated that thrusting does not necessarily need to be in the optimal direction. Many times during the DS1 mission, the thrust was pointed in a direction defined by the convenience of the

mission, instead of the optimal trajectory direction, without a significant penalty.

The mission-design process resulted in a linearized trajectory indicating the trajectory state, thrust, and thrust direction on one-day centers. The trajectory incorporates the effects of thrust-duty cycles, coast periods, and periodic hydrazine drop-off mass. The navigation team used this as a preliminary trajectory to begin the detailed navigation trajectory development.

DS1 used a low-thrust trajectory program called SEPTOP for the preliminary mission design. The program inputs are models of power (solar range, aging, and radiation dose), IPS performance (thrust and mass flow as a function of IPS input power), spacecraft power (as a function of solar range), and initial launch state (position, velocity, and mass) away from the gravitational attraction of Earth. The models for IPS performance are continuous and characterized in SEPTOP as coefficients of a fourth-order polynomial. When the program has an optimal solution, it outputs the power level, thrust, thrust direction, mass flow, and spacecraft state in 1-day increments. Because the inputs into SEPTOP are continuous curves (as defined by the polynomials), the output is also continuous. However, since the IPS has quantized operation, this translation must be done by the navigation software (auto-navigation). The IPS–mission, throttle–table values are used by auto-navigation to select the proper throttle profile (throttle level) over the mission after trajectory has been optimized by SEPTOP. The mission throttle table uses the end-of-life (EOL) value for power, flow rate, and thrust. The mass flow rate and thrust do not change as the thruster ages, so only the IPS input power increases with thruster age.

Computer Algorithm for Trajectory Optimization (CATO) is an intermediate-level trajectory program that could add further fidelity to the trajectory design. It has the capability of adding the gravitational effects of the Earth and the Moon. CATO was used to generate the launch state used by SEPTOP. It was also used to test the fidelity of the SEPTOP trajectory. It was found that the optimization process using a detailed trajectory design was time consuming and did not offer any major benefits.

There are three major navigation tasks: 1) convert the preliminary trajectory received from mission design into a detailed flyable trajectory, 2) determine the current spacecraft and target state (position and velocity) using Doppler, ranging, and optical navigation, and 3) determine the maneuver file needed to fly to the target.

The flight-navigation software is important to IPS validation because in addition to the control of the IPS thrust level and spacecraft thrust vector, it is used to autonomously plan maneuvers over the entire mission. The maneuver plan takes into account the effects of IPS-burn errors, spacecraft-pointing errors, solar pressure, and hydrazine attitude-control maneuvers.

There are a number of mission-margin elements, all of which are interrelated. The major elements are available IPS power, available xenon propellant, thrust profile, and thrust duty cycle. This is somewhat different than chemical propulsion systems, where propellant, interstellar probe (Isp), thrust, and burn time are mission margin elements.

The thrust duty cycle is used as the major control of mission margin. Instead of assuming in the trajectory design that thrusting occurs when permitted, each thrust segment is assumed to have a duty cycle less than 100%. A shortfall of thrust impulse would be corrected by increasing the duty cycle. The duty cycle used by DS1 varied from 90% to 92%. The remaining 8% to 10% is not all usable since a portion is used for optical navigation and downlink of data.

A second mission margin tool is the use of forced ballistic coasts during very efficient thrust periods. The trajectory design program is made to perform coasts during normally optimum-thrust periods. This results in a mission penalty, but ensures that an IPS anomaly that temporarily disrupts thrusting will not threaten the mission.

2.5.8 Contamination—Risking spacecraft contamination by the ion engine’s non-propellant efflux has always impeded the use of ion propulsion. Consequently, the NSTAR project included since its inception the development of a diagnostics package of contamination-monitoring instrumentation to fly with the engine. The location of the NSTAR diagnostic package (NDP) instrumentation relative to the ion engine on DS1 is shown in Figure 8. The NDP

contamination-monitoring instruments include quartz crystal microbalances (QCM) and calorimeters packaged together in a remote sensors unit (RSU). The RSU is located 75 cm from the centerline of the ion thruster’s exhaust beam. One pair of contamination monitors (QCM0, CAL0) has a direct line-of-sight view of the ion engine’s accelerator grid (~85° from the thrust centerline). The other pair (QCM1, CAL1) is shadowed from the ion engine’s accelerator grid by the launch vehicle interface ring on the propulsion module assembly.




Figure 8. Location of Diagnostics Hardware on the DS1 Propulsion Module

The data from QCM0 and CAL0 are consistent with the collection of a total of 250 angstroms of molybdenum from launch through November 1999. These data have not been corrected for solar-illumination and temperature effects on the QCM beat frequency. However, these effects are believed to be minor for QCM0 because the observed change in frequency ($\Delta f > 5,000$ Hz since launch) is much greater than either the solar-illumination effect ($\Delta f < 250$ Hz shadow to maximum illumination) or the thermal effect ($\Delta f < 50$ Hz for $\Delta T \leq 60^\circ$ C in the range $+20^\circ$ C to $+80^\circ$ C). These effects are relatively more important for QCM1 since it has indicated $\Delta f < 500$ Hz since launch.

Of the 250 angstroms of molybdenum collected by QCM0, 100 angstroms were collected in the first 750 hours of NSTAR operation. The deposition rate appears to be well correlated with the Mission Throttle Level, as indicated in Figure 9 and Figure 10.




Figure 9. QCM Deposition Rate vs. Time and Mission-Throttle Level

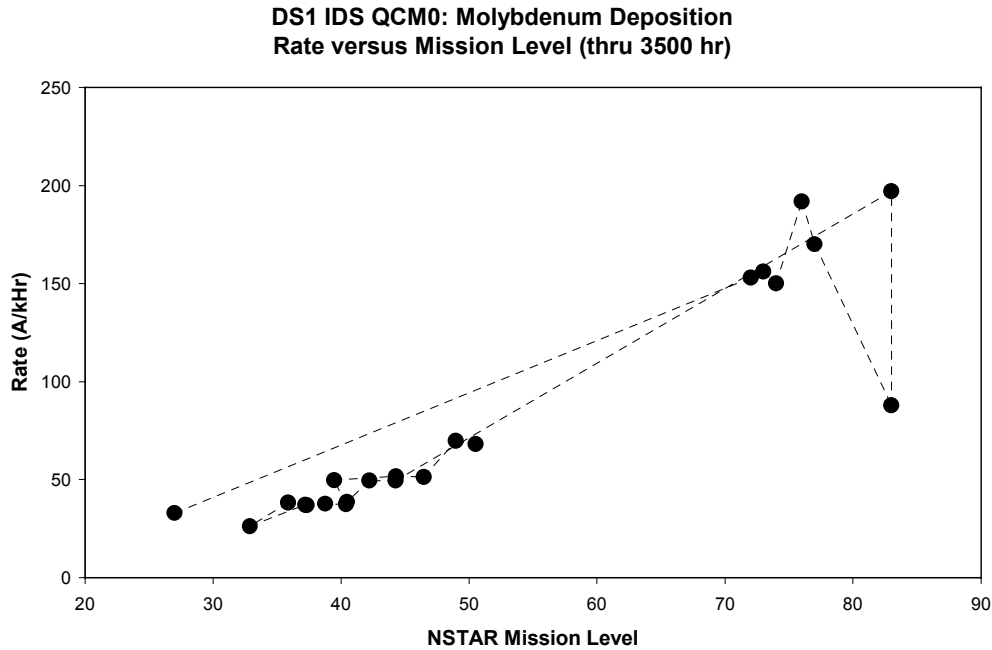


Figure 10. QCM Deposition Rate Increases with Mission-Throttle Level

Based on preliminary analyses of the results from the witness monitors from the 8,000-hr life demonstration test, it is estimated that the molybdenum collection rate during the ground test at the location corresponding to the position of QCM0 is approximately 160 angstroms/kWh. The average molybdenum collection rate for QCM0 on DS1 is

70 angstroms/kWh. Since the average engine power on DS1 is approximately half that of full power (the 8,000-hr test was run at full power) and since the grid erosion rates are expected to scale with engine power, it appears that ground test and flight test deposition rates are comparable.

2.6 Test Program

The NSTAR test program employed an extensive ground-test activity together with the flight test on DS1 to validate the ion propulsion technology.

2.6.1 Ground Test Program—The NSTAR ground-test program was planned around the use of engineering-model thrusters (EMT) build by NASA GRC and, eventually, flight model thrusters fabricated by HED. A total of four EMTs and two flight thrusters were fabricated and tested. The principal objective of the ground-test program was to demonstrate that the NSTAR thruster design had sufficient total-impulse capability and reliability to accomplish deep-space and near-Earth-space missions of near-term interest. The NSTAR project originally included a sequence of four major tests labeled NPT1 through NPT4, as indicated in Table 7. Between NPT1 and NPT3, however, the actual project ground-test history included three other series of tests termed development tests (DTs), engineering

development tests (EDTs), and characterization tests (CTs). These test series were inserted into the NSTAR project in order to provide sufficient information to be confident that the NSTAR thruster and the NSTAR IPS designs would function as promised and with high reliability.

The long duration tests shown in Table 7 were designed to identify unexpected failure modes, characterize the parameters that drive known failure mechanisms, and determine the effect of engine wear on performance. The first test, NPT1, was planned to be 2,000 hours of operation at the full-power point. Failure of a non-flight-type propellant isolator resulted in the test being divided into two test segments: NPT1 and NPT1A. Several potential failure mechanisms were identified in these test segments (see References [17,18] for details). These failure mechanisms were studied in the subsequent shorter duration DTs listed in Table 8.

Table 7. NSTAR Project Tests (NPT)

Test	Purpose	Description	Thruster	Duration (hrs)	Xenon Throughput (kg)
NPT 1	Wear	First 2K	EMT1	867	9.4
NPT 1A	Wear	Finish 2K	EMT1	1163	12.6
NPT 2a	FIT A	PPU integration test	EMT2	21	N/A
NPT 2b	FIT B	PPU integration test	EMT3a	12	N/A
NPT 3	LDT	Life Demonstration Test	EMT2	8194	88
NPT 4	ELT	Extended Life Test	FT2	>12,000*	125*

*Planned

Table 8. NSTAR Development Tests

Test	Purpose	Description	Thruster	Duration (hrs)	Location
DT 1	erosion rate	floating & grounded Screen Grid (SG)	EMT1	37	GRC 5
DT 2	erosion rate	grounded SG	EMT1	50	GRC 5
DT 3	erosion rate	floating SG	EMT1	51	GRC 5
DT 6c	technq accuracy	floating SG—measurement accuracy	EMT1	0.25+	GRC 3
DT 7	mass loss	grounded SG	EMT1	100	GRC 3
DT 16	performance	new grids, backup badges	EMT1a	12	GRC 3
DT 9c	low power perf.	@ low power w/ margin testing	EMT1a	168	GRC 3
DT 18	perf. & margins	second part of old DT 17	EMT1b	50	GRC 5
DT 8a	facility check	with flow sensitivity	FMT	21	JPL148
DT 9b	low power perf.	@ low power w/ margin testing	FMT	870	JPL149
DT 15	revalidation	redesigns for NPT1 issues	EMT1b	1011	JPL148
DT 19	chamber check	replaces DT 17a	J-Series	24	JPL148

As a result of these tests several design changes were made to the engineering model thrusters. The effectiveness of these design changes in eliminating the failure modes identified in NPT1 was then validated in DT15 using EMT1b, which incorporated the design changes. This development test was planned to be a 1,000-hour wear test at the full-power point. Since the failure modes were originally observed in both of the approximately 1,000-hr long NPT1 test segments, the duration of DT15 was selected to be 1,000 hours, with the expectation that this was the shortest test duration that could provide confidence that the failure modes had been eliminated. It was essential to have this confidence prior to starting the endurance test for the full 8,000-hour design life. The development test DT15 was successfully executed and the test was voluntarily terminated after 1,011 hours of operation at full power. Post-test inspection of the thruster indicated that the design changes had successfully eliminated the failure modes observed in NPT1 [18].

2.6.1.1 8,000-hr Life Demonstration Test—Following DT15, the NSTAR project test NPT3, which was designed to demonstrate the full 8,000-hr thruster life, was carried out. This life demonstration test (LDT) used the second engineering-model thruster, EMT2, and was the most successful endurance test of a high-power ion engine ever performed (details of this test are given in [5,6]). A total of 8,192 hours of operation was achieved at the 2.3 kW full-power point before the test was voluntarily terminated. A total of 88 kg of xenon propellant was processed, demonstrating a total impulse of 2.73×10^6 N-s.

Thrust measurements taken over the entire-throttling range at the beginning of the test are shown with calculated beginning-of-life (BOL) values and calculated values at the

end of the 8,000-hr test in Figure 11. The difference between the measured and calculated thrust is less than 1 mN. The calculated thrust is essentially constant as a function of time because the engine conditions that effect the thrust calculation are controlled. The total engine efficiency is given as a function of time over the 8,000-hr test for six throttle levels in Figure 12. These data indicate a slight decrease in engine efficiency over the first 4,000 hours of the test and very little efficiency change over the second half of the test.

Demonstrating adequate life of the neutralizer cathode was one of the key objectives of the 8,000-hr test. To achieve adequate service life of the neutralizer, its operation must be kept in what is referred to as the “spot mode.” This mode of operation is characterized by a relatively low neutralizer-keeper voltage with low-amplitude voltage oscillations. The neutralizer can also operate in what is known as the “plume mode” characterized by a higher neutralizer-keeper voltage and higher amplitude-keeper voltage oscillations. Operation in the plume mode is believed to result in a significantly shortened neutralizer-service life. The operating mode for the neutralizer is determined by the flow rate for a given emission current. The neutralizer operation as a function of flow rate was characterized periodically over the entire throttling range to monitor changes in the flow-rate margin. A certain minimum flow rate and total-emission current are required to prevent plume-mode operation. The flow-rate boundary between stable spot-mode operation and plume mode for the neutralizer over the entire NSTAR throttling range is shown in Figure 13. The difference between the flow rate corresponding to the plume/spot mode boundary and the flow rate specified in the throttle table is the flow-rate margin.




Figure 11. Comparison of Measured BOL Thrust with Calculated Thrust at BOL and EOL




Figure 12. Engine Efficiency as a Function of Time and Power Level during the 8,000-hr Test of EMT2




Figure 13. Margin Between the NSTAR Neutralizer Flow Rates and the Transition from Spot to Plume Mode

The neutralizer cathode was also disassembled and examined for signs of wear and material transport. The only significant wear site was the neutralizer-cathode orifice. The upstream-orifice diameter was essentially unchanged from the pretest value of 0.280 mm, while the downstream end of the orifice increased by 70 percent to 0.48 mm. The surface of the chamfer was observed to be heavily textured from ion bombardment, but no significant dimensional changes have occurred. Small tungsten deposits up to about 10 μm in diameter were found inside the orifice near the upstream entrance. The upstream face of the orifice plate showed no signs of erosion, although a ring of barium deposits was found around the orifice. There was only slight surface texturing on the downstream face of the cathode-orifice

plate and no damage to the weld between the plate and the cathode tube.

The neutralizer-keeper electrode also experienced very little wear. The downstream face and weld show no evidence of sputter damage. The upstream face of the molybdenum keeper has a thin deposit of tungsten around the orifice; this might have come from the neutralizer orifice. A portion of the tantalum-keeper tube was exposed to high-angle-beam ions and shows some surface texturing, but no significant mass loss.

A number of ion-optics performance parameters were measured periodically during the 8,000-hr test at the

nominal- and throttled-operating points. After the test, the grids were examined for signs of wear, including sectioning and detailed SEM measurements of erosion-site geometry. The beam-current density and potential distributions measured about 2.5 cm downstream of the exit plane are shown in Figure 14 and Figure 15. The beam-current density distribution is strongly peaked on the centerline, but drops sharply at a radius of 12 to 13 cm, which is 1 to 2 cm radially in from the periphery of the hole pattern. These profiles did not change significantly over the test and yield average flatness parameters ranging from 0.32 at the minimum power point to 0.46 at full power. The peak-beam potential ranges from 3.2 to 4.9 V and is largest for intermediate power levels. Both distributions show peak offsets from the thruster centerline. This phenomenon was quite repeatable and evidently represents a true deviation from axis symmetry in the beam.

The 8,000-hr test identified electron-backstreaming as one of the key potential-failure modes for the engine. Electron-backstreaming refers to the phenomenon in which the space potential in the centers of the accelerator-grid apertures is insufficiently negative to prevent electrons in the beam plasma from streaming backwards into the engine. This phenomenon can result in a substantial performance loss for the engine, as well as the potential to damage the thruster by over heating. The accelerator-grid voltage at which electron-backstreaming occurs was measured periodically throughout the 8,000-hr test and is shown in Figure 16 for operation at the full-power point. The increase in the magnitude of the accelerator-grid voltage required to prevent electron-backstreaming observed over the 8,000-hr test results from the enlargement of the accelerator-grid apertures due to sputtering by charge-exchange ions. Post-test measurements of the accelerator-grid apertures as a function of the radial




Figure 14. Beam Current Density Distribution Measured at BOL for Six Throttle Levels (EMT2)




Figure 15. Beam Potential Measurements at BOL for Six Throttle Levels (EMT2)




Figure 16. Variation in Electron-Backstreaming Voltage at Full Power over the Course of the 8,000-hr Test (EMT2)

position on the grid are given in Figure 17. The pre-test accelerator-grid-aperture diameters are 1.14 mm. These data indicate a significant increase in the aperture diameter in the center region of the grid. The electron-backstreaming voltage margin at the end of the 8,000-hr test is given in Figure 18 over the NSTAR throttling range. While these voltage margins appear to be small, the accelerator grid could easily be operated at voltages more negative than those in the throttle table late in the engine life with essentially no adverse effects. The NSTAR PPU can provide accelerator-grid voltages as negative as -250 V.

The 8,000-hr test also provided a wealth of information regarding the details of other potential wear-out modes, including: erosion on the downstream side of the accelerator grid, erosion of the screen grid, erosion of the cathode keeper electrode, erosion of the cathode-orifice plate, and the thicknesses of sputter-deposited material films throughout the thruster [6]. Only one new potential failure mode was identified by this test. This failure mode results from material that is sputtered from the cathode-orifice plate and deposited on the upstream side of the cathode keeper electrode. If this sputter-deposited material becomes sufficiently thick, it could flake off and electrically short the cathode to the keeper. The thickest material deposits found anywhere in the thruster were on the upstream side of the cathode keeper. The separation distance between the cathode and the keeper is only 0.51 mm (0.020 inches), a distance that can easily be bridged by a flake of sputter-deposited material.

The data from the 8,000-hr test is being used in the development of models of the engine's principal wear-out failure modes. These models are being used in a probabilistic framework to quantitatively assess the engine failure-risk as a function of propellant throughput (or total

impulse) [19 to 24]. This modeling activity is a key part of the NSTAR program to validate the service life of the ion engine.

2.6.1.2 Extended Lifetime Test—After the successful completion of the 8,000-hr test, the last major test in the NSTAR project plan is to demonstrate 150% of the engine-design life using the DS1 flight spare engine (FT2) fabricated by HED. The engine-design life is most easily expressed in terms of the total amount of xenon propellant that the thruster can process. For the NSTAR project, the engine-design life is 83 kg of xenon, which corresponds to 8,000 hours of operation at full power. To demonstrate 150% of the engine life, therefore, requires a test in which 125 kg of xenon is processed by the engine. This test, designated NPT4 in the project plan, was originally designed to follow a representative mission-throttling profile; therefore, some of the test documentation still makes reference to a mission profile test (MPT). The test was later renamed the extended lifetime test (ELT) when it became clear that following a mission profile would not provide as much information about the engine-wearout modes at throttled conditions as a less complicated throttling plan. A secondary objective of this test is to demonstrate extended operation at throttled conditions since the previous project-level life tests had all been performed at the full-power point. It is believed that the full-power point is the most stressing to the engine; however, the ELT is designed to obtain the data necessary to support this assertion.

As of this writing (March, 2000), the ELT has operated FT2 for more than 9,400 hours. The first 500 hours of the test were performed at NSTAR throttle level 12 (TH12). From 500 hours through 5,000 hours, the engine was operated at full power. At 5,000 hours, the thruster was throttled to TH8, which is approximately 63% of full power. The test




Figure 17. Accelerator Grid Aperture Diameters Measured after the 8,000-hr Test Indicate Significant Enlargement from Their Original 1.14-mm Diameter Values




Figure 18. Electron-Backstreaming Voltage Margin at the End of the 8,000-hr Test (EMT2)

plan calls for the thruster to process 30 kg of xenon at this throttle level. The overall efficiency of FT2 over the first 5,400 hours of the ELT is given in Figure 19 for the entire engine-throttling range.

Electron-backstreaming data for FT2 versus run time is compared to that for EMT2 from the 8,000-hr test in Figure 20. The data for FT2 is systematically above that for EMT2 for operation at full power (TH15). This is believed to be a result of separation between the grids of the ion accelerator system being smaller at operating temperature in FT2 than in EMT2. A smaller grid separation requires a more negative accelerator grid to prevent electron backstreaming. The data at TH8 in Figure 20 exhibits a step-function change in the electron-backstreaming limit, even though the beam voltage is the same for both throttle levels. This step-function change is a result of the lower beam-current density for operation at TH8. The higher density of positive ions at full power increases the local space charge between the

grids more than at TH8 and, consequently, a more negative accelerator-grid voltage is required to prevent electron-backstreaming at full power.

The purveyance margin for the ion-accelerator system on FT2 is compared in Figure 21 to data taken on EMT2 over the 8,000-hr test. The purveyance limit defines the lower boundary of the engine-throttling envelope as shown in Figure 4 and is qualitatively defined as the beam voltage (for a fixed accelerator-grid voltage and beam current) at which direction impingement on the accelerator grid begins. The purveyance margin is the difference between the purveyance limit and the throttle table-set point for the beam voltage, which is 1100 V at both TH15 and TH8. The purveyance margin data in Figure 1 for FT2 agrees well with that for EMT2 at TH15. The purveyance margin increases at TH15 because of the lower beam current at this throttle level.




Figure 19. FT2 Efficiency Versus Power during ELT




Figure 20. Comparison of Electron-Backstreaming Limits for FT2 and EMT2

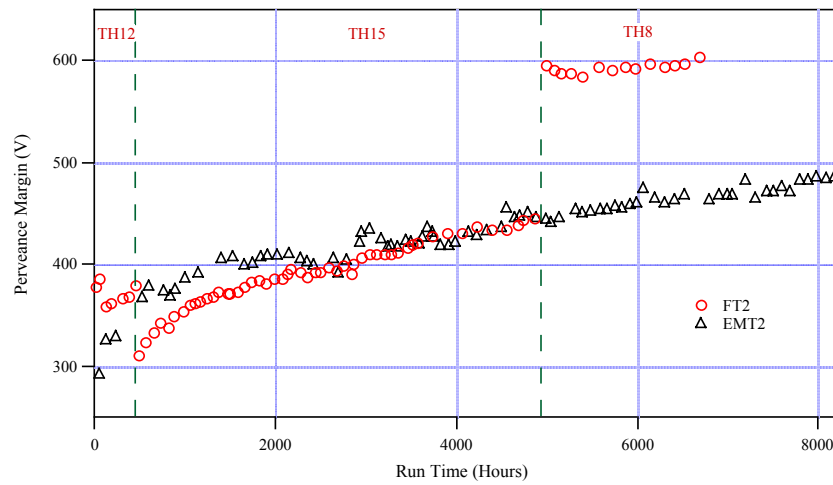


Figure 21. Comparison of the Long-Term Behavior of the Purveyance Margin for FT2 and EMT2

The discharge voltage is a key independent thruster-operating parameter that is used as an indicator of the health of the cathode and strongly affects key thruster-wearout modes. The long-term behavior of the discharge voltage for FT2 is compared to EMT2 in Figure 22. These data indicate excellent agreement for operation at TH15. This good agreement disappeared, as expected, when FT2 was throttled to TH8. Operation at throttled conditions typically results in higher-discharge voltages.

These data and the data given in [25] indicate that the operating behavior of the flight spare ion engine is very similar to that of the engineering-model thruster, EMT2. Since EMT2 exhibited excellent erosion characteristics (i.e., very little erosion), it is anticipated that the flight thrusters will exhibit similar life characteristics. The success of the ELT so far helps verify one of the key assumptions made during the design and fabrication of the flight thrusters: the engine structural and thermal designs could be improved without impacting the engine-service life as long as the critical components (which include the magnetic-field configuration, the cathode and neutralizer, and the ion-accelerator system) were unchanged from the engineering-model thrusters.

2.6.1.3 Characterization Tests—During the time that the 8,000-hr test was being conducted, many questions regarding other details of the thruster operation, behavior of the IPS components at the system level, and interface issues required a series of characterization tests (CTs). A total of 39 CTs were proposed. From this list, 18 of the highest priority tests were selected and executed. Table 9 lists the CTs which were actually performed.

One of the most significant CTs was CT31b, the end-to-end test of key elements of the IPS with the spacecraft power system. This test used an engineering model engine, a breadboard PPU, a breadboard DCIU, a solar-array simulator and the high voltage power conditioning unit (HVPCU) from the spacecraft’s power system. This test verified that there were no stability problems associated with handling the large power load represented by the IPS. This test is highly recommended for any future program planning the use of ion propulsion.

2.6.1.4 Engineering Development Tests—To address still further issues associated with the design of the flight engines, another series of tests was developed. This series, called engineering development tests (EDTs), was designed to address primarily structural and thermal issues associated with the engine design. The list of EDTs performed under the NSTAR project is given in Table 10.

2.6.2 Flight Test Program—The validation objectives of the IPS flight test on DS1 include demonstrating the functionality and performance of the system in an environment similar to what will be encountered by future users, the compatibility of the IPS with the spacecraft and science instruments, and autonomous navigation and control of the IPS with minimum ground-mission-operations support.

2.6.2.1 Operating Modes—The DCIU software is designed to perform the functions described briefly in this section. The system also has a number of fault-recovery functions that are defined in [26]. Only a few of those will be discussed here.




Figure 22. Comparison of the Long-Term Behavior of the Discharge Voltage for FT2 and EMT2

Table 9. NSTAR Characterization Tests

Thruster	Test	Purpose	Description	Duration (hrs)	Location
EMT1b	CT1	plasma screen grounding		1	GRC 5
EMT2	CT19'-1	pre-transport sensitivity	abbreviated TP: only 2 op points	12.2	GRC 5
	CT18	AC frequency components		n/a	JPL148
	CT13	magnetic map	various distances from thruster	n/a	JPL233
	CT19'-2	pre-LDT sensitivity		~5	JPL148
EMT3	CT5	low flow start–3 sccm		1	GRC 5
	CT6	single plena		6	GRC 5
	CT14	empirical thermal measmts	part of EDT2b	9	GRC 5
	CT22b	measure PPU in power quality	BBPPU during recycle	~8	JPL148
	CT27b	PPU input impedance		2	GRC 5
EMT4	CT31b	system end-to-end power stability	includes HVPCU	~25	JPL149
	CT36b	SAS IF verification		~16	JPL149
	CT36c	diode mode trial		1	JPL149
SPOT	CT31b	system end-to-end power stability	includes HVPCU	n/a	JPL149
n/a	CT33	DCIU-XEM1 c/o		n/a	JPL233
SPOT	CT22a	same as CT22b	BBPPU during recycle	n/a	JPL148
SPOT	CT24	PPU start circuit	effects on DS1	n/a	GRC 5
SPOT	CT27a	PPU input impedance		n/a	GRC 5

Cathode Conditioning—After launch, the cathodes are heated for several hours to help drive off oxidizing impurities from the inserts. This sequence is initiated by a single command and controlled by the DCIU.

Thruster Ignition—This operating mode begins with pressurizing the plenum tanks to the proper values, starting propellant flow to the engine, and preheating the cathodes prior to ignition of the neutralizer discharge. After 210 seconds of heating, the neutralizer high-voltage-pulse ignitor is started. After neutralizer-keeper current is detected, the heater and ignitors are turned off and the discharge is ignited. When both discharges have successfully lit, the high voltage is turned on at the minimum power level and the engine is throttled to the final setpoint. The accelerator-grid voltage is set to –250 V for two hours after ignition, then is increased to the correct throttle-point value.

Steady State Operation—The DCIU is capable of operating the thruster at any one of 16 discrete throttle levels from a throttling table stored in memory. This table contains the setpoints for the PPU power supplies and the XFS pressures and can be modified by ground command. The NSTAR 16-level-throttle table showing the entire range of operation is listed in Table 11. The DCIU commands the PPU power supplies to deliver these values and controls the XFS valves to maintain the desired pressures in steady-state operation. The beam-current setpoint is maintained by closed-loop control of the discharge current.

Throttling—When a new throttle level is commanded, the DCIU ramps the XFS pressures and the PPU outputs to the

new values. If the power level is being increased, the flows are raised before the engine power is changed. To throttle down, the electrical parameters are changed first, then the flow rates.

Thruster Power Down—In this operating mode the power supplies are turned off and all XFS valves are closed.

Continuous Recycling Fault Mode—The DCIU monitors the number of recycle events initiated by the PPU under high-voltage fault conditions. If 25 or more are recorded in a 90-second time period, the engine is shut off and a fault flag is set.

Grid Clear Fault Recovery—In the event of a physical short between the grids that cannot be cleared by recycling or mechanical methods, the DCIU can be commanded to execute a grid-clear operation. In this operating mode, internal relays in the PPU are closed to apply the discharge supply to the ion optics. The supply is then turned on at a pre-determined current level for a specified period of time in an attempt to resistively heat and to vaporize the short.

These DCIU functions can be called with ground commands. In addition, the spacecraft can generate commands to the IPS to perform certain operations. The IPS is throttled autonomously by the spacecraft to track the solar-array output. DS1 also includes an autonomous system (AutoNav) to navigate the spacecraft to the next encounter target. This system contains an optimized trajectory that was computed on the ground and a catalog of ephemerides for a number of stars, asteroids, planets, and DS1 target bodies.

Table 10. NSTAR Engineering Development Tests

Test	Purpose	Description	Thruster	Duration (hrs)	Location
EDT1a	initial vibe		EMT1b	n/a	NTS
EDT1b	follow-up vibe	with 3rd mounting pt	EMT1c	n/a	NTS
EDT1c	TGA vibe @ .2 g2/Hz	practice for FT#1	EMT1d	n/a	JPL144
EDT2a	cold start, etc.	downstream open	EMT3a	29	GRC 5
EDT2b	2nd phase thermal	+ downstream cover	EMT3a	65	GRC 5
EDT2c	3rd phase thermal	+ gimbal sim plate	EMT3b	334	GRC 5
EDT2d	4th thermal	+ DS1 thermal shield	EMT4	20	GRC 5
EDT2e	final thermal	same as 2d	PFT	41	GRC 5
EDT5	thrust stand performance	w/ modified ExB	EMT3	28	GRC 5
EDT6	500 hr cathode erosion		EMT3	500	GRC
EDT7a	Internal B field		EMT3	n/a	GRC
EDT9	mesh separation		EMT4	8	GRC
EDT12	screen grid saturation		EMT3	4	GRC
EDT16a	shorted discharge keeper		EMT3	3	GRC
EDT20a	plume tests		EMT3	12	GRC

Table 11. Flight Throttle Table of Parameters Controlled by the DCIU

NSTAR Throttle Level	Mission Throttle Level	Beam Supply Voltage (V)	Beam Supply Current (A)	Accelerator Grid Voltage (V)	Neutralizer Keeper Current (A)	Main Plenum Pressure (psia)	Cathode Plenum Pressure (psia)
15	111	1100	1.76	-180	1.5	87.55	50.21
14	104	1100	1.67	-180	1.5	84.72	47.50
13	97	1100	1.58	-180	1.5	81.85	45.18
12	90	1100	1.49	-180	1.5	79.29	43.80
11	83	1100	1.40	-180	1.5	76.06	42.38
10	76	1100	1.30	-180	1.5	72.90	41.03
9	69	1100	1.20	-180	1.5	69.80	40.26
8	62	1100	1.10	-180	1.5	65.75	40.26
7	55	1100	1.00	-150	2.0	61.70	40.26
6	48	1100	0.91	-150	2.0	57.31	40.26
5	41	1100	0.81	-150	2.0	52.86	40.26
4	34	1100	0.71	-150	2.0	48.08	40.26
3	27	1100	0.61	-150	2.0	43.18	40.26
2	20	1100	0.52	-150	2.0	39.22	40.26
1	13	850	0.53	-150	2.0	39.41	40.26
0	6	650	0.51	-150	2.0	40.01	40.26

Periodically (one-to-three times per week) during a burn, the system automatically turns the spacecraft to optically observe the positions of a number of these bodies against the stellar background and calculates the spacecraft position. The heliocentric orbit is then determined and the trajectory propagated to the next target. Required course changes are generated by the maneuver design element and accomplished by varying the IPS-thrust direction and duration. When enabled, this technology dramatically reduces the need for mission operations support, as described below.

2.6.2.2 *The NSTAR Throttle Table*—The NSTAR 16-point throttle table contains the IPS setpoints required to operate the system over a chosen throttling range. A corresponding mission-throttle table containing the flow rates, thrust, and PPU input- and output-power levels is maintained in spacecraft memory to enable the mission-trajectory calculations performed by the Nav Manager. The complete NSTAR mission table is shown in Table 4. The development of these throttle tables is described in this section.

Power throttling is accomplished by varying the beam voltage and current. The engine-throttling envelope with

lines of constant-beam power is shown in. The boundaries of this envelope represent the maximum beam voltage and current capabilities, the minimum-beam current (which is determined primarily by the minimum-discharge current) and the beam-voltage-purveyance limit. The NSTAR throttle table was designed to maximize the specific impulse; therefore, the power is varied with the beam’s current throttling over most of the range. The lowest-power levels are achieved by operating at the minimum beam current and throttling the beam voltage.

The discharge-chamber-flow rate was selected to give the propellant utilization shown in Figure 23. The propellant efficiency of 0.9 was selected at high power levels as a compromise between maximizing total engine efficiency and minimizing double ion production, which can drive internal-erosion rates. A propellant efficiency of 0.90 to 0.91 is maintained over most of the range. At the lowest powers, the double-to-single ion-current ratio is low; therefore, the propellant efficiency was chosen to give a discharge loss that yielded the correct total power at that point.

The thrust in the mission-throttle table is calculated from the engine’s electrical setpoints,

$$T = \alpha F_t J_b \left(V_s - V_g \right)^{1/2} \left(\frac{2M}{e} \right)^{1/2} \quad (1)$$

where J_b is the beam current, V_s is the beam power-supply voltage, V_g is the coupling voltage between neutralizer common and the facility ground or ambient-space plasma, M is the mass of a xenon ion, and e is the charge of an electron. The factors α and F_t correct for the doubly-charged ion content of the beam and thrust loss due to non-axial ion velocities [5]. A constant value of 0.98 for F_t based

on earlier 30-cm thruster ground tests and a value of α based on a curve fit to centerline double ion-current measurements as a function of propellant utilization efficiency in a 30-cm, ring-cusp inert-gas thruster [27] were used. Earlier direct measurements of thrust from the LDT agreed well with the calculated value [5,6]. More recent measurements with the flight thrusters were somewhat lower than the calculated values for intermediate throttle levels. The difference between the measured thrust and the table values is shown in Figure 24.

The power required for a given thrust level increases over the engine lifetime due to wear [5,6]; therefore, two tables representing beginning-of-life (BOL) and end-of-life (EOL) were developed. These have the same engine setpoints shown in Table 11 but different engine-power levels. The BOL table was developed primarily through testing with engineering-model thrusters and updated with data from pre-flight measurements with FT1. The EOL table was based largely on measurements from the 8200-hour test of EMT2. The power at the lowest throttle levels was extrapolated from performance curves obtained after about 6500 hours of operation. The extrapolations were based on sensitivity data, which were used to correct for slight differences in some of the controlled parameters. The difference between BOL- and EOL-engine power is plotted in Figure 25. Additional measurements taken at some of these throttle levels after about 6900 hours of operation in the LDT are also shown. They suggest that the EOL power at some of the lower throttle levels is overestimated in the throttling table. BOL data obtained with the two flight thrusters demonstrates that their initial performance agrees well with the table values.

The PPU input power corresponding to a given engine power is determined by the PPU efficiency. The flight-PPU efficiency of was characterized as a function of input-bus

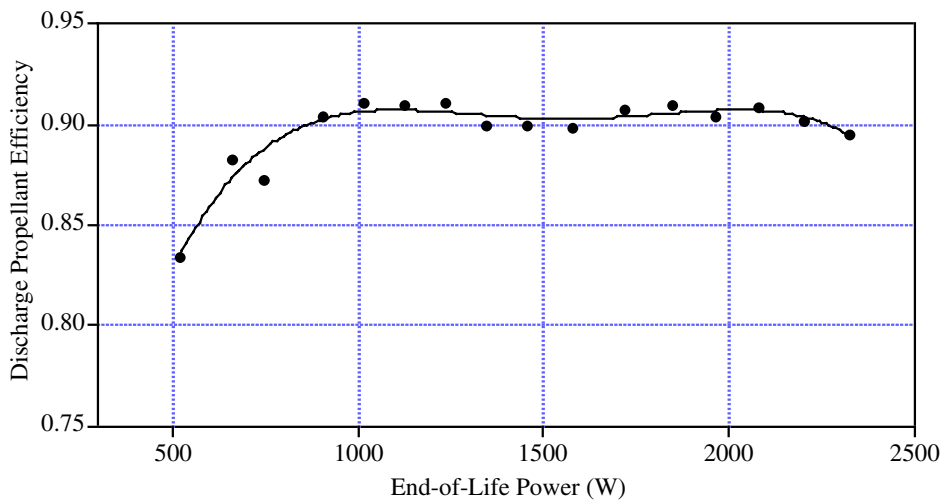


Figure 23. NSTAR Ion-Thruster Discharge-Propellant Utilization Efficiency




Figure 24. The Difference Between Measured and Calculated Thrust over the NSTAR Throttle Range




Figure 25. Difference Between a Given Power Level and the Beginning-of-Life Power

voltage and temperature in several ground tests, as shown in Figure 26. The lowest measured values over this range of parameters were used to define the lowermost line in the figure. This conservative estimate of PPU efficiency was used to generate the PPU input powers in the throttle table.

In order to make finer steps in power throttling to more closely track the solar-array peak power, a 112-point throttle table was also developed for use in flight. Power throttling between the 16 NSTAR throttle points is accomplished by varying the beam voltage to give steps that are approximately 20 W apart. A 16-point subset of this table is loaded into the DCIU to provide fine throttle control over a restricted power range for a given mission phase.

2.6.2.3 Post-Launch IPS Operation and Validation Activities—Operation of the ion propulsion system during the DS1 primary mission can be organized into several phases, which are summarized in this section.

Decontamination—The first IPS in-space activity was a bakeout of the downstream portion of the propellant-feed system that occurred six days after launch. Prior to this, the thruster axis was oriented 90° away from the Sun and the thruster front-mask temperature was -45° C. The spacecraft was turned so that the angle between the axis and the Sun was 30° to warm the thruster and feed system. Over a 29-hour period the thruster temperature exceeded 110° C and the XFS lines reached more than 45° C. This was done to help remove any residual contaminants in the portions of the feed system that had been exposed to air prior to launch. The cathode-conditioning sequence was then executed to bakeout the cathode inserts. Finally, 16 days after launch, the discharges were operated for four hours at high power levels to further bakeout the engine prior to application of high voltage.




Figure 26. In-Flight Measurements of PPU Efficiency Compared to Ground Test Data

Initial Start and Grid Short—The following day the first engine ignition occurred. Both cathodes lit properly and the engine ran nominally at the minimum-power point for 4.5 minutes before continuous recycling caused a thruster shutdown. A short between the grids was suspected, but at this point a failure of one of the high-voltage supplies could not be ruled out. Fourteen additional start attempts were made under various engine-thermal conditions (created by spacecraft turns toward or away from the Sun); all ended in continuous recycling when the high voltage was applied.

Troubleshooting—Taking advantage of the flexibility of the thrusting start date, a detailed investigation of the problem was undertaken. Several options were identified, including: attempting a grid-clear command, thermally cycling the engine to force a mechanical separation of the grids that might dislodge a particle, running additional recycles, and developing additional diagnostics to help identify the fault.

The NSTAR PPU is designed to deliver 4 A into a grid short to clear those that are not cleared by recycles. However, this system was designed primarily to clear thin molybdenum flakes generated by spalling of sputter-deposited films inside the discharge chamber after many thousands of hours of operation. Grid shorting this early in a mission was more likely due to particulates from the launch-vehicle payload fairing or generated during the payload preparation, which could be much larger than films from the discharge chamber. The risk of permanently welding a large particulate between the grids with the standard-grid clear circuit was not known, so an experimental and theoretical effort to characterize the grid-clear process was undertaken prior to using it under these circumstances. The results of this investigation are reported in [28].

Thermal and structural models of the ion optics were also coupled during this period to determine the mechanical effect of thermally cycling the grids. This modeling showed

that significant transient changes in the grid spacing can be achieved by turning the spacecraft to heat or cool the grids. This technique was used to clear grid shorts on the SERT II flight experiment [29] and appeared to have a very minimal risk. During the two-week problem-investigation period, the spacecraft was turned several times; this thermally cycled the grids over greater than a 100° C range.

The IPS is designed with hardware interlocks that prevent operation of the high-voltage supplies before the discharges are ignited; therefore, it was not possible to command these supplies to turn on separately to test them. The DCIU software was modified to provide brief bursts of high-speed data for various PPU electrical parameters during recycles to help diagnose which supplies were affected. Finally, a test involving operation of the discharge supply only, with no propellant flows (which is allowed by the system), was developed. If the grids are shorted, the accelerator-grid-voltage telemetry will change when the discharge-open-circuit voltage is applied; otherwise it remains close to zero. This is a clear discriminator between open circuits and shorts on the ion optics.

Recovery Start—Thirty-one days after launch, the discharge-only test was executed; the results suggested that the grids were not shorted. Another start attempt was then made, primarily with the intent to gather high-speed engine data during continuous recycling to help diagnose the fault. Fortunately, the engine started properly this time and has continued to run flawlessly since this point. Apparently, the thermal cycling successfully cleared debris lodged between the grids.

The origin of the surmised debris cannot be conclusively identified, but the event itself points to the importance of contamination control on the engine pre-launch. Much care was taken to launch with a dust- and debris-free thruster, in both design and handling. An especially concentrated effort

was devoted to the nearby solid rocket motor (SRM) dome surfaces and the spacecraft-separation system, with design changes actually implemented once the contamination analysis identified possible sources for debris in the original plan.

In the future, all reasonable origins for debris should be studied and identified and appropriate protection should be implemented.

First Performance Test—Over the next 335 hours, the engine was operated at power levels ranging from 0.48 to 1.94 kW to characterize the BOL performance. This burn was used to contribute to the required spacecraft ΔV , but was not controlled by AutoNav. The throttle levels were dictated primarily by the validation objectives. This test was designated IPS Acceptance Test 1 (IAT1).

Deterministic Thrusting—IAT1 was followed by 95 hours of thrusting at power levels ranging from 1.7 to 1.86 kW. These initial operations also contributed to the required total impulse, but were executed with ground commands. These were followed by a coast period of 74 days and seven navigational burns (NBURNS) totaling 912 hours of operation. These maneuvers were executed autonomously by AutoNav and used automatic-peak-power tracking to determine the maximum achievable throttle level. The first of these, NBURN 0, did not use the optical navigation for spacecraft-position determination; however, all subsequent NBURNS have exercised the full AutoNav capability. This part of the mission is on an outbound portion of the trajectory, so the available array power decreased continuously. NBURN 0 was run with engine power levels ranging from 1.73 to 1.62 kW, while the following six NBURNS were performed with power levels of 1.18 to 0.71 kW. These burns completed the deterministic thrusting required for the encounter with asteroid Braille.

Second Performance Test—After another coast period of 21 days, a second throttling test was performed. This brief test, designated IAT2, was restricted to power levels ranging from 0.49 to 0.98 kW by total solar-array power.

2.6.2.4 In-Flight System Performance—One of the primary objectives of the flight-validation activity is to verify that the system performs in space as it does on the ground. The parameters of interest to future mission planners are those in the mission-throttle table: thrust and mass flow rate as a function of PPU input power. In this section, the system power, thrust, and mass-flow-rate behavior will be evaluated in terms of the throttle table.

PPU Power Input Requirements—The PPU input power is determined by the PPU output power (engine-power requirement) and the PPU efficiency. The difference between the in-flight engine, input power and the BOL

throttle-table power is shown in Figure 25. These power values are based on the individual power-supply current and voltage-telemetry readings. The total engine power consumed during the IAT1 throttle test and initial operations differed from the table values by only about 2 W on average, although the uncertainties are much larger than this, as shown by representative error bars on the figure. The engine-power requirement increased by 12 to 15 W with time, however, as the data from NBURNS 1 to 3 and IAT2 show. This is a normal consequence of engine aging [5,6], and the total power at this point in the mission is still less than the EOL power used in the throttle table, which is represented by the solid line in Figure 25. This increased power demand is due primarily to increased discharge-power losses, as discussed in the next section.

In-flight measurements of the PPU efficiency suggest that it is higher than that measured in ground tests, as shown in Figure 26. These values are based on the total engine power and PPU high-voltage bus current and voltage telemetry with an additional 15 W assumed for the low voltage-bus-input power. There is no telemetry for the low voltage bus; however, ground testing showed a 15 W loss for all conditions. The efficiency is sensitive to the line voltage and the temperature, as the ground data show. The in-flight measurements were taken with line voltages of 95 ± 5 V and baseplate temperatures ranging from 0 to 37° C, so they should be compared with the solid line in the center of the preflight data and the highest dashed line. The range of uncertainty in these measurements encompasses the ground test data; however, the in-space measurements appear to be higher systematically by about one percentage point. This apparent performance gain is not understood and may be due to a systematic error in the ground or flight measurements.

If the PPU efficiency is actually higher than anticipated, it more than offsets the increased output-power requirements observed so far in the primary mission. Figure 27 displays the difference between the observed PPU-input power and the BOL-input power from the throttle table. The input power required early in the mission was approximately 20 W lower than expected because of the higher PPU efficiency. The data from the NBURNS and IAT2 show that the input power is just now approaching the BOL throttle-table value.

IPS Thrust—The acceleration of the spacecraft is measured most accurately from changes in the Doppler shift of the telecommunications signals. With models of the spacecraft mass as a function of time, the Doppler residual data can be used to measure the thrust of the IPS with an uncertainty of less than 0.5 mN. Preliminary thrust measurements have been obtained so far from IAT1, the initial operations, and NBURN 0. The flight-beam voltage and current values, which determine to a large extent what the thrust is, are

slightly different from the setpoints in the table. The flight-thrust measurements are, therefore, compared to the thrust calculated from the actual electrical parameters rather than the table values. The difference in the measured and calculated thrust is shown in Figure 28, with the curve fits to similar data obtained with a thrust balance in ground tests. The ground and flight data agree well with the calculated values at low power levels, but are lower at intermediate powers. The flight data suggest that the difference in true thrust and calculated thrust grows linearly with power, peaking at 1.6 mN lower than expected at mission level 83 (1.82 kW engine power). The error bars shown in this figure are based on the uncertainty in the measured thrust and do not include errors in the calculated thrust.

This discrepancy may also be due to a systematic error in the flight telemetry, although the agreement with ground data argues against that conclusion. As Equation (1) shows, the true thrust might be lower than calculated because of a higher double-ion content, greater beam divergence than observed in the previous 30-cm thruster tests, or differences in the coupling voltage in space compared to ground tests. Additional measurements and analysis will be required to resolve this issue.

Although the actual thrust appears to be slightly lower than expected, at the beginning of the mission the overall system performance was still very close to the BOL throttle-table




Figure 27. Difference Between a Given Input Power to the Flight PPU on DS1 and the Corresponding Throttle Table BOL Value




Figure 28. Difference Between Measured and Calculated Thrust in Flight Compared to Ground Measurements

level in terms of thrust for a given PPU input power. Figure 29 shows that at the beginning of the mission the higher PPU efficiency largely compensated for the lower thrust. In this comparison, the thrust is within 0.5 mN of the table values. The gap between the two widens as the engine wears and the total engine-power requirement for a given throttle level grows. The PPU input power required for the thrust levels measured during NBURN 0 has exceeded the EOL throttle-table power for an equivalent thrust.

The thrust-vector behavior in-flight is similar to that observed in ground tests. The engine is mounted on a two-axis gimbal with range of $\pm 5^\circ$. When the IPS is not operating, a hydrazine attitude-control system is used for 3-axis stabilization. After ignition of the ion thruster, control in two axes is transferred to the IPS gimbal system. Potentiometers on each axis of the gimbal provide a measure of the thrust-vector stability during IPS operation. There is a brief transient after transfer of control; however, after that the mean value of the gimbal angle appears to be stable over long periods of time. The thrust vector of the flight engine relative to the thruster axis was measured using a thrust-vector probe [16] prior to integration and alignment on the spacecraft. The gimbal-angle data in Figure 30 show that this alignment was excellent. They also demonstrate that the thrust vector changes slightly with throttle level, as shown in previous ground tests [16].

Propellant Flow Rates—The performance of the xenon feed system is discussed in detail in [9]. In general, the performance has been excellent, although the flow rates are slightly higher than the throttle-table values. The mean value of the main flow is 0.05 to 0.14 sccm (about 0.4 to 1.0 percent) high, while that of the two cathode flows is 0.03 sccm (about 1.0 percent) high. This is in part

intentional. As Figure 31 shows, the XFS bang-bang regulators result in a sawtooth-pressure profile. The control system is designed so that the minimum pressure in this sawtooth yields the throttle table flow-rate values. In addition to this deliberate conservatism, there is a slight bias in both regulators because one of each of the three pressure transducers on the two plena had a slight offset after launch.

Overall System Performance—The propulsion system performance can be summarized in terms of specific impulse and efficiency. At the beginning of the mission, the Isp was about 60 seconds lower than expected and the engine efficiency was 2 to 2.5 percentage points lower than the throttle-table values. The measured performance was still excellent, with a measured efficiency of 0.42 to 0.60 at Isp's ranging from 1960 to 3125 seconds over an engine-throttling range of 478 to 1935 W. Measured mission-planning performance parameters are listed in Table 12.

2.6.2.5 Engine Behavior In-Flight—The engine behavior in space has been very similar to that observed in ground testing. The detailed operating characteristics of the engine are discussed in this section.

Engine Ignitions—A total of 32 successful engine ignitions have occurred in the first 1791 hours of the primary mission with only one failure to achieve beam extraction (due to the initial grid short discussed above). The data from the first 25 ignitions are reviewed here. The nominal heater-current value is 8.5 A; the actual cathode and neutralizer-heater currents in flight have been constant at 8.444 A and 8.375 A, respectively. The time history of the heater voltages, which are an indicator of heater health, are plotted in Figure 32. The uncertainty in these measurements is about 12%. The first 15 ignitions include the first successful engine start




Figure 29. Thrust Measured in Flight as a Function of PPU Input Power Compared to the Throttle Table Values




Figure 30. In-Flight Gimbal Positions as a Function of Mission Throttle Level




Figure 31. Example of Flow Rate Throttling

and 14 start attempts after continuous recycling shut the thruster off. The peak-heater voltage is a function of the heater impedance, current, and temperature. The data show that the heater voltage increases in any rapid sequence of ignitions because the conductor is hotter at the beginning of each consecutive start. The subsequent data show that the heater voltage is also higher when the initial thruster temperature (indicated by the front-mask temperature in the plot) is higher. The scatter in the peak voltages under similar temperature conditions is low and very similar to that observed in ground tests.

The time required for the cathodes to ignite after the 210 seconds heat phase and application of the high voltage-ignitor pulses is plotted in Figure 33. The neutralizer ignition delays show trends that also follow initial

temperature, with 20 to 80 second delays observed for the lowest temperatures. Delays of up to 86 seconds were also observed during ground-thermal tests at the lowest temperatures [13] and are not considered to be a concern. In all cases, the discharge cathode has ignited 5 to 6 seconds after successful neutralizer ignition, which reflects delays in the start sequence. Its ignition reliability may be higher because it has a slightly higher heater current and because it automatically goes through a longer heat phase when the neutralizer ignition is delayed.

Throttling Characteristics—The throttling sequences were in all cases executed properly by the DCIU after receiving ground commands. An example of the throttling sequence is shown in Figure 31 and Figure 34. The IPS Manager software onboard the spacecraft is also designed to

Table 12. Flight Engine Performance Measured in Space

NSTAR Throttle Level	Mission Throttle Level	PPU Input Power (kW)	Engine Input Power (kW)	Measured Thrust (mN)	Main Flow Rate (sccm)	Cathode Flow Rate (sccm)	Neutralizer Flow Rate (sccm)	Specific Impulse (s)	Total Efficiency
12	85	1.99	1.86	75.34	19.99	2.91	2.82	3035	0.602
11	83	1.94	1.82	72.55	18.63	2.75	2.67	3125	0.610
11	83	1.96	1.83	72.63	18.62	2.75	2.67	3131	0.609
10	77	1.84	1.72	69.54	18.59	2.75	2.67	3000	0.594
10	76	1.82	1.70	67.21	17.31	2.58	2.51	3109	0.602
10	75	1.79	1.68	66.81	17.33	2.58	2.51	3087	0.601
10	74	1.77	1.66	66.11	17.33	2.59	2.51	3054	0.595
10	73	1.75	1.65	65.64	17.31	2.59	2.51	3035	0.594
10	72	1.73	1.63	65.15	17.31	2.59	2.51	3012	0.592
9	69	1.67	1.57	62.27	16.08	2.50	2.43	3070	0.597
6	48	1.29	1.22	47.43	11.42	2.50	2.42	3006	0.573
6	48	1.29	1.22	47.39	11.44	2.49	2.42	3004	0.571
3	27	0.89	0.84	31.70	6.93	2.50	2.43	2770	0.511
0	6	0.50	0.48	20.77	6.05	2.50	2.43	1961	0.418




Figure 32. Time History of Peak Cathode and Neutralizer Heater Voltages in Flight

autonomously throttle the engine to track the peak power available from the array. The engine is initially throttled up until auxiliary battery power drain is observed and then decreased until no battery power is required. Anytime battery operation is detected as available array-power drops or the spacecraft’s power needs increase, the IPS is commanded to throttle down to accommodate the reduced power. This function was successfully demonstrated in all of the NBURNS, which were accomplished with no ground control required over the detailed engine operations.

Steady-State Setpoint Accuracy—As mentioned above, the flight-flow rates are slightly higher than the throttle-table setpoints. In addition, the beam current is 4 to 13 mA high over a range of 0.51 to 1.49 A. The beam current is controlled in flight to within +2 mA by varying the discharge current in a closed loop. This variation is driven

primarily by the flow-rate sawtooth, as shown in Figure 34. The neutralizer-keeper current is 17 mA low at the 2 A setpoint and 10 mA low at 1.5 A. The accelerator-grid voltage is 2 V higher than the setpoint at all operating points. The beam voltage is on average about 3 V lower than the setpoints. The offsets in beam-power supply settings result in slightly higher beam-power levels than the throttling tables assume. This is largely offset by lower neutralizer-power levels, as explained below. All of these parameters are well within the specified flight-system tolerances.

Discharge Performance—As indicated in the previous section, the difference between the total engine power and the throttle-table values is dominated by the discharge-power difference. Discharge performance is summarized in terms of the ion-energy cost (eV/ion) plotted in Figure 35.




Figure 33. Time History of Cathode and Neutralizer Ignition Delays in Flight




Figure 34. Example of In-Flight Throttle-up and Throttle-down Sequences




Figure 35. Discharge Loss Measured in Flight Compared to the Throttle Table Values

The standard error of these measurements is 1.5 percent. This plot shows the beginning and end-of-life discharge loss as a function of mission-throttle level. The data from early in the DS1 mission are quite close to the throttle-table values except in the middle of the range (throttle levels 40 to 60), where the flight data are higher. This appeared to be true of the ground measurements as well, suggesting that the BOL throttle-table discharge loss and total power are low by about 10 W in this range. The data from NBURNS 1 to 3 and IAT2 indicate that the discharge losses are increasing with time as a consequence of engine wear [5,6]. The lowest throttle levels are particularly sensitive to engine wear and show the largest increases in flight, up to 40 W. However, all of the data are still bounded by the throttle-table BOL and EOL values.

The discharge voltage and current are compared with the throttle-table values in Figure 36 and Figure 37. The voltages measured in flight are typically within 2% of the throttle-table voltages. The ground-test data are also plotted in this figure and tend to be slightly higher, although some of these measurements have not been corrected for voltage drops in the ground-facility power cables. There is very little drift in the discharge voltage over the course of the flight, which is consistent with long duration ground-test data [5,6]. The discharge current is also close to the BOL table values initially, with the exception of measurements at mission level 48. This is in the range where the table values appear to underestimate true BOL behavior. Unlike the voltage, the discharge current increases with time and drives the discharge power toward the EOL values.




Figure 36. Discharge Voltage Measured in Flight Compared to the Throttle Table Values and Ground Test Measurements




Figure 37. Discharge Current Measured in Flight Compared to the Throttle Table Values

Data on the sensitivity of discharge losses, voltage and current to small variations in flow rates, and beam current from the ongoing extended life test were used to examine the effect of setpoint errors on the flight-discharge parameters. The effects compete and result in negligible changes in these parameters due to the small flow- and beam-current errors.

Ion Optics Performance—The ion optics appear to be performing very well so far in flight. The accelerator-grid-impingement current as a function of beam current is compared to ground-test data in Figure 38. The standard error of these measurements is about 0.03 mA. The data obtained in the ground-test facilities are higher because they include a contribution from charge-exchange reactions with residual tank gas. The flight impingement-current levels in space are about 0.4 mA lower at 0.51 A and 1.7 mA lower at 1.5 A compared to pre-flight measurements in the JPL

endurance-test facility, which operates at pressure levels of $2\text{--}5 \times 10^{-4}$ Pa ($1.5\text{--}4 \times 10^{-6}$ Torr) over the full-throttle range. Accelerator grid erosion measurements obtained in long duration tests in this facility are, therefore, conservative. Data obtained in VF5 at NASA GRC, which has a residual-gas pressure about three times lower than that at JPL, show impingement currents that are about 0.4 mA greater than the space values. The ratio of impingement current to beam current is shown as a function of beam current in Figure 39. This parameter, which is used in some probabilistic models of accelerator-grid erosion [19,21,23,24], ranges from 0.17 percent at 0.51 A to 0.28 percent at 1.5 A with a standard deviation of 0.012 percent. A total of 88 high-voltage faults have occurred during 1791 hours of engine operation (excluding those that occurred as a result of the initial grid short). There has been no evidence of electron back-streaming. The discharge loss has consistently increased




Figure 38. Accelerator Grid Impingement Current Measured in Space Compared to Ground Test Measurements




Figure 39. In-Space Ratio of Accelerator Grid Impingement Current to Beam Current

slightly when the accelerator-grid voltage is raised from -250 V after ignition to the throttle setpoint, which is the nominal behavior. This transition is monitored for decreases in the discharge loss, which could signal the loss of electron backstreaming margin.

Neutralizer Performance—The neutralizer-power consumption has been 4 to 7 W lower than the BOL throttle-table values due to a lower neutralizer-keeper voltage, shown in Figure 40. This power savings roughly compensates for a higher beam-power demand due to the beam-current offset. The voltage dropped by about 0.5 V over several days before many of these data were taken in IAT1. The IAT1 data show that at that point in the mission, the keeper voltage was up to 2 V less than the pre-test values. This difference is not yet understood. The voltage has continued to decrease with time, as the data from the initial operations and the NBURNS show. This behavior has been observed in

ground tests [5,6] and is an indication of improving emitter-surface conditions.

There is no instrumentation on the DS1 spacecraft that allows the true neutralizer-coupling voltage to be easily determined. The voltage of neutralizer common with respect to the spacecraft ground is metered, and the behavior is shown in Figure 41. To properly compare this with the ground measurements of coupling voltage, also shown in this plot, the spacecraft potential with respect to the ambient plasma must be known. It may be possible to estimate this from the onboard plasma diagnostics; however, this analysis is not yet complete. It is interesting to note that the voltage variation with throttle level has the same slope as that of the coupling voltage in ground measurements and that the magnitude is decreasing with time, which also occurs in ground tests.

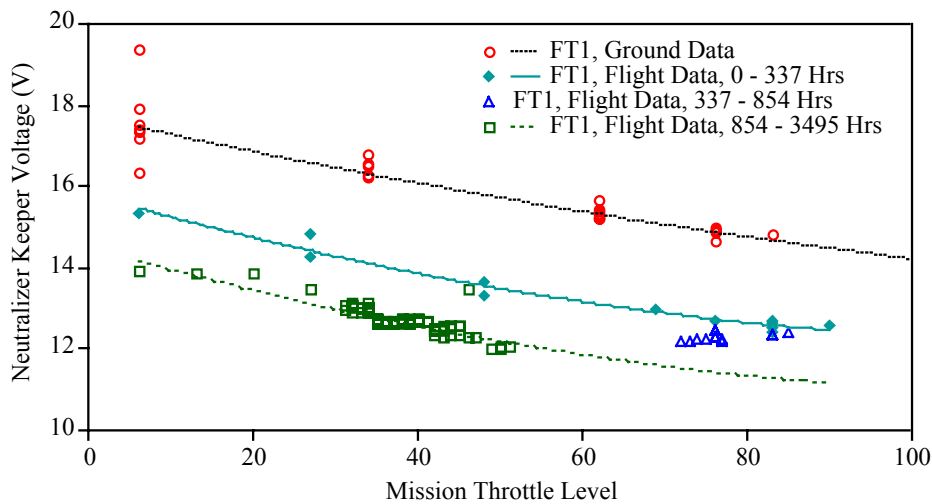


Figure 40. Neutralizer Keeper Voltage Measured in Space and in Ground Tests




Figure 41. Neutralizer Common Voltage Measured with-Respect-to Spacecraft Ground in Space and with-Respect-to Facility Ground in Ground Tests

2.6.2.6 Mission Operations—Although the total thruster-operating time so far has been orders-of-magnitude longer than that required by impulsive propulsion systems, the mission-operations demands from the IPS have been reasonably minimal (once account is taken for this flight being the first experience with low-thrust navigation and the consequent conservativeness for the sequencing and activity-review processes). Once confidence in the IPS operation was gained, the mission-operations process was streamlined as originally intended.

This was largely due to the successful implementation of a high degree of spacecraft autonomy. Autonomous navigation has significantly reduced the demands on the navigation- and trajectory-design teams. Spacecraft control of the IPS relieves the ground controllers considerably. In the initial phase of the mission, a number of propulsion engineers were involved in mission operations and validation. However, the final NBURNs have become sufficiently routine at this point that not much workforce is assigned to this area. The flight-data dissemination and analysis has also been largely automated. During Deep Space Network coverage, the spacecraft telemetry is displayed in real time on a Web site that can be accessed by the flight team. Data are also stored in the JPL ground-data system and automatic queries to this system generate files of IPS data periodically that are sent via FTP to all flight team members. A series of macros written in Igor Pro software are used to automatically load, analyze, and plot these data.

The success in reducing mission-operations requirements with automation is an extremely significant result because the fear of excessive operations costs has been a major barrier to the acceptance of ion propulsion for planetary missions. It now appears that the mission-operations costs for SEP-driven spacecraft are similar to those for conventional spacecraft or possibly less in cases where the use of ion propulsion results in shorter trip times.

3.0 TECHNOLOGY VALIDATION SUMMARY

The following key risks were retired by the NSTAR project, and the flight of the ion propulsion system on DS1:

- Adequate engine life—Prior to the NSTAR project, no ion engine intended for primary propulsion had ever been successfully operated for its full design life. The NSTAR project did this and is in the process of demonstrating 150% of the engine design life.
- Guidance, Navigation and Control of an SEP spacecraft—The low-thrust nature of SEP made this a risk area. The operation of the SEP system on DS1 demonstrated that GN&C is not more difficult with an SEP spacecraft, just different.
- Mission-operation costs—Requiring the propulsion system to operate continuously led some to project that

a standing army of propulsion and power engineers would be required to operate the spacecraft. However, the electrical nature of SEP lends itself well to autonomous operation, resulting in essentially no significant increase in mission operations cost for SEP vehicles.

- Spacecraft contamination by the SEP system—Slow erosion of the engine results in a non-propellant efflux from the thruster that could contaminate sensitive spacecraft surfaces. Data from DS1 indicates that this efflux travels essentially line-of-sight from the engine and poses no health risk to the spacecraft.
- SEP impacts on science instruments—The charge-exchange plasma generated by the operation of the SEP system is easily detected by onboard plasma instruments. DS1 showed that this low-energy plasma does not interfere with measurements of the much more energetic solar-wind plasma.
- SEP impacts on communication—The charge-exchange plasma generated by the operation of the SEP system could affect the transmission or reception of electromagnetic waves. However, no impact of the SEP system on communications with DS1 could be detected.
- Electromagnetic compatibility (EMC) of the SEP system with the spacecraft—The high-power nature of SEP and the use of strong permanent magnets in the ion engines could make it difficult for the SEP system to be electromagnetically compatible with the spacecraft. DS1 showed that while this issue requires careful engineering, it is an easily tractable problem.

4.0 FUTURE APPLICATIONS

Many missions have been identified by JPL's advanced mission planning activity as being either enabled or strongly enhanced by the use of solar-electric propulsion based on NSTAR or derivatives of the NSTAR ion-propulsion technology, including: Comet Nucleus Sample Return, Mercury Orbiter, Neptune Orbiter, Titan Explorer, Saturn Ring Observer, Europa Lander, and Venus Sample Return. In addition, it is anticipated that several Discovery Mission proposals will baseline the use of NSTAR-based ion propulsion systems to reduce the cost of going to scientifically interesting but propulsively difficult destinations.

To illustrate the benefits enabled by the use of an NSTAR-derivative SEP system for a Comet Nucleus Sample Return (CNSR) mission, the performance of this mission with SEP for the target-comet 46P/Wirtanen is compared to ESA's chemical-propulsion-based Rosetta mission to the same comet. The Rosetta spacecraft has an initial wet mass of 2,900 kg and must be launched on an Ariane 5. This spacecraft takes more than 9 years to reach the comet, arrives with a net spacecraft mass of 1300 kg, and is not

capable of returning a sample from the comet. The SEP-based CNSR spacecraft, on the other hand, has an initial-wet mass of 1830 kg and is launched on a Delta IV medium launch vehicle. The spacecraft takes only 2.6 years to reach the comet with a delivered mass of over 1300 kg and takes an additional 4.5 years to return a sample to Earth. Thus, the SEP-based CNSR spacecraft can travel to the comet and return to Earth in less time than it takes for the Rosetta spacecraft to fly to the comet!

Future deep-space missions will require multi-engine SEP systems, instead of the single-engine system used on DS1, with up to 4 engines operating at a time and processing up to 10 kW of power. In addition, these systems will require a significantly enhanced engine-throughput capability, operation at higher power levels per engine, and operation at higher specific impulses. The NSTAR service life assessment activity, which includes a combination of long-duration testing [5,6,16 to 18,25] and analyses [19 to 24] of the critical engine-wear-out-failure modes, indicates that the NSTAR engine can process a total propellant throughput of 130 kg with a low failure risk. Further analyses and extended testing of the DS1 flight-spare engine are planned to extend this throughput capability to larger values.

5.0 ACKNOWLEDGMENT

This publication was prepared at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

6.0 LIST OF REFERENCES

- [1] R. Kakuda, "NSTAR Flight Validation Report," MOR-602, JPL D-18219, internal document, November 1, 1999.
- [2] J. S. Sovey, et al., "Development of an Ion Thruster and Power Processor for New Millennium's Deep Space 1 Mission," AIAA-97-2778, presented at the 33rd AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Seattle, WA, June 1997.
- [3] J. S. Sovey, V. K. Rawlin, and M. J. Patterson, "A Synopsis of Ion Propulsion Development Projects in the US: SERT I to DS1," AIAA-99-2270, presented at the 35th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Los Angeles, CA, June 1999.
- [4] R. Kakuda et al., "NSTAR User Performance Needs Report," JPL D-12672, May 2, 1995.
- [5] J. E. Polk, et al., "The Effect of Engine Wear on Performance in the NSTAR 8,000-Hour Ion Engine Endurance Test," AIAA-97-3387, presented at the 33rd AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Seattle, WA, June 1997.
- [6] J. E. Polk, et al., "An Overview of the Results from an 8,200-Hour Wear Test of the NSTAR Ion Thruster," AIAA-99-2446, presented at the 35th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Los Angeles, CA, June 1999.
- [7] J. E. Polk, et al., "Validation of the NSTAR Ion Propulsion System on the Deep Space One Mission: Overview and Initial Results," AIAA-99-2274, presented at the 35th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Los Angeles, CA, June 1999.
- [8] M. G. Marcucci and J. E. Polk, "NSTAR Xenon Ion Thruster on DS1: Ground and Flight Tests," I-18, presented at the 8th International Conference on Ion Sources, Kyoto, Japan, Sept. 6 to 10, 1999.
- [9] G. B. Ganapathi and C. S. Engelbrecht, "Post Launch Performance Characterization of the Xenon Feed System on Deep Space One," AIAA-99-2273, presented at the 35th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Los Angeles, CA, June 1999.
- [10] J. A. Christensen, et al., "The NSTAR Ion Propulsion Subsystem for DS1," AIAA-99-2972, presented at the 35th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Los Angeles, CA, June 1999.
- [11] J. A. Hamley, et al., "The Design and Performance Characteristics of the NSTAR PPU and DCIU," AIAA-98-3938, presented at the 34th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Cleveland, OH, July 1998.
- [12] J. A. Christensen, et al., "Design and Fabrication of a Flight Model 2.3-kW Ion Thruster for the Deep Space 1 Mission," AIAA-98-3327, presented at the 34th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Cleveland, OH, July 1998.
- [13] V. K. Rawlin, "NSTAR Flight Thruster Qualification Testing," AIAA-98-3936, presented at the 34th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Cleveland, OH, July 1998.
- [14] J. J. Wang, et al., "Deep Space One Investigations of Ion Propulsion Plasma Interactions: Overview and Initial Results," AIAA-99-2971, presented at the 35th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Los Angeles, CA, June 1999.
- [15] V. Rawlin, et al., "Thermal Environmental Testing of NSTAR Engineering Model Ion Thrusters," IEPC-97-051, presented at the 25th International Electric Propulsion Conference, Cleveland, OH, August 1997.
- [16] J. E. Polk, et al., "Behavior of the Thrust Vector in the NSTAR Ion Thruster," AIAA-98-3940, presented at the 34th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Cleveland, OH, July 1998.
- [17] M. J. Patterson, et al., "2.3-kW Ion Thruster Wear Test," AIAA-95-2516, presented at the 31st AIAA/ASME/SAE/ASEE Joint Propulsion Conference, San Diego, CA, July 1995.
- [18] J. E. Polk, et al., "A 1,000-hr Wear Test of the NASA NSTAR Ion Thruster," AIAA-96-2717, presented at the 32nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Lake Buena Vista, FL, July 1996.

- [19] J. E. Polk, et al., “Probabilistic Analyses of Ion Engine Accelerator Grid Life,” IEPC-93-176, presented at the 23rd International Electric Propulsion Conference, Seattle, WA, September 1993.
- [20] J. E. Polk, et al., “Spatial and Temporal Distribution of Ion Engine Accelerator Grid Erosion,” AIAA-95-2924, presented at the 31st AIAA/ASME/SAE/ASEE Joint Propulsion Conference, San Diego, CA, July 1995.
- [21] J. R. Brophy, et al., “Ion Engine Service Life Validation by Analysis and Testing,” AIAA-96-2715, presented at the 32nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Lake Buena Vista, FL, July 1996.
- [22] J. E. Polk, et al., “In Situ, Time-Resolved Accelerator Grid Erosion Measurement in the NSTAR 8,000-hr Ion Engine Wear Test,” IEPC-97-047, presented at the 25th International Electric Propulsion Conference, Cleveland, OH, August 1997.
- [23] J. R. Anderson, et al., “Service Life Assessment for Ion Engines,” IEPC-97-049, presented at the 25th International Electric Propulsion Conference, Cleveland, OH, August 1997.
- [24] J. R. Brophy, et al., “The Ion Propulsion System on NASA’s Space Technology 4/Challengion Comet Rendezvous Mission,” AIAA-99-2856, presented at the 35th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Los Angeles, CA, June 1999.
- [25] J. R. Anderson, et al., “Results of an On-Going Long-Duration Ground Test of the DS1 Flight Spare Ion Engine,” AIAA-99-2857, presented at the 35th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Los Angeles, CA, June 1999.
- [26] J. A. Hamley, “NSTAR Thruster Element Technical Requirements Document,” Technical Report D13638, Jet Propulsion Laboratory, Internal Document, 1997.
- [27] M. J. Patterson, et al., “NASA 30-cm Ion Thruster Development Status,” AIAA-94-2849, presented at the 30th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Indianapolis, IN, July 1994.
- [28] K. D. Goodfellow, G. B. Ganapathi, and J. F. Stocky, “An Experimental and Theoretical Analysis of the Grid Clearing Capability of the NSTAR Ion Propulsion system,” AIAA-99-2859, presented at the 35th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Los Angeles, CA, June 1999.
- [29] W. R. Kerslake and L. R. Ignaczak, “SERT II 1980 Extended Flight Thruster Experiments,” IEPC-81-665, presented at the 18th International Electric Propulsion Conference, Las Vegas, 1981.
- [30] M. D. Rayman, et al., “Results from the Deep Space 1 Technology Validation Mission,” IAA-99-11.2.01, presented at the 50th International Astronautical Congress, Amsterdam, The Netherlands, October 1999.

Appendix A. List of Telemetry Channels and Names

The IPS and spacecraft-data channels used for IPS diagnostics, trending analysis, and NSTAR archive storage are given in Table A1

Table A1. IGOR Data Channels

Channel	Title of Parameter	Channel	Title of Parameter	Channel	Title of Parameter
SCET		V0128	Shutdown Mode	V0198	PPU Status Word #2
ERT		V0129	Code Checksum	V0199	# of Recycles
V0001	EHA DCIU XIPS Mode	V0130	XFS Operating Mode	V0200	XFS Status Word
V0002	EHA PPU Status Word 1	V0131	Software Version #	V0201	Valve Status Word
V0003	EHA XFS Status Word	V0132	PPU Data Packet ID	V0202	# SV3 Cycles
V0004	EHA Mgr. Talking?	V0133	Accel Current	V0203	# SV4 Cycles
V0005	EHA Manager DCIU state	V0134	Accel Voltage	V0204	Continuous Dump Offset
V0006	EHA Last Command sent	V0135	Beam Current	V0205	Continuous Dump Segment
V0008	EHA DCIU state	V0136	Beam Voltage	V0206	Continuous Dump #0
V0009	EHA XIPS Mode	V0137	Discharge Current	V0207	Continuous Dump #1
V0010	EHA Thrust Mode	V0138	Dischrg Voltage	V0208	Continuous Dump #2
V0011	EHA Startup Mode	V0139	Discharge Heater Current	V0209	Continuous Dump #3
V0012	EHA Throttle Mode	V0140	Discharge Heater Voltage	V0210	Continuous Dump #4
V0013	EHA Accel Current	V0141	HV line current	V0211	Continuous Dump #5
V0014	EHA Beam Current	V0142	HV line voltage	V0212	Continuous Dump #6
V0015	EHA Beam Voltage	V0143	Neutralizer Current	V0213	Continuous Dump #7
V0016	EHA Discharge Current	V0144	Neu. Voltage	V0218	Continuous Dump #8
V0017	EHA Discharge Voltage	V0145	Neutralizer Heater Current	V0219	Continuous Dump #9
V0018	EHA Neutralizer Voltage	V0146	Neutralizer Heater Voltage	V0220	Peek Memory Offset
V0019	EHA Neutralizer Common	V0147	Neutralizer Common	V0221	Peek Memory Segment
V0020	EHA PT1 Pressure	V0148	+5V Ref	V0222	Peek Memory #0
V0021	EHA XFS Temperature TP1	V0149	PPU [RT-1] Temp	V0223	Peek Memory #1
V0022	EHA XFS Temperature TP4	V0150	PPU Temp. [RT-2, Neu. Sw., Q1]	V0224	Peek Memory #2
V0023	EHA Measured Press. 1	V0151	PPU Temp. #3 [RT-3, Screen]	V0225	Peek Memory #3
V0024	EHA Measured Press. 2	V0152	PPU Temp. #4 [RT-4, Disc. Rect.]	V0226	Peek Memory #4
V0025	EHA Echo DCIU command	V0153	+5V PPU	V0227	Peek Memory #5
V0026	# of IPS commands received	V0154	+15V PPU	V0228	Peek Memory #6
V0027	# of 1553 commands pending	V0155	-15V PPU	V0229	Peek Memory #7
V0028	Greatest # 1553 commands pending	V0156	Discharge Cmd Level	V2510	Gimbal Pot Voltage
V0029	IPS telemetry period	V0157	Discharge Heater Cmd Level	V2512	Gimbal 1 (+X+Y)
V0030	Lower mission power level	V0158	Neutralizer Cmd Level	V2520	Gimbal 2 Pot Voltage
V0031	Upper mission power level	V0159	Neutralizer Heater Cmd Level	V2522	Gimbal 2 (+X-Y)
V0032	DCIU thrust level	V0160	Screen Cmd Level	V3100	Boot Load Mode
V0033	Desired thrust duration (s)	V0161	Accelerator Cmd Level	V3101	Safe Mode Status
V0034	Thrusting?	V0162	PPU Digital Input: Bit 0 = Recycle Flag	V3102	Standby Mode
V0035	Thrust period cum.	V0163	PPU Digital Output	V3103	Grid Clear Mode
V0036	Cum. since last update	V0164	XFS Data Packet ID	V3104	Cathode Cond. Mode
V0037	Accumulated thrust mag.	V0165	PT1 Pressure	V3105	Thrust Mode
V0038	# of packets since last DCIU telem.	V0166	PA1 Pressure	V3106	XFS ON Mode

Deep Space 1 Technology Validation Report—*Ion Propulsion System (NSTAR)*

Channel	Title of Parameter	Channel	Title of Parameter	Channel	Title of Parameter
V0039	Processing recycle?	V0167	PA2 Pressure	V3107	XFS Initialization
V0040	DCIU heartbeat	V0168	PA3 Pressure	V3116	Recycle Flag
V0100	DCIU Data Packet ID	V0169	PA4 Pressure	V3132	Neutralizer Htr Enable
V0101	DCIU Time	V0170	PA5 Pressure	V3133	Discharge Htr Enable
V0103	DCIU command accepted counter.	V0171	PA6 Pressure	V3134	Neutralizer Enable
V0104	# Cmd rejected	V0172	XFS Temp TP1	V3135	Discharge Enable
V0105	Power Level Checksum	V0173	XFS Temp TP2	V3136	Beam Enable
V0106	Command #0	V0174	XFS Temp TP3	V3140	Thruster A Select
V0107	Command #1	V0175	XFS Temp TP4	V3141	Thruster B Select
V0108	Command #2	V0176	XFS Temp TP5	V3142	Grid Clear Enable
V0109	Command #3	V0177	XFS Temp TP6	V3143	Recycle Clear
V0110	Error #0	V0178	Main Flow Temp TJ1	V3148	Neutralizer Lit
V0111	Error #1	V0179	Cathode Flow Temp TJ2	V3149	Discharge Lit
V0112	Error #2	V0180	Neutralizer Flow Temp TJ3	V3150	Beam Supply Lit
V0113	Error #3	V0181	Regulator 1 Temp TR1	V3151	Grid Clear Required
V0114	# of Errors	V0182	Regulator 2 Temp TR2	V3152	Neutralizer Heater Open
V0115	DCIU +5V	V0183	SV1/3 Pulse Width	V3153	Discharge Heater Open
V0116	DCIU +15V	V0184	SV2/4 Pulse Width	V3154	Grid Clear Fail
V0117	DCIU -15V	V0185	SV1/2,3/4 Delay Width	V3156	Thruster A Status
V0118	+28V Bus Current	V0186	SV2/1,4/3 Delay Width	V3157	Thruster B Status
V0119	Processing Time	V0187	Latch Valve Width	V3158	Neutralizer Failed to Light
V0120	Power Level	V0188	Measured Pressure 1	V3159	Discharge Failed to Light
V0121	XIPS Mode	V0189	Required Pressure 1	V3160	Multiple Recycle Flag
V0122	Safe Mode	V0190	Measured Pressure 2	V3161	Continuous Recycle Flag
V0123	Grid Clear Mode	V0191	Required Pressure 2	V3162	Beam Control Enable
V0124	Cathode Conditioning Mode	V0192	Number SV1 Cycles	V3163	Diode Mode Enable
V0125	Thrust Mode	V0194	Number SV2 Cycles	V3164	Beam Voltage 5% error
V0126	Startup Mode	V0196	Status Data Packet	V3165	Beam Current 5% error
V0127	Throttle Mode	V0197	PPU Status Word 1	V3166	Accel Voltage 5% error
V3167	Accel Current 5% error	V3300	Shutdown Heaters Off	V4068	DSEU1 temp.
V3168	Discharge Voltage 5% error	V3301	XSHCLSVL	V4069	DSEU2 temp.
V3169	Discharge Current 5% error	V3319	XFS Initialization Mode	A0945	Pulses X3
V3170	Neutralizer Voltage 5% error	V3320	XFS Run Mode Status	A0947	Pulses X4
V3171	Neutralizer Current 5% error	V3329	Software Version - Minor Revision	A0949	Pulses Z1
V3172	Beam Voltage 10% error	V3330	Software Version - Major Revision	A0952	Pulses Z2
V3173	Beam Current 10% error	V3401	Ingested mass flow	A0954	Pulses Z3
V3174	Accel Voltage 10% error	V3402	Main flow rate	A0956	Pulses Z4
V3175	Accel Current 10% error	V3403	Cathode flow rate	A0958	Pulses X1
V3176	Discharge Voltage 10% error	V3404	Neutralizer flow rate	A0961	Pulses X2
V3177	Discharge Current 10% error	V3405	Total flow rate	A1401	Sun from X axis (Cos)
V3178	Neutralizer Voltage 10% error	V3406	Total main flow rate	A1402	Sun from Y axis (Cos)
V3179	Neutralizer Current 10% error	V3407	Total mass flow	A1403	Sun from Z axis (Cos)
V3180	XFS Normal Mode	V3408	Beam voltage	A1640	X3 RCS on-time
V3181	XFS Single Plenum Mode	V3409	Beam current	A1646	X4 RCS on-time
V3182	Single Main	V3410	Total Eng Pwr	A1650	Z1 RCS on-time
V3183	Single Cathode	V3411	Discharge loss	A1658	Z2 RCS on-time
V3184	Fault Protection Ena/Dis	V3412	Total prop. util. eff.	A1666	Z3 RCS on-time
V3185	XFS Initialized	V3413	Discharge prop. util. eff.	A1676	Z4 RCS on-time

Deep Space 1 Technology Validation Report—*Ion Propulsion System (NSTAR)*

Channel	Title of Parameter	Channel	Title of Parameter	Channel	Title of Parameter
V3196	Latch Valve #1 Open/Close	V3414	Xe double ion fraction	A1685	X1 RCS on-time
V3197	Latch Valve #2 Open/Close	V3415	Thrust loss factor	A1692	X2 RCS on-time
V3198	Latch Valve #3 Open/Close	V3416	Thrust	P2030	Solar Array Voltage
V3199	Latch Valve #4 Open/Close	V3417	Specific impulse	P2040	Solar Array 1 Current
V3200	Latch Valve #5 Open/Close	V3418	Overall thrust eff.	P2050	Solar Array 2 Current
V3201	Safe Mode Start	V3419	EHA mission power level	P2060	Essential Bus Current
V3202	Safe Mode Shutdown	V3420	EHA/IPS mission power level	P2061	Essential Bus Voltage
V3203	Safe Mode Close Valves	V3421	Mssn Th Chck Sum	P2062	Bus 1 Current
V3217	Grid Clear Start	V3422	P1 Measured - Req.	P2063	Bus 1S Current
V3218	Grid Clear Light Discharge	V3423	P2 Measured - Req.	P2064	Bus 2 Current
V3219	Grid Clear Check Jb	V3424	Vb Meas - Tbl	P2065	Bus 3 Current
V3220	Grid Clear Terminate	V3425	Vb Meas - Tbl	P3072	PPU Input Power
V3221	Grid Clear Reset	V3426	Va Meas - Tbl		
V3233	Cathode Conditioning Start	V3427	Ja Meas - Tbl		
V3234	Cathode Conditioning Heat 1	V3428	Vd Meas - Tbl		
V3235	Cathode Conditioning Cool 1	V3429	Jd Meas - Tbl		
V3236	Cathode Conditioning Heat 2	V3430	Vn Meas - Tbl		
V3237	Cathode Conditioning Cool 2	V3431	Jn Meas - Tbl		
V3238	Cathode Conditioning Terminate	V3435	Main Err. SV1 - SV2 Cycles		
V3239	Cathode Conditioning Reset	V3436	Cathode Err. SV3 - SV4 Cycles		
V3249	Thrust Startup	V3437	Set Beam Voltage		
V3250	Thrust Throttle	V3438	Set Beam Current		
V3251	Thrust Steady State	V3439	Set Accel Voltage		
V3252	Thrust Shutdown	V3440	Set Accel Current		
V3253	Thrust Shutdown XFS	V3441	Set Discharge Voltage		
V3265	Startup Start	V3442	Set Discharge Current		
V3266	Startup XFS Init	V3443	Set Neutralizer Voltage		
V3267	Startup Preheat Both	V3444	Set Neutralizer Current		
V3268	Startup Preheat Discharge	V3445	Set Main Pressure		
V3269	Startup Ignite Neutralizer	V3446	Set Cathode Pressure		
V3270	Startup Ignite Discharge	V3447	Set Single Plenum Pressure		
V3271	Startup Cool Both	V3448	Req. Cathode flow		
V3272	Startup Cool Discharge	V3449	Req. Neut. flow		
V3273	Startup High Voltage On	V3450	Main Flow Error		
V3274	Startup Ignition Failure	V3451	Main Cathode Error		
V3281	Throttle Start	V3452	Neutralizer Error		
V3282	Throttle Down Neutralizer	V3453	Req. Main flow		
V3283	Throttle Down Discharge	V4002	Temp		
V3284	Throttle Down Beam	V4051	DCIU Temp 1		
V3285	Throttle Down Accelerator	V4052	PPU Temp 1		
V3286	Throttle Down XFS	V4053	PPU Temperature 2		
V3287	Throttle Up Neutralizer	V4054	Xenon Temp		
V3288	Throttle Up Discharge	V4061	Gimbal 1 (+X+Y) Temp.		
V3289	Throttle Up Beam	V4062	Gimbal 2 Temp.		
V3290	Throttle Up Accelerator	V4063	Gim Brckt Temp		
V3291	Throttle Up XFS	V4064	Thrstr Msk Temp		
V3297	Shutdown Start (Beam Off)	V4065	Xe tank temp		
V3298	Shutdown Discharge Off	V4066	DCIU temp.		
V3299	Shutdown Neutralizer Off	V4067	Thruster Temp.		

Appendix B. Date of Turn-on/off and Frequency of Data Capture

DATE OF TURN-ON/OFF

Below is the list of the IPS technology validation activities and beam on and off times that took place on DS1. The total

accumulated hours as of 1999-30T00:00 is 3575 hours.
(Ken Fujii, 12/16/99.)

Table B1. Beam On/Off Time

Beam On Time	Beam Off Time	Duration (hr)	Event
1998-314T193426	1998-314T193926	0.08	Initial IAT Attempt
1998-328T225224	1998-342T220440	335.20	IAT0
1998-346T004902	1998-346T025300	2.07	IPS arc 1.1
1998-348T221838	1998-352T214040	95.37	IPS arc 1.1
1998-352T225317	1998-355T205537	70.04	IPS arc 1.1
1998-356T011959	1998-356T204419	19.41	IPS arc 1.1
1998-356T215719	1999-005T160009	330.05	IPS arc 1.1
1999-022T213604	1999-022T221636	0.68	SPEak
1999-057T231116	1999-058T001100	1.00	IPS Readiness Test
1999-075T071448	1999-081T195503	156.67	IPS arc 1.2 (C1A NBURN1)
1999-082T130932	1999-088T113958	142.51	IPS arc 1.2 (C1A NBURN2)
1999-089T040828	1999-095T160458	155.94	IPS arc 1.2 (C1A NBURN3)
1999-096T171034	1999-102T162959	143.32	IPS arc 1.2 (C1B NBURN1)
1999-103T090017	1999-109T162957	151.49	IPS arc 1.2 (C1B NBURN2)
1999-110T090642	1999-116T120458	146.97	IPS arc 1.2 (C1B NBURN3)
1999-117T042258	1999-117T173458	13.20	IPS arc 1.2 (C1B NBURN4)
1999-138T095155	1999-139T001015	14.31	RAX
1999-148T090818	1999-148T222257	13.24	IAT2
1999-165T160444	1999-165T201604	4.19	IPS Test TCM 1
1999-166T041229	1999-166T082604	4.23	IPS Test TCM 2
1999-204T225503	1999-205T011918	2.40	ACA-5 day TCM
1999-211T160802	1999-214T044617	60.64	Post Encounter NBURN
1999-214T220803	1999-221T065954	152.86	E1C NBURN
1999-222T042520	1999-228T075521	147.50	IPS arc 2.1 (C2A NBURN1)
1999-228T204023	1999-235T062023	153.67	IPS arc 2.1 (C2A NBURN2)
1999-237T025731	1999-242T211021	138.21	IPS arc 2.1 (C2A NBURN3)
1999-243T064953	1999-249T152033	152.51	IPS arc 2.1 (C2A NBURN4)
1999-250T025452	1999-256T174922	158.91	IPS arc 2.1 (C2B NBURN1)
1999-257T031431	1999-263T190922	159.91	IPS arc 2.1 (C2B NBURN2)
1999-264T085352	1999-270T183922	153.76	IPS arc 2.1 (C2B NBURN3)
1999-271T075502	1999-277T163922	152.74	IPS arc 2.1 (C2B NBURN4)
1999-278T054921	1999-284T162753	154.64	IPS arc 2.1 (C2C NBURN1)
1999-285T054831	1999-291T152752	153.66	IPS arc 2.1 (C2C NBURN2)
1999-292T061351	1999-293T114252	29.48	IPS arc 2.1 (C2C NBURN3)
1999-310T010846	1999-310T050757	3.99	MICAS Pointing Test

FREQUENCY OF DATA CAPTURE

The IPS telemetry rate was limited by the speed of spacecraft software, the size of the spacecraft memory, the spacecraft telemetry rate as a function of Earth distance, spacecraft orientation, the selected DSN station, and the needs of other competing users.

The maximum IPS data rate was 2048 bits per second. This occurred when all of the IPS data was sampled once every second. By selecting a smaller subset of data and sampling at a lower rate, the IPS data rate was varied from 2048 bits per second to 2 bits per second when the IPS was thrusting.


The limited speed of the spacecraft's telemetry system limited the maximum average IPS data-sample rate to once every two seconds (although one sample per second rate was used for short periods of time). Most of the early IPS

telemetry was at a 10 seconds per sample rate, or 200 bits per second.

After initial IPS checkout, spacecraft telemetry was greatly reduced because of reduced link performance and DSN-station passes. The IPS-sample rate was reduced to one sample every 5 minute. This reduced the IPS data rate to less than 7 bits per second.

As the Earth distance increased, it was necessary to further reduce spacecraft telemetry. The IPS was sampled once every 15 minutes, resulting in a data rate of 2 bits per second. It is expected that the data rate will be reduced to 1/2 bit per second during the latter portion of the mission.

By using proper data selection, the data rate could be easily reduced by a factor of four. It is envisioned that, using onboard logic, future missions will not need to communicate with the IPS unless there is a fault.



Ion Propulsion Subsystem Environmental Effects on Deep Space 1: Initial Results from the IPS Diagnostic Subsystem DS1 Technology Validation Report

David E. Brinza, Michael D. Henry, Anthony T. Mactutis, Kenneth P. McCarty, Joel D. Rademacher, Thomas R. van Zandt, Joseph J. Wang and Bruce T. Tsurutani

*Jet Propulsion Laboratory, California Institute of Technology
Pasadena, California 91109*

Ira Katz and Victoria A. Davis
*Maxwell Technologies
San Diego, California 92123*

Guenter Musmann, Falko Kuhnke, Ingo Richter, Carsten Othmer,
Karl-Heinz Glassmeier
*Institute for Geophysics and Meteorology, Technical University of
Braunschweig
Braunschweig, Germany*

Stewart Moses
*TRW
Redondo Beach, California 90278*



Table of Contents

<u>Section</u>	<u>Page</u>
Extended Abstract	vi
Overview	vi
Background	vi
IPS Effects on Spacecraft Potential	vi
Contamination from IPS	vi
IPS-Generated Plasma Noise and EMI	vii
DC Magnetic Fields from IPS	vii
Ion Propulsion Subsystem Environmental Effects on Deep Space 1: Initial Results from the IPS Diagnostics Subsystem	1
DS1 Technology Validation Report	1
Abstract	1
1.0 Introduction	1
1.1 Plasma Environment	1
1.2 Fields Environment	2
1.3 Contamination Environment	3
2.0 Diagnostics Element Description	3
2.1 Ground Test Diagnostics	3
2.2 Modeling Tools	4
2.3 IPS Diagnostics Subsystem on DS1	4
3.0 Charge-Exchange Plasma	7
3.1 DS1 Chassis Potential Without Langmuir Probe Bias Voltage	7
3.2 RPA Current as a Function of IPS Mission Level	10
3.3 Variation of RPA Current with Ring Langmuir Probe Bias	10
3.4 Expansion of Langmuir Probe Bias Potential onto Black Kapton	11
3.5 Expansion of Charge-Exchange Plasma Around DS1	11
4.0 Contamination Assessment	13
4.1 Ground Test Contamination Results	13
4.2 Flight Contamination Results	14
5.0 IPS Plasma Wave & EMI Characteristics	19
5.1 Plasma Wave Electric-field Measurements	19
5.2 AC Magnetic Fields (EMI)	21
5.3 Plasma Wave Transient Signals	22
6.0 IPS DC-Magnetic Fields	25
6.1 Ground Magnetic Field Mapping	25
6.2 Flight Measurements	26
7.0 Conclusions	31
8.0 Acknowledgments	32
9.0 References	32
Appendix A. List of Telemetry Channels and Names	34
Appendix B. Date of Turn-on/off and Frequency of Data Capture	34
Appendix C. List of Acronyms and Abbreviations	35

Figures

<u>Figures</u>	<u>Page</u>
Figure 1. IPS Diagnostics Subsystem Hardware	5
Figure 2. IPS Diagnostics Subsystem Block Diagram	5
Figure 3. Major Contributors to Current Balance on DS1	8
Figure 4. Computed Ion Density Contours for NSTAR Ion Engine at Full Power (dimensional scale is meters).....	8
Figure 5. Self-consistent Potential Computed for NSTAR Thruster Operating at Full Power (dimension scale in meters).....	8
Figure 6. Variation of IPS Neutralizer Common with IDS Langmuir Probe Bias Voltage in IAT2	9
Figure 7a–d. RPA Sweeps Obtained at Each Langmuir Probe Voltage Level.....	9
Figure 8. Estimated Langmuir Probe Potential Versus Probe Bias Voltage	10
Figure 9. Computed RPA Ion Current as a Function of IPS Mission Levels (measured currents from IAT2 also shown)	10
Figure 10. Effect of Langmuir Probe Bias on RPA Current	11
Figure 11. Ion Focusing by RPA Langmuir Probe.....	11
Figure 12. Expansion of Langmuir Probe Bias onto RSU Thermal Blanket.....	11
Figure 13. Langmuir Probe Least Squares Fit ($\theta = 1.8$ eV)	11
Figure 14. Model Geometry for PIC Simulation.....	12
Figure 15. Contour Plots for: (a) Beam Ion Density, (b) Neutral Xenon Density, and (c) Charge-Exchange Ion Production Rates	12
Figure 16. Results of the DS1 PIC Simulation: (a) Electric Field Potential, (b) Electric Field Direction, (c) Charge-Exchange Ion Densities, and (d) Charge-Exchange Ion Flow Directions	13
Figure 17. Geometry of Collimated Contamination Monitors for the NSTAR LDT (drawing is <u>not</u> to scale).....	14
Figure 18. Molybdenum Accumulation for NSTAR LDT Witness Monitors Versus Angle from Thruster Axis	14
Figure 19. QCM1 Early Mission Response.....	15
Figure 20. Early Mission DS1 Sun Orientation.....	15
Figure 21a-h. QCM0 and QCM1 Data for IPS Operations of the First Year of Flight for DS1.....	16
Figure 22. Mo Deposition Rates Versus Mission Level (QCM0 is the line-of-sight sensor).....	18
Figure 23. Ratio of Non-Line-of-Sight to Line-of-Sight Deposition Rates as a Function of Mission Level	18
Figure 24. Response of IDS Line-of-Sight Calorimeter During Initial IPS Operations.....	18
Figure 25. Solar Irradiance History for the Line-of-Sight Calorimeter During Initial IPS Operations	19
Figure 26. Plasma Wave Spectrum for ICT Thrust Levels	20
Figure 27. Plasma Wave Spectrum for S-Peak Thrust Levels.....	20
Figure 28. Plasma Wave Spectra for IAT1 Mission Levels.....	20
Figure 29. Plasma Wave Spectrum for IAT2 Mission Levels.....	21
Figure 30. Plasma Wave Spectrogram for IPS Transition from ML20 to ML27	21
Figure 31. Plasma Wave Spectrogram for IPS Transition from ML83 to ML90.....	21
Figure 32. Response of the Search Coil Magnetometer to IPS Start During Ground Testing.....	21
Figure 33. AC Magnetic Spectra for IAT1 Mission Levels	22
Figure 34. AC Magnetic Field Spectra for IAT2 Mission Levels	22
Figure 35. IPS Ignition in CT36 Ground Test (monopole)	22
Figure 36. IPS Ignition in CT36 Ground Test (PWA dipole)	23
Figure 37. E-field Transient Signal for Flight IPS Ignition.....	23
Figure 38. B-Field Transient Signal for Flight IPS Ignition.....	23
Figure 39. E-field Signature for IPS Recycle at $t = -0.45$	23
Figure 40. B-field Signature for IPS Recycle at $t = -0.45$	23
Figure 41. E-field Signature for RCS Thrusters Firing at $t = 0$	24
Figure 42. B-field Signature for RCS Thrusters Firing at $t = 0$	24
Figure 43. E-field Signature for Particle Impact at $t = 0$	24
Figure 44. B-field Recording for Particle Impact at $t = 0$	24
Figure 45. Magnetic Field Model for the NSTAR Ion Engine.....	25
Figure 46. Pre-LDT Magnetic Map of EMT#2 Thruster.....	25
Figure 47. The 24-hour Averaged, Calibrated FGM Data in DS1 Coordinates	26
Figure 48. The 24-hour Averaged Temperature Data from Thruster and FGM Sensors.....	27
Figure 49. The Linear Temperature Model of the IPS Thruster Magnetic Field.....	28

Deep Space 1 Technology Validation Report—*Ion Propulsion Subsystem Environmental Effects on
Deep Space 1: Initial Results from the IPS Diagnostics Subsystem*

Figure 50. Residual Magnetic Field Data After Temperature Correction 29
Figure 51. Temperature-corrected FGM Field Data versus FGM Temperature..... 30

Tables

<u>Tables</u>	<u>Page</u>
Table 1. Effect of Langmuir Probe Bias on Ion Energy and Neutralizer Common.....	9
Table 2. Relevant IPS Operating Conditions for Mission Levels 6 and 13.....	10
Table 3. Selected Parameters for Estimating the Change in Thermo-optical Properties.....	19

EXTENDED ABSTRACT

Overview

The Deep Space 1 (DS1) mission has successfully validated the use of ion propulsion technology for interplanetary spacecraft. The NASA Solar Electric Propulsion (SEP) Technology Applications Readiness (NSTAR) Project developed the Ion Propulsion Subsystem (IPS) for DS1. As part of the IPS validation effort, the NSTAR Project included a Diagnostics Element to characterize the local environment produced during IPS operations and its effects on spacecraft subsystems and science instruments. An integrated, comprehensive set of instrumentation was developed and flown on DS1 as the IPS Diagnostics Sensors (IDS) subsystem. During the technology validation phase of the DS1 mission, data were collected from the IDS under a variety of IPS operating conditions. IDS characterized the local plasma and contamination environments, electrostatic and electromagnetic noise, and magnetic fields associated with IPS.

Background

The DS1 IPS generates thrust by ejecting a beam of high-velocity (>30 km/s) xenon ions from the thruster. Ions are created within the discharge chamber of the engine via electron impact and are accelerated through ion optic grids to form the ion beam (see Figure 1). The fraction of xenon ionized in the discharge chamber is 80% to 90%. The xenon atoms that are not ionized in the discharge chamber diffuse through the grid and into space. The high-velocity beam ions and thermal-velocity atoms interact via a process referred to as resonant charge exchange in which an electron is transferred to the beam ion from the neutral xenon atom outside of the engine. This charge-exchange xenon (CEX) ion is accelerated by the electrostatic potential in the region where it was created. Electrons from the neutralizer balance the electric charge due to the beam and CEX ions. CEX ions strongly affect the chassis potential, the local contamination environment, and the plasma wave noise produced by IPS.

IPS Effects on Spacecraft Potential

The CEX ions formed downstream of the IPS engine grids are pushed by the electrostatic potential within the ion beam plume. Some of the CEX ions are accelerated roughly perpendicular to the thrust vector. The paths of these ions are influenced by electric fields around DS1. As a result, a relatively cold (1 to 2 eV) flowing plasma surrounds the DS1 spacecraft. Most of the current from the ion engine is collected by the grounded thruster “mask” near the grids. The major components that affect IPS current balance are shown in Figure 2. IPS current balance establishes the spacecraft potential. IDS has determined CEX plasma ion energies (12 to 21 eV), densities (10^{12} to 10^{13} m $^{-3}$) and electron temperatures (1.2 to 2.0 eV). The results were used to estimate the spacecraft potential. Depending on IPS operating conditions, the potential of the DS1 chassis is -6 eV to -10 eV with respect to solar wind “ground.” The

potential causes CEX ions to follow curved paths and even “orbit” the DS1 spacecraft. Mounted on the opposite side of DS1, the Plasma Experiment for Planetary Exploration (PEPE) instrument detected CEX ions in addition to solar wind protons during IPS operations.




Figure 1. Principal Elements of Ion Engine Operation




Figure 2. Major Components for Current Balance on DS1

Contamination from IPS

Significant amounts of CEX ions are formed very near the grid, where the neutral density and beam currents are highest. These CEX ions are accelerated into the outer engine grid with sufficient energy to physically knock atoms (molybdenum) from the grid via a process called sputtering. This leads to grid erosion, a wear mechanism that can continue until mechanical failure of the grid. The sputtered molybdenum atoms from the grid are ejected in a broad pattern from the engine and, due to their low-volatility, represent a contamination risk for sensitive surfaces on the spacecraft. The IDS has measured the contamination environment at the Remote Sensors Unit (RSU) and has found that the direct line-of-sight deposition rates of

molybdenum correlate reasonably well with ground test experience (Figure 3). Non-line-of-sight transport, due to ionized molybdenum ions, was also characterized in flight, a measurement that is made difficult in ground test because of chamber effects. The IPS logged 3,500 operating hours in the first year of flight with 250 Å (25 nm) of molybdenum deposited on line-of-sight contamination monitors; only 25 Å accumulated on nearby sensors shadowed from direct view of the engine grid.

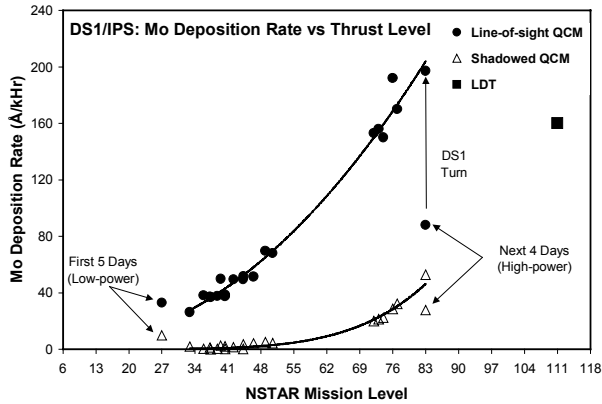


Figure 3. Mo Deposition Rates on Line-of-Sight and Shadowed Monitors During IPS Operations

IPS-Generated Plasma Noise and EMI

Ground tests and flight experiments show that hollow cathode devices produce substantial noise in the low-frequency (<50 MHz) regime. Electrical noise produced within the discharge of the neutralizer is conducted by the CEX plasma medium. IDS has measured the plasma noise and electromagnetic fields associated with IPS operations. Noise spectra for selected operating levels are shown in Figure 4. Transient voltage spikes (<2 V/m) due to IPS “arcing” events are comparable to those observed for hydrazine thruster firings. The largest amplitude EMI, based on search coil measurements, is from engine gimbal actuators used for thrust vector control. The IPS plume does not affect the telecommunications link.




Figure 4. Plasma Noise for Selected IPS Thrust Levels

DC Magnetic Fields from IPS

The NSTAR engine utilizes rare-Earth permanent magnet rings to improve the ionization efficiency within the discharge chamber. The magnetic fields from IPS are substantial (12,000 nT at 1 m) and are symmetric about the thrust axis. IPS magnetic field configuration is shown in Figure 5. IDS has determined the temperature dependence of the IPS magnetic fields. Analysis of the residual field after temperature correction and gimbal position to assess long-term field stability is in progress. Temporal stability of the IPS field would permit background subtraction, thereby allowing external fields to be determined.





Figure 5. DC Magnetic Field Map for IPS Engine



NMP DS1 FACT SHEET

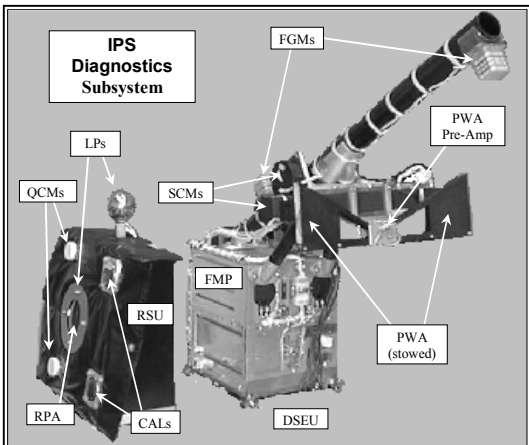



NSTAR

IPS Diagnostic Sensors (IDS)

Goal
Understand the *in-situ* (local) environment of a spacecraft using an ion propulsion system (IPS).

- Approach**
- Perform ground and spaceflight measurements of the following critical IPS environmental factors:
 - Plasma, contamination
 - AC/DC electric, magnetic fields
 - Develop & validate predictive models for future ion propulsion missions



Instrument Description
Twelve environmental sensors in two interconnected units: (Mass: 8 kg, Power: 21 W)

Remote Sensors Unit (RSU):
Plasma: two Langmuir Probes (LPs), Retarding Potential Analyzer (RPA)
Contamination: two Quartz Crystal Microbalances (QCMs), two Calorimeters (CALs)
Diagnostic Sensor Electronics Unit/Fields Measurement Processor (DSEU/FMP):
Electrostatic Fields: 2-m dipole Plasma Wave Antenna (PWA) with pre-amplifier
Electromagnetic Waves: two Search Coil Magnetometers (SCMs); one failed
DC Magnetic Fields: two ea. three-axis Flux-Gate Magnetometers (FGMs)

IDS Partners:

Jet Propulsion Laboratory:	Systems Engineering, FMP, PWA, SCM, Structure, I&T, Mission Operations
Physical Sciences, Inc.:	DSEU Electronics, Calorimeters
Maxwell Technologies:	Plume modeling
QCM Research:	Quartz Crystal Microbalances
Technical University of Braunschweig:	Flux-Gate Magnetometers
TRW:	Plasma Wave Spectrometer, Pre-amp

Sensor Specifications:

Sensor	Measurement	Range	Resolution
QCMs	Mass/area	0 to 500 $\mu\text{g}/\text{cm}^2$	0.005 $\mu\text{g}/\text{cm}^2$
CALs	Solar Absorptance (α)	$\alpha = 0.08$ (BOL) to 0.99	$\Delta\alpha = 0.01$
	Hemi. Emittance (ϵ)	$\epsilon = 0.05$ to 0.85 (BOL)	$\Delta\epsilon = 0.01$
LPs	Probe Current	$I = -0.4$ to 40 mA	1%
	Probe Voltage	$V = -11$ to +11 VDC	1%
RPA	Current (Gain Select)	$I = 0.01, 1, 10, 100\mu\text{A}$	1%
	Grid Bias Voltage	$V = 0$ to +100 VDC	0.4V
PWA	E-field (Adjust. Gain)	50 to 160 dB $\mu\text{V}/\text{m}$	± 3 dB $\mu\text{V}/\text{m}$
	24 Freq. Channels *	10 Hz to 30 MHz (4/decade)	$\pm 40\%$ (-3 dB)
SCM	B-field (Adjust. Gain)	80 to 160 dBpT	± 3 dBpT
	16 Freq. Channels *	10 Hz to 100 kHz (4/decade)	$\pm 40\%$ (-3 dB)
FGMs	Magnetic Field Vector **	$\pm 25,000$ nT	0.5 nT

* 20 kHz waveform capture (1 sec)
** 20 Hz B-vector waveform capture (up to 55 sec)

Programmatic:
Funded by the NSTAR Project with deeply appreciated support from JPL/TAP, DARA, TRW and NMP

Point-of-contact:
David.E.Brinza@jpl.nasa.gov
Jet Propulsion Laboratory 125-177
4800 Oak Grove Drive
Pasadena, CA 91109
(818)354-6836

- Key Findings:**
- IPS plasma drives DS1 chassis -6 to -10 V with respect to solar wind "ground"
 - Chamber tests can permit electrical "short" between chassis and IPS plume potentials
 - Line-of-sight contamination from IPS molybdenum grids comparable to ground measurement
 - Plasma waves <120 dB $\mu\text{V}/\text{m}$; IPS transients comparable to DS1 hydrazine thruster events
 - IPS permanent magnetic field vs temperature determined; field stability not yet verified (Jan.'00)

Ion Propulsion Subsystem Environmental Effects on Deep Space 1: Initial Results from the IPS Diagnostics Subsystem DS1 Technology Validation Report

*David E. Brinza, Michael D. Henry, Anthony T. Mactutis, Kenneth P. McCarty, Joel D. Rademacher, Thomas R. van Zandt, Joseph J. Wang and Bruce T. Tsurutani
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California*

*Ira Katz and Victoria A. Davis
Maxwell Technologies, San Diego, California*

*Stewart Moses
TRW, One Space Park, Redondo Beach, California*

*Guenter Musmann, Falko Kuhnke, Ingo Richter, Carsten Othmer, Karl-Heinz Glassmeier
Institute for Geophysics and Meteorology, Technical University of Braunschweig, Braunschweig, Germany*

ABSTRACT

The Deep Space 1 (DS1) mission has successfully validated the use of ion propulsion technology for interplanetary spacecraft. The NASA Solar Electric Propulsion (SEP) Technology Applications Readiness (NSTAR) Project developed the Ion Propulsion Subsystem (IPS) for DS1. As part of the NSTAR validation effort, the NSTAR Project included a diagnostics element to characterize the local environment produced during IPS operations and its effects on spacecraft subsystems and science instruments. An integrated, comprehensive set of diagnostics, the NSTAR Diagnostics Package (NDP) was developed and operated on DS1 to characterize the IPS environment. The DS1 Spacecraft Team officially assigned the name “IPS Diagnostics Subsystem (IDS)” to the NDP for the DS1 mission. During the technology validation phase of the DS1 mission, a large amount of data was collected from the IDS under a variety of IPS operating conditions. IDS was able to characterize the contamination environment, charge-exchange xenon ion and electron population and energies, plasma noise and electromagnetic noise, and magnetic fields associated with IPS. The initial results presented here describe the charge-exchange plasma, contamination, plasma wave/EMI, and DC magnetic environments critical to designers of future space missions using ion propulsion.


1.0 INTRODUCTION

This introduction is intended to provide the reader with a brief overview of Ion Propulsion Subsystem (IPS) environmental perturbations considered important for spacecraft and science operations. The objective for the IPS Diagnostics Subsystem (IDS) flown on Deep Space 1 (DS1) is to characterize these environments within significant resource constraints. The technical requirements for IDS measurements are based upon the results from the NASA/USAF Workshop on Environmental Diagnostics for ELITE/STAR[1].

The NASA Solar Electric Propulsion (SEP) Technology Applications Readiness (NSTAR) ion thruster operating aboard DS1 generates a local environment that includes electrostatic, magnetic and electromagnetic fields, charged particles, and neutral particles. The thruster environmental components, in combination with the natural space environment and the space vehicle, produce the “induced environment.” The induced environment has the potential of impacting the performance of spacecraft subsystems or science sensors. Based on the operating experience thus far on DS1, the IPS induced environment is benign to spacecraft subsystems.

1.1 Plasma Environment

The operation of the ion thruster with neutralizer generates a plasma flow about the spacecraft[2]. The primary beam (1 kV) xenon ions interact with thermal energy xenon atoms diffusing from the thruster via a resonant charge exchange process to generate low-energy ions in the plume:



The total charge-exchange ion current generated is estimated to be less than 5 mA for the NSTAR thruster. These charge-exchange ions are accelerated by electric field gradients in the vicinity of the thruster, moving radially at energies up to 20 eV. Electrons are emitted from a hollow cathode neutralizer similar in design to the plasma contactor to be used on the International Space Station. The electrons from the neutralizer associate with the charge-exchange ions to create a cold, flowing plasma. This cold, flowing plasma effects the spacecraft in ways described in the paragraphs that follow.

1.1.1 Spacecraft Potential—Thruster operation might be expected to “clamp” the spacecraft potential to the local space plasma potential. The electron temperature (expected to be 1 to 3 eV) is expected to drive the spacecraft potential to no more than -10 V[3]. In the interplanetary environment, the Debye length is typically greater than 1

km; thus, direct measurement of spacecraft potential cannot be performed by Langmuir probe sensors. Electron temperature measurements, coupled with ion current and energy knowledge, and a reliable ion plasma-plume modeling tool are used to estimate the spacecraft potential.

1.1.2 Current Balance—The plasma flow produced by the thruster can provide a path for parasitic current loss from the solar arrays[4]. The extent of this current drain is determined by solar array design, spacecraft ground convention, solar array potential, and the plasma densities associated with the thruster. The currents from the solar array, through the ion thruster system, and the spacecraft bus were monitored as part of the DS1 engineering measurements. Analyses of these measurements are required to understand current balance within the spacecraft. The NSTAR IPS has an internal ground (neutralizer common) that is virtually isolated from the DS1 spacecraft ground. Potential measurements of the neutralizer common with respect to spacecraft ground provide information regarding current flow between the IPS and DS1. Effects of Langmuir Probe operation on IPS neutralizer common provide additional insight into current balance on DS1 during IPS operations.

1.1.3 Charge-Exchange Ion Interference—The density of charge-exchange ions from the NSTAR ion engine can present a risk to sensitive particle-detection instruments. Mass spectrometers designed to operate in solar-wind environments are typically particle-counting instruments with high-gain channel electron multipliers or other sensitive detectors. Measurements of the charge-exchange ion flux near the NSTAR engine is made with a retarding potential analyzer. The Plasma Experiment for Planetary Environments (PEPE) particle spectrometer measures electron and charge-exchange ion densities on the opposite side of the DS1 spacecraft.

1.1.4 Energetic Ion Impingement—The ion plume contains energetic ions (1 keV) that would erode surfaces exposed to direct impingement via sputtering. These ions are emitted from the thruster primarily (95%) in a cone with a half angle of about 45° about the thrust axis. Measurable energetic ion flux at higher off-axis angles may be found; however, their risk to spacecraft subsystems is low. Charge exchange ions may also sputter coatings; however, a current of less than 1 $\mu\text{A}/\text{cm}^2$ of low-energy ions (<20 eV) is expected at 75-cm distance from the thruster (at the exit plane). This charge-exchange ion flux is not expected to sputter material from spacecraft surfaces. Ground measurements of erosion (and contamination) were performed for long-duration tests. Flight measurements include a retarding potential analyzer with sub-nA sensitivity and bias voltages up to +100 VDC.

1.2 Fields Environment

The NSTAR thruster produces static electric and magnetic fields and electromagnetic disturbances during routine operation. The design of the thruster, neutralizer, and power processor unit (PPU) considers the conducted and emitted electromagnetic interference (EMI) effects. The interaction of the plume and charge exchange plasma with the natural environment and spacecraft power system can also generate electrostatic, magnetic, and electromagnetic fields. The following sections describe electric, magnetic, and electromagnetic fields effects induced by the NSTAR IPS on DS1.

1.2.1 Electrostatic Fields—Charge-exchange plasma associated with the NSTAR engine provides a conductive medium for time-varying electrostatic fields[5]. Plasma waves are generated in the region of the neutralizer by temporal instabilities in the hollow cathode discharge. The electron plasma frequency (ν) varies with the square root of electron density (n_e)[6]:

$$\nu \approx 8.98 (\text{Hz}\cdot\text{m}^3)^{1/2} n_e^{1/2}$$

The plasma density is expected to decrease from $10^{15}/\text{m}^3$ in the plume just outside the thruster to less than $10^{13}/\text{m}^3$ at one meter from the engine. Plasma waves have been measured for ion thrusters from very low frequencies (a few kilohertz) up to tens of megahertz. Due to locally strong magnetic fields, the plume is also a source of cyclotron electric fields. Flight measurements with an electric field antenna sensitive over the frequency range of 10 Hz to 30 MHz and a search-coil magnetometer from 10 Hz to 50 kHz are performed aboard DS1.

1.2.2 Electromagnetic Fields—The primary electromagnetic interference (EMI) concern with an IPS is its impact on the spacecraft communications system. In interplanetary missions, attenuation and phase delay due to the plume/plasma density may occur along the link path. Measurements in ground test have provided data for effective modeling of plume effects on RF electromagnetic wave propagation. Flight measurements utilizing the on-board telecommunications system were performed on DS1. The DS1 Mission Operations Team incorporated maneuvers with telecom operations to provide through-the-plume geometry for assessment of worst-case effects of ion thruster operations with spacecraft communications. No detectable change in telecom signal strength could be observed in this measurement.

High-level electromagnetic fields may arise from thruster operation from current fluctuations in the NSTAR propulsion system. The PPU was subjected to electromagnetic compatibility testing (such as RE101 from MIL-STD-461D) with the unit operating with a characteristic thruster load. Strong AC fields can impact scientific

instruments and possibly spacecraft subsystems. AC fields from the IPS will interfere with fields measurements; therefore, science fields measurements should be made only while the IPS is not thrusting. The space science community has interest in lower frequency EMI characteristics of ion propulsion system operations. Thrust-phase portions of the mission may limit particles and fields measurements. It was expected that the thruster beam would produce waves due to beam instabilities induced by the ambient environment. The search-coil magnetometer detects EMI over the frequency range of 10 Hz to 50 kHz; however, signals due to ion-beam solar wind have not been uniquely identified. Other EMI sources, such as the engine gimbal assembly (EGA), and solar array actuators have been detected on DS1.

1.2.3 Magnetic Fields—DC magnetic fields arise from permanent magnets used in the thruster design. The permanent magnets in the ion thruster are configured to maximize ionization efficiency[7]. The thruster body is constructed of titanium; therefore, DC magnetic fields surround the thruster. Measurements of the magnetic field pattern for the NSTAR ion engine indicate fields of nearly 5000 nT are expected at one meter from the thruster. Electrical currents through the spacecraft power system and ion propulsion system produce other stray magnetic fields. DC fields are a significant consideration in science missions where the magnetic fields are measured with high sensitivity. In a typical science mission, magnetometers are generally exposed to DC fields due to the spacecraft subsystems of less than 1 nT. For the DS1 technology validation mission, magnetic cleanliness of the spacecraft was not a major consideration. The DC magnetic fields measured on DS1 contains contributions from the NSTAR IPS and the rest of the DS1 spacecraft (heaters, solar arrays, other subsystems). As a goal, the flight magnetic measurements are intended to distinguish spacecraft fields from thruster-generated fields with better than 1-nT sensitivity.

1.3 Contamination Environment

The xenon propellant used in the NSTAR thruster is a non-contaminating species. One of the wear mechanisms for the thruster involves gradual sputtering of the molybdenum accelerator grids, eventually leading to mechanical failure of the grid structure[8]. Sputtered neutral molybdenum atoms are emitted in the general direction of the plume. Charge-exchange of the sputtered molybdenum with primary ion beams will occur (albeit with much smaller cross section than for resonant charge exchange of xenon). The charge-exchange molybdenum ions may be transported to surfaces “upstream” of the thruster. The upstream deposition rates are expected to be very low, even in the immediate vicinity of the thruster. However, even very thin coatings on the order of a few Angstroms (\AA , $1\text{\AA} = 10^{-10}$ m) can produce significant effects in thermo-optical (solar absorptance and emittance) properties of thermal control materials or

transmission of solar radiation through solar cell cover glasses.

The results from diagnostic sensors are useful from two perspectives: (1) The in-flight data provides a spacecraft systems engineer information for modeling environments on future spacecraft and (2) the data, when correlated with ground test, can help assess engine health. Contamination measurements can provide an indication of grid wear. The flight measurement will rely upon a calorimetric measurement of thermo-optical properties of a space-stable optical solar reflector supplemented with rate measurements via a quartz crystal microbalance (QCM).

2.0 DIAGNOSTICS ELEMENT DESCRIPTION

The NSTAR diagnostics effort includes ground test, modeling, and flight measurements to assess the environmental impact of ion-thruster operations on spacecraft payloads (instruments) and sub-systems. The validation of performance of the ion thruster sub-system includes direct measurement of phenomenology associated with the interactions described in the introduction. The ground test, modeling, and flight measurement approaches are described below.

2.1 Ground Test Diagnostics

The NSTAR thruster element included development and test of engineering model thruster (EMT) and flight thruster systems. The NSTAR contractor, Hughes Electron Dynamics Division, delivered flight thrusters with significant design heritage to the 30-cm xenon ion thrusters developed by the NASA Glenn Research Center[9]. Various ground tests were conducted throughout the NSTAR project, culminating with flight thruster compatibility tests with the DS1 spacecraft prior to launch. The following sections describe these NSTAR tests in the context of diagnostic measurements.

2.1.1 Early EMT Testing—The early EMT tests were moderate in duration (hundreds of hours up to 2000 hr) to characterize erosion characteristics, thruster performance, etc. During this phase, design details and operating points of the NSTAR thruster were adjusted to enhance thruster reliability and performance for long duration operation. Since minor changes to thruster design may substantially alter the contamination, EMI, or plasma conditions associated with the thruster, very few quantitative diagnostic tests were planned. A few witness materials were examined and qualitative measurements of EMI were performed; however, these tests remain geared to thruster evaluation.

2.1.2 Life Demonstration Test—The NSTAR program performed a life demonstration test (LDT) of an ion engine that successfully demonstrated the ability of the NSTAR EMT to operate at full power for more than 8000 hours[10].

The LDT afforded an excellent opportunity to collect contamination data and to establish flight plasma sensor design and performance requirements. Ground tests produced “chamber effects” that can interfere with the measurement of the relevant environments. Interaction measurements; however, there were mitigation approaches that provided useful data. Much of the data gathered from the LDT was of comparative nature: before and after grid mass, thrust vector stability, engine efficiency, etc. The NSTAR diagnostics element characterized the magnitude and stability of the DC magnetic field produced by the EMT before and after the LDT.

Contamination measurements in the LDT were considered valuable since the magnitude of erosion and deposition measurements scale with operating times, especially for witness specimen measurements. The NSTAR diagnostic element performed a contamination assessment during the LDT to quantify deposition amounts and/or erosion effects while providing an estimate of the contribution of chamber effects. “Collimated” witness specimens (fused silica windows) were located at various angles with respect to the plume axis. “Un-collimated” witness specimens were mounted in equivalent location to assess chamber effects. The post-LDT analyses determined composition of deposits as well as the thickness as a function of angle from the beam.

The LDT provided the opportunity to perform periodic plasma probe tests, including Langmuir probe, plasma wave antenna, retarding potential analysis, and even ion/neutral mass spectrometry. Simple model sensors were installed within the LDT test chamber, with major consideration given to minimizing risk to the thruster or the facility. The NSTAR Project would not accept significant technical nor schedule risk from diagnostics in the execution of the LDT.

2.1.3 EMI/EMC—As part of the acceptance process, the flight units underwent characterization of DC magnetic fields, measurement of DC and AC magnetic fields during operation, measurement of AC electric fields during operation, and assessment of plume effect on RF communications. These tests were performed at JPL and at the NASA Glenn Research Center. Included in this test was a spacecraft-level test in which the NSTAR PPU was operated into a resistive load.

DS1 IPS Compatibility Test—The full flight system functional test of the IPS on DS1 was conducted in vacuum following spacecraft thermal vacuum testing. This test also provided an opportunity to characterize plasma and electric/magnetic fields associated with operation of the ion thruster in flight configuration. IDS hardware was integrated and fully operational for the IPS compatibility test. Although the IPS operating time was limited, IDS successfully captured plasma and fields data in this test.

Correlation with flight data provides insight into chamber effects on potential and EMI measurements.

2.2 Modeling Tools

The NSTAR Project has invested significant effort in developing plume models to predict local environments on spacecraft utilizing ion propulsion. These models were used extensively to aid in establishing measurement requirements for the IPS Diagnostics Subsystem. Results from analysis of the IDS flight data will be compared with model predictions to update the modeling tools.

2.2.1 Direct Simulation Techniques—Monte-Carlo particle-in-cell (PIC) codes[11] were developed and executed for electrostatic and electromagnetic characteristics of the NSTAR ion thruster in various environments (free-space, chamber, DS1 spacecraft with simple boundary conditions). The computations simulate plumes due to the NSTAR ion engine using accurate characteristics for engine operations (primary-beam voltage, current, and spatial distributions, propellant utilization, neutralizer conditions, etc.). The generation and propagation of charge-exchange ions are based on a purely physical model that includes particle densities and velocities, accurate collision cross sections, and Coulombic and Lorentz forces. These codes were hosted on massively parallel processors to allow statistically meaningful simulations to be performed in reasonable amounts of time. The characteristics of the charge-exchange ion flow were useful to determine the orientation of the NSTAR diagnostic sensors and to estimate the anticipated magnitudes of charge-exchange currents, plasma densities, and temperatures.

2.2.2 Semi-empirical Modeling—The Environment Work Bench (EWB) modeling tool developed at Maxwell Technologies was employed for estimating system-level interactions associated with the NSTAR ion engine operating on the DS1 spacecraft[12]. The ion engine plume model used in EWB was initially based on laboratory data and PIC code simulations of the NSTAR ion engine. The plume model will be updated with refined modeling and flight data results in order to provide a useful tool for design of future ion propulsion based missions. In the future, systems engineers, mission planners, and principal investigators can utilize this system-level modeling tool on conventional (desktop or laptop) computers.

2.3 IPS Diagnostics Subsystem on DS1

A suite of 12 diagnostic sensors was integrated into the IDS shown in Figure 1. IDS was located adjacent to the NSTAR ion engine on the DS1 spacecraft.

2.3.1 IDS Architecture—IDS consists of two interconnected hardware units: the Diagnostics Sensors Electronics Unit (DSEU) and the Remote Sensors Unit (RSU). The DSEU component of the IDS has considerable heritage to

SAMMES, a modular instrument architecture developed by BMDO[13,14]. A block diagram for the IDS is shown in Figure 2. The IDS is a highly integrated instrument package with a single +28 VDC power and dual MIL-STD-1553 serial communications interface to the DS1 spacecraft. The compact IDS instrumentation package weighed just 8 kg and required 21 W for full operation.

The IDS contains two separate processor elements: the DSEU microprocessor and the fields measurement processor (FMP)[15]. The DSEU microprocessor supports the communications interface with DS1, controls serial communications with the FMP, and digitizes and controls the sensors within the RSU. The IDS operates as a remote terminal on the DS1 MIL-STD-1553 serial bus. Telemetry from the RSU sensors is collected on 2-second intervals and placed in selected 1553 subaddresses for transmission to DS1. Configuration messages are transmitted to the DSEU to select active sensors within the RSU and FMP and to establish sweep ranges and gains for these sensors. Configuration messages to the FMP are passed through the DSEU to the FMP directly. The DSEU polls the FMP for data at half-second intervals. In the typical FMP “scan” mode operation, a block of sensor data is transmitted at 16-second intervals. Occasionally, the FMP will transmit

1-second waveforms sampled at 20 kHz from the plasma wave and search sensors and 20 Hz from the flux-gate magnetometers. These “burst” events can be commanded or initiated via internal triggering within the FMP.



Figure 1. IPS Diagnostics Subsystem Hardware

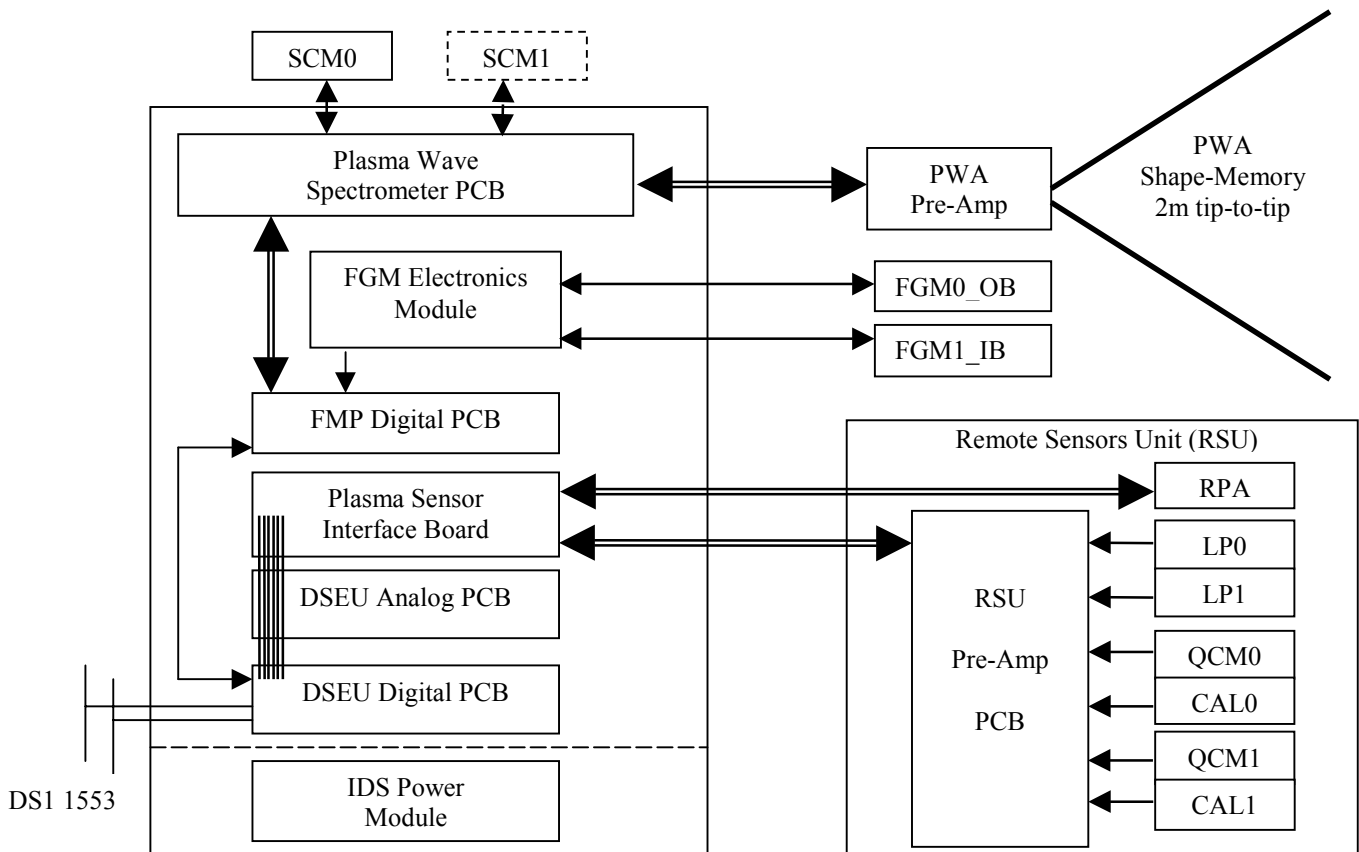


Figure 2. IPS Diagnostics Subsystem Block Diagram

The highly integrated design approach greatly simplified spacecraft interface design, integration, and mission operations for IDS. The interface control document was developed in a very straightforward process with the greatest issue involving positioning of the IDS hardware to avoid interferences with the launch vehicle upper stage. Mechanical and electrical integration of IDS was accomplished within 2 hours. Mode changes during mission operations was accomplished by transmitting single 1553 messages (64-byte) at the desired time. These commands were readily integrated into operations sequences.

2.3.2 Contamination Monitors—Two QCM and calorimeter pairs were integrated in the RSU to characterize mass-deposition rates and contamination effects on surface thermo-optical properties. One pair of sensors is oriented to a direct line-of-sight view of the NSTAR ion engine. The DS1 propulsion module shadows the other contamination monitor pair from direct view of the NSTAR engine. The QCMs detect mass variations on the sensor surface via the induced frequency change in the oscillating-quartz crystal sensor. The calorimeters provide indirect knowledge of solar absorptance and hemispherical emittance by temperature measurement of the thermally isolated sensor surface.

Each QCM (Mark 16 flight sensors procured from QCM Research, Laguna Beach, California) provides very high sensitivity measurement ($<10 \text{ ng/cm}^2$) of mass accumulation on the sensor[16]. The long-term drift of the QCM should not exceed 50 ng/cm^2 per month, which corresponds to a minimum detectable molybdenum deposit rate of one monolayer per year. Temperature changes and solar illumination of the sense crystal affect QCM response. For substantial mass accumulation, the temperature and solar illumination effects on the QCM measurement are minor.

The calorimeters[17] can determine solar absorptance changes to better than 0.01 and emissivity changes to better than 0.01. The calorimeters use the Sun as a stimulus for determination of solar absorptance. The calorimeters include a controlled heater to permit measurement of the hemispherical emissivity of the surface. Spacecraft surfaces in the field of view of the calorimeter complicate data analysis because of the uncertain heat loads that these surfaces provide to the sensor surface.

The data from the QCM and calorimeter sensors are reduced, analyzed, and correlated with NSTAR ion engine operations. The QCM with direct line-of-sight to the NSTAR ion engine was expected to accumulate readily detectable amounts of sputtered molybdenum. Pre-flight estimates indicated the deposition rate on non-line-of-sight surfaces near the thruster from ionized molybdenum will be very low.

2.3.3 Charge Exchange Plasma Sensors—IDS includes a retarding potential analyzer (RPA) and two Langmuir probes to characterize the charge-exchange plasma produced by the NSTAR thruster. The RPA measures the charge-exchange ion energy distribution over the range of 0 to +100 eV near the thruster exit plane. The RPA sensor axis is co-aligned with the predicted charge-exchange ion flow direction expected at the RPA location. Langmuir probes are used to measure the electron temperature and the density of the plasma near the NSTAR thruster.

The RPA used in the IDS was salvaged from the Ion Auxiliary Propulsion System on P80-1 (Teal Ruby). These units were fabricated and qualified for flight by Hughes Electronics in 1978[18]. Extensive performance and calibration data have been obtained for the flight units. The RPA is a four-grid design with screen and suppressor grids operated at -12 VDC and the bias-grid voltage adjustable from 0 to +100 VDC. An RPA sweep consists of sixteen voltage steps, within the 0 to +100-V range, with a minimum step size of 0.39 V. The currents for the biasing voltages applied to the grids within the RPA will be monitored and included in the RPA telemetry stream to permit detailed analysis of the charge-exchange plasma near the engine. The ion collector includes a pre-amplifier with selectable full-scale detection ranges from 10^{-9} to 10^{-3} A . In the case of the IDS, the full-scale selectable gains for the RPA are 10 nA, 1 μA , 10 μA , and 100 μA . The entrance aperture to the RPA is 5 cm in diameter.

Two Langmuir probe sensors were included in the IDS: LP0, a spherical probe (4-cm diameter), and LP1, planar ring (50 cm^2) on a conductive MLI blanket. The probes were independently biased (swept or constant voltage range) from -7 VDC to +11 VDC. Langmuir probe current measurement range extends from $-500 \mu\text{A}$ to +40 mA. The Langmuir probe-support circuitry was designed and fabricated by Sentran Corporation, Goleta, California.

2.3.4 Fields Measurements—The baseline diagnostic sensor package for NSTAR did not include electric or magnetic field sensors. The presence of high-density-field permanent magnets in the NSTAR thruster warranted investigation as to the long-term stability of these fields. An augmentation to the IDS for fields measurement was made possible by the participation of Technical University of Braunschweig (TUB), TRW, and the Jet Propulsion Laboratory (JPL) Integrated Space Physics Instrument team. Measurement of the DC magnetic fields was performed by two three-axis flux-gate magnetometers, each mounted on a short boom extending from the spacecraft. Measurements of low frequency AC magnetic fields (10 Hz to 50 kHz) characterize the electromagnetic interference (EMI) produced by the engine. In addition, it was possible that electromagnetic waves induced by plasma stream instabilities within the plume and by plume interactions with

the solar wind could be detected. Two search coil magnetometer sensors were mounted to the boom to measure these electromagnetic waves. The plasma wave environment produced by the thruster was expected to be similar for emissions that have been measured for the Space Station Plasma Contractor hollow cathode source. The emissions are very broadband, from essentially DC to about 10 MHz, with interference with spacecraft operations highly improbable. A 2-m tip-to-tip dipole antenna with adjustable gain pre-amplifier on the boom measured the plasma wave environment over the frequency range of 10 Hz to 30 MHz.

2.3.5 Flux-Gate Magnetometers—Two sensitive, three-axis flux-gate magnetometers designed and built by TUB were mounted on the boom near the NSTAR ion engine. The inboard magnetometer is located in a high-density-field region (9,000 nT). The outboard magnetometer was positioned to place the sensor in a somewhat weaker field (less than 3,000 nT). The magnetometer sensitivity is better than 1 nT with $\pm 25,000$ -nT full-scale range. The maximum sampling rate of the flux-gate magnetometers is 20 Hz.

2.3.6 Search Coil Magnetometers—Two single-axis search coil magnetometers were mounted on the boom. One search coil is a new technology miniaturized sensor developed in the JPL MicroDevices Laboratory that uses a field rebalance technique for measurement. The second search coil was a build-to-print of the Orbiting Geophysical Observatory (OGO-6) single-axis sensor manufactured by Space Instruments, Inc., Irvine, California. The second search coil sensor was apparently damaged at the launch site by large AC fields and was inoperable during the DS1 mission. Flight measurements were performed with a measurement bandwidth over 10 Hz to 50 kHz. The full-scale range at 200 Hz is 100 nT with a resolution of 1 pT. The AC magnetic fields were characterized as a discrete power spectrum with four measurement intervals per decade. The transient waveform for “events” was also captured with a sampling rate of 20 kHz for 20-msec windows. The transient recorder utilized a circular buffer with a threshold trigger to capture events. The threshold parameters are capable of being updated via ground command through the DS1 spacecraft.

2.3.7 Plasma Wave Antenna—A simply deployed dipole plasma wave antenna (PWA) with adjustable-gain pre-amplifier was mounted onto the boom. The PWA is a pair of low-mass Ni-Ti shape-memory alloy (SMA) metallic strips with a tip-to-tip separation of 2 m. The PWA deployment occurs upon exposure of the stowed SMA coiled ribbon to the Sun. Within 2 hours, the PWA antenna slowly extends to its deployed position. The PWA is connected to a low-noise preamplifier co-located on the boom that was designed and built by TRW, Redondo Beach, California. The amplified PWA output is processed by the plasma wave spectrometer (PWS), also designed and built by TRW, to

provide a spectrum analysis in a low-frequency domain of 10 Hz to 100 kHz and a high-frequency domain of 100 kHz to 30 MHz. The low-frequency domain is characterized by a voltage-swept band pass filter with a minimum of four measurements per decade with an amplitude range of 100 μ V/m to 1000 mV/m. The high-frequency domain is characterized with a minimum of four measurements per decade with the same amplitude range as the low-frequency domain. Transient waveform measurements will be performed at a 20-kHz sampling rate with a 20-msec circular buffer. Threshold parameters for trigger and downlink of transient waveforms are capable of being uploaded from the DS1 spacecraft.

3.0 CHARGE-EXCHANGE PLASMA

The electrostatic potential of the DS1 spacecraft with respect to the ambient space plasma is determined by current balance[11,12]. Charge-exchange plasma from the NSTAR ion engine drives the current balance on the spacecraft. The amount of charge-exchange plasma produced by the NSTAR ion engine varies with the engine operating conditions. Electric probes, such as the IDS Langmuir probes, are capable of sinking large amounts of current. The perturbations by the IDS Langmuir probes can substantially effect the DS1 spacecraft potential. The following sections describe the current understanding of the spacecraft potential, charge-exchange ion variation with engine thrust level, and effects produced by the IDS Langmuir Probes.

3.1 DS1 Chassis Potential Without Langmuir Probe Bias Voltage

At equilibrium, spacecraft chassis ground potential is determined by the fact that the net current to the exposed conductors (thruster-mask ring around engine, Langmuir probe with black Kapton on RPA box, see Figure 3) is zero. In the interplanetary space plasma environment, the Debye length is much larger than the spacecraft dimensions. The plasma density from the ion engine is many orders of magnitude larger than the space plasma. The following analysis assumes the charge-exchange ions collected are orbit limited, which may be a questionable assumption.

The surface area of the conductors and the plasma density at the conductor determines the relative contribution of current collection. The surface area of the thruster mask ring is 0.085 m² (inner radius 0.15 m, outer radius 0.2225 m). The surface area of the ring Langmuir probe is 0.0050 m², without considering the black Kapton outer blanket of the RSU. It will be shown later that the effective collection area is approximately double when the conductive black Kapton is included. Plasma density estimates were computed via PIC code simulation[11,12] (Figure 4). The plasma density at the thruster mask ring is in excess of 10¹⁴ m⁻³; the density at the RPA Langmuir probe is less than 10¹³ m⁻³. Current

balance at the thruster-mask ring, therefore, dominates the chassis potential with no bias on the Langmuir probe.




Figure 3. Major Contributors to Current Balance on DS1




Figure 4. Computed Ion Density Contours for NSTAR Ion Engine at Full Power (dimensional scale is meters)

The equality of electron and ion current to the thruster-mask ring determines the relationship of the chassis potential (ϕ) to the plasma-electron temperature (θ).

$$I_e = en\sqrt{\frac{e\theta}{2\pi m_e}} \exp(\phi/\theta) = I_i = en\sqrt{\frac{e\theta}{2\pi m_i}} (1 - \phi/\theta)$$

Rearranging and simplifying gives:

$$\exp(\phi/\theta) = \sqrt{\frac{m_e}{m_i}} \left(1 - \frac{\phi}{\theta}\right)$$

This equation is solved numerically for ϕ/θ and is satisfied with a value of -4.5 . Chassis potential is related to the plasma potential (ϕ) by:

$$\phi = -4.5\theta + \phi$$

PIC computations performed prior to flight (Figure 5) predict that the plasma potential (ϕ) near the thruster mask ring is approximately 1.25 V and the electron temperature confirmed by measurement, $\theta = 1.8$ eV. As a result, the chassis potential for DS1 during NSTAR operations is estimated at -6.75 V.




Figure 5. Self-consistent Potential Computed for NSTAR Thruster Operating at Full Power (dimensional scale in meters)

As a consistency check for the estimated chassis potential, the variation of the in-flight measured voltage of the IPS internal ground (neutralizer common) with the Langmuir probe bias is compared to ion energies measured by the IDS RPA. During the second IPS performance acceptance test in flight (IAT2) conducted on May 28, 1999, the IDS Langmuir probe sensors were held at four voltage levels (-7 V, -1 V, $+5$ V, and $+11$ V, with respect to chassis ground) for a few minutes at each IPS thrust level. The effect on the IPS internal ground is shown in Figure 6.

RPA sweeps obtained at each Langmuir probe voltage level are shown in Figures 7a through 7d. Note the increasing mean ion energy with increasing Langmuir probe bias. The important results from Figures 6 and 7 are summarized in Table 1.

Note that when the Langmuir probe bias is at $+11$ V, the neutralizer common is 1.75 V higher than when the Langmuir probe bias is near ground. This implies that the

DS1 chassis ground is driven -1.75 V due to electron collection by the Langmuir probe at $+11$ V. During IPS operations, the Langmuir probe is able to drive the DS1 chassis potential from -6.75 V (no bias) to -8.50 V (bias = $+11$ V).

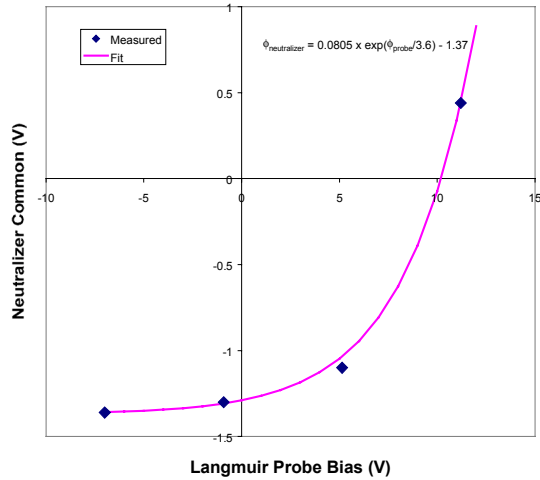


Figure 6. Variation of IPS Neutralizer Common with IDS Langmuir Probe Bias Voltage in IAT2




Figure 7a. RPA Sweep at -7 -V Langmuir Probe Bias




Figure 7b. RPA Sweep at -1 -V Langmuir Probe Bias




Figure 7c. RPA Sweep at $+5$ -V Langmuir Probe Bias




Figure 7d. RPA Sweep at $+11$ -V Langmuir Probe Bias

Table 1. Effect of Langmuir Probe Bias on Ion Energy and Neutralizer Common

Langmuir Probe Bias (V)	Ion Energy (eV)	IPS Neutralizer Common (V)
-7	12.9	-1.35
-1	13.0	-1.3
$+5$	12.9	-1.1
$+11$	14.6	$+0.45$

The estimated net current collection at the thruster mask ring with the Langmuir probe bias at $+11$ V is 2.75 mA (assuming $\phi/\theta = 5.5$):

$$I_i = Aen\sqrt{\frac{e\theta}{2\pi m_i}}(1 + \phi/\theta) = 4.0 \text{ mA}$$

$$I_e = Aen\sqrt{\frac{e\theta}{2\pi m_e}}\exp(-\phi/\theta) = -1.25 \text{ mA}$$

The measured Langmuir probe current is about 2 mA when biased at +11 V; this is in fairly good agreement with the above-calculated net ion collection at the thruster mask ring.

The voltage of the Langmuir probe with respect to the local plasma potential as a function of Langmuir probe bias is shown in Figure 8 below. Note that the Langmuir probe will not collect substantial electron current from the plasma until the probe bias has reached approximately +7.5 V.




Figure 8. Estimated Langmuir Probe Potential Versus Probe Bias Voltage

3.2 RPA Current as a Function of IPS Mission Level

Charge-exchange ion production is expected to depend upon IPS operating conditions, since charge-exchange ions are formed by the interaction of beam ions and neutral xenon escaping from the IPS discharge chamber. The expected charge-exchange ion current at the IDS RPA has been calculated for the NSTAR ion engine operating conditions reported in “Engine Table Q.” The calculations use velocity-dependent resonant charge-exchange cross sections computed from the formula provided by Sakabe and Izawa[19]. A transmission factor of 0.27 for the four-grid RPA is based on an individual grid transparency of 0.72. The results of the calculation, with measured RPA currents from IAT2, are illustrated in Figure 9.

A curious feature in the data shown in Figure 9 is the larger ion current observed at IPS mission level 6 than at higher mission levels (up to 34). In fact, the RPA ion current is 40% higher for mission level 6 than mission level 13. The reason for the enhanced charge-exchange ion production at mission level 6 is the higher relative xenon flow rate in the discharge chamber than the conditions for mission level 13. The excess or residual xenon escaping from the discharge chamber accounts for the higher charge-exchange ion production. Table 2 compares the operating conditions for mission levels 6 and 13.




Figure 9. Computed RPA Ion Current as a Function of IPS Mission Levels (measured currents from IAT2 also shown)

Table 2. Relevant IPS Operating Conditions for Mission Levels 6 and 13

Quantity (units)	ML6	ML13
Total chamber flow (sccm)	8.450	8.290
Total chamber flow (Amps equiv.)	0.606	0.594
Beam current (Amps)	0.509	0.529
Residual Xe flow (Amps equiv.)	0.097	0.065
Beam* residual Xe (Amps ²)	0.049	0.035

The ratio of the product of the beam current and residual Xe for mission level 6 versus 13 is 1.4. This ratio is in good agreement with the measured charge-exchange current ratios for mission levels 6 and 13.

3.3 Variation of RPA Current with Ring Langmuir Probe Bias

The placement of the Langmuir probe at the entrance to the RPA causes the probe bias voltage to effect the path of ions approaching the RPA. Figure 10 shows the variation of RPA current with Langmuir probe bias.

The variation in the ion current is attributed to a focusing effect due to Langmuir probe bias. The potential contour and trajectories for ions approaching the RPA with surrounding Langmuir probe at +11 V were computed (see Figure 11). The potential is expressed in terms of the local plasma potential; hence, the entrance to the RPA (chassis ground) is approximately -9.5 V and the ring Langmuir probe bias is +2.5 V. The plasma conditions for this calculation assumes a density of 10^{12} m^{-3} and a temperature of 1.8 eV. The trajectories for 5 eV xenon ions are shown to illustrate the focusing effect of the Langmuir probe.

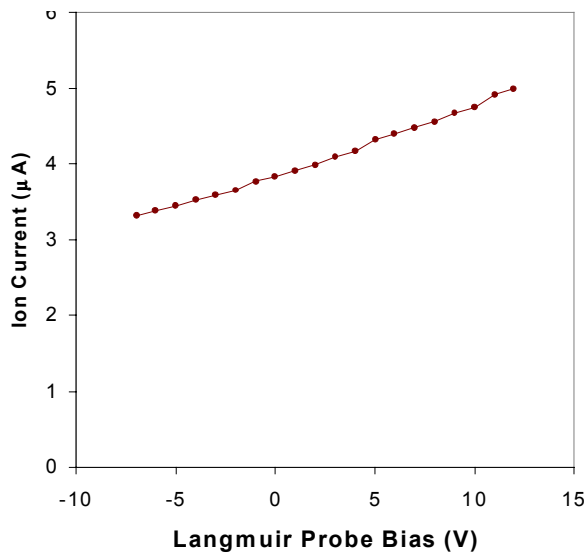


Figure 10. Effect of Langmuir Probe Bias on RPA Current

the Langmuir probe current is easily removed to allow temperature determination as shown in Figure 13.

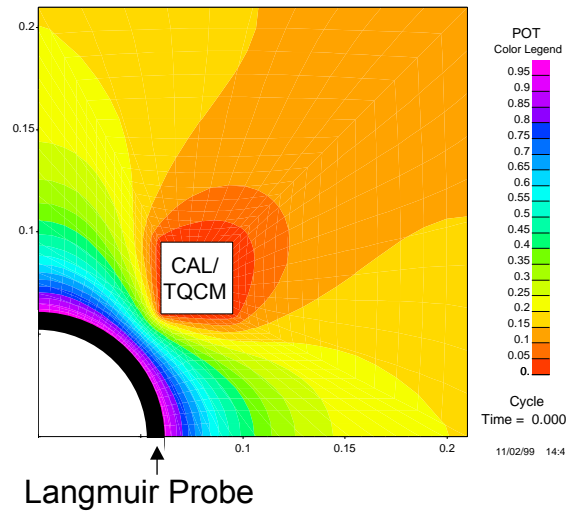


Figure 12. Expansion of Langmuir Probe Bias onto RSU Thermal Blanket

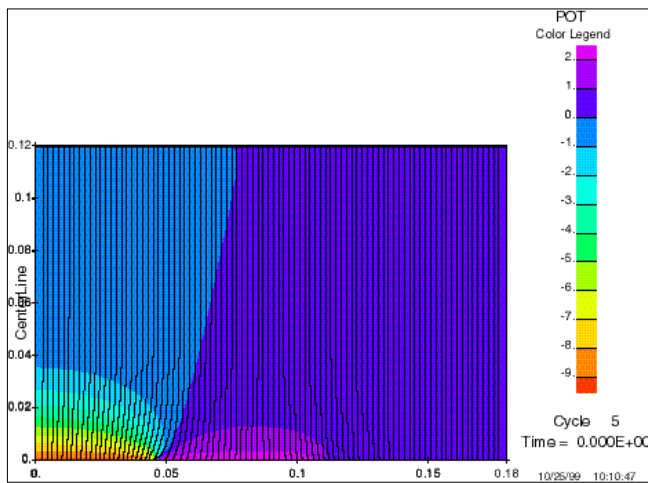


Figure 11. Ion Focusing by RPA Langmuir Probe

3.4 Expansion of Langmuir Probe Bias Potential onto Black Kapton

The RPA Langmuir probe is in direct contact with the RSU thermal blanket. The outer layer of this thermal blanket is fabricated from conductive, carbon-filled Kapton film. This black Kapton material provides a resistive path from the Langmuir probe to the spacecraft chassis (ground). The effect of the blanket surface on effective probe size was calculated. The expansion of the probe bias onto the blanket surface is shown in Figure 12. The conductive blanket effectively doubles the size of the RPA Langmuir probe.

Over course of the mission, the effective resistance from the RPA Langmuir probe to the spacecraft chassis decreased due to deposition of molybdenum sputtered from the ion engine grid. At the time of IAT2, the effective resistivity of the film was 17 kΩ per square. The resistive component to




Figure 13. Langmuir Probe Least Squares Fit ($\theta = 1.8$ eV)

3.5 Expansion of Charge-Exchange Plasma Around DS1

This subsection describes results obtained by computer modeling of the expanding charge-exchange ion cloud around the DS1 spacecraft[11]. The charge-exchange plasma produced near the IPS thruster exit was easily detected by the PEPE instrument, located at the opposite end of the DS1 spacecraft. A particle-in-cell (PIC) computer model was constructed to simulate the charge-exchange ion plasma environment surrounding the DS1 spacecraft, especially in the backflow region (upstream of the thruster plume). The physics of the charge-exchange plasma backflow is similar to that of plasma expanding into a vacuum or wake. The expansion fan is a pre-sheath for the spacecraft, which turns the trajectories of the ions into the upstream direction until they enter the sheath of the spacecraft.

The model is a full three-dimensional PIC simulation in which the DS1 spacecraft and solar-array elements are included, as shown in Figure 14. For efficiency in computation, the $43 \times 43 \times 71$ grid cells used were uniform in size ($d \approx 6$ cm). Approximately 5 million particles were simulated in steady-state conditions. The electrons were included as a fluid with a Boltzmann distribution based on the electron temperature measured by the IDS (approximately 2 eV). Inputs to the simulation include the beam ion density, neutral density, and charge-exchange ion-production rate near the thruster exit. Figure 15 illustrates the beam-ion density, neutral-xenon density, and charge-exchange production rate downstream of the DS1 ion engine. The beam and neutral-density plot contours are normalized to the peak densities at the engine exit plane at uniform intervals of 0.05. The charge-exchange ion production rate is also normalized to the peak rate at the engine exit; however, the contours are given on intervals of 0.005, 0.01, 0.05, 0.1, 0.5, and 1.0.




Figure 14. Model Geometry for PIC Simulation

The results of the PIC simulation are illustrated in Figure 16. Figure 16 provides plots of (a) the plasma-electric potential, (b) normalized electric-field vectors, (c) charge-exchange density and (d) charge-exchange ion-flow field vectors around the DS1 spacecraft.

The peak potential is 19 V with respect to spacecraft ground and is shown in Figure 16a at 1-V intervals. The direction of the electric-field gradients, illustrated in Figure 16b, clearly shows how the charge-exchange ions are accelerated into the backflow region. The PIC simulation estimates the charge-exchange ion density to be approximately 10^6 cm^{-3} near the IDS, decreasing to 10^4 cm^{-3} near PEPE. Figure 16c shows the charge-exchange ion density distribution around DS1 and the direction of flow of the charge-exchange ions. The charge-exchange density near DS1 during IPS operations is at least three orders of magnitude greater than the ambient solar-wind plasma density.

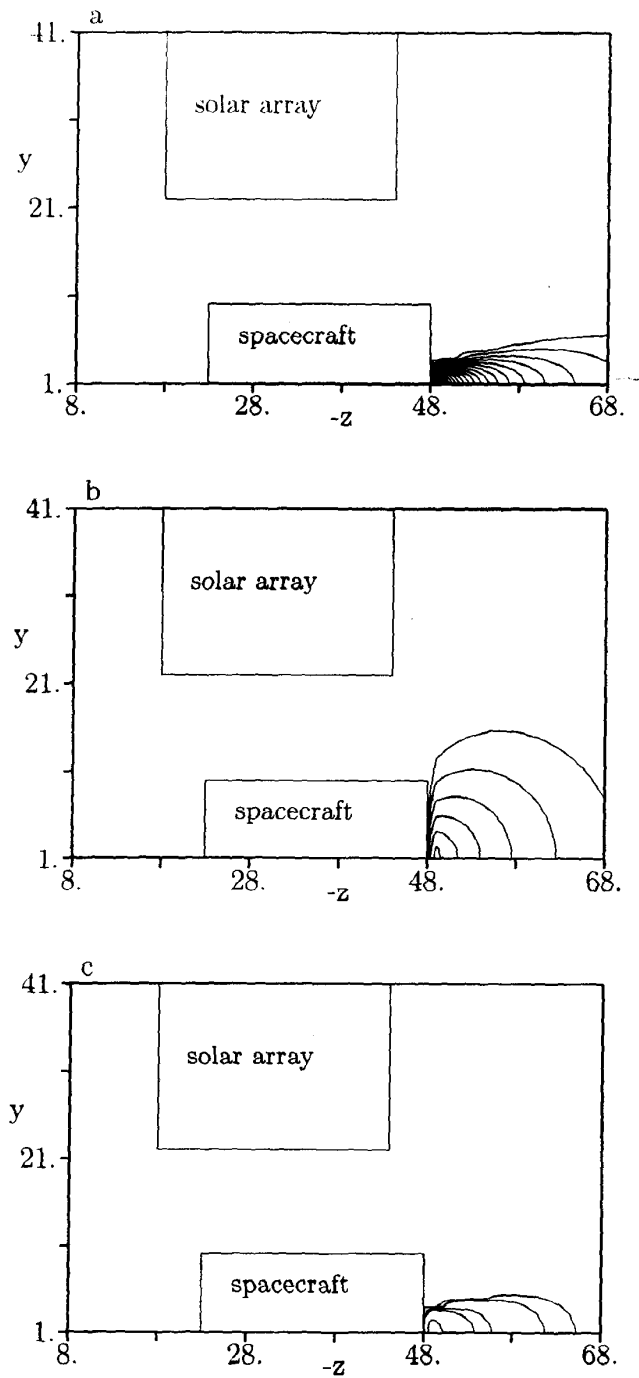


Figure 15. Contour Plots for: (a) Beam Ion Density, (b) Neutral Xenon Density, and (c) Charge-Exchange Ion Production Rates




Figure 16. Results of the DS1 PIC Simulation: (a) Electric Field Potential, (b) Electric Field Direction, (c) Charge-Exchange Ion Densities, and (d) Charge-Exchange Ion Flow Directions

4.0 CONTAMINATION ASSESSMENT

The NSTAR Diagnostics Element has produced useful data regarding the IPS contamination environment. The 8000-hour Life Demonstration Test (LDT) afforded the opportunity to measure the thickness and composition of deposits accumulated from extended operation of an ion engine. The IDS flight-contamination monitors functioned properly and provided high-quality data regarding deposition rates as a function on IPS thrust level.

4.1 Ground Test Contamination Results

The NSTAR 8000-hour LDT was performed at JPL to validate the long life of the NSTAR thruster. A fundamental purpose of the LDT was to assess the effects of extended operation on the engine, especially the grids and cathodes. The grid wear-out mechanism is loss-of-grid material (molybdenum) via sputtering by charge-exchange xenon ions[8]. A significant portion of the sputtered grid material is emitted outward from the engine. In the ground-test

environment, chamber effects can strongly effect the results for contamination-witness specimens. The LDT chamber walls were lined with graphite plates to reduce the amount of material sputtered back onto the engine[10]. The contamination monitors described below were designed to minimize effects from material sputtered from the chamber walls.

4.1.1 LDT Contamination Monitors—A series of collimated 1-inch diameter fused silica windows were mounted on a curved support beam 46 inches (1.2 m) from the engine (see Figure 17). The witness monitors were placed at angles from 40° to 110° from the thrust axis, at 10° intervals. To avoid collection of sputtered chamber material, the witness windows were place in long (25 cm) tubes lined with tantalum foil. At the entrance of the tube, a collimating aperture was positioned to limit the witness field-of-view to the ion-engine grid. Shadow wires (tungsten) were positioned on the windows to facilitate profilometry measurements.




Figure 17. Geometry of Collimated Contamination Monitors for the NSTAR LDT (drawing is not to scale).

Subsequent to the completion of the 8000-hour LDT, the contamination witnesses were removed from the collimation tubes for analyses. Visual inspection of the windows clearly showed a metallic film for witnesses located between 60° and 110° from the thruster beam axis. The metal films appeared hazy or crazed, not highly specular as a uniform flat coating would appear. Attempts to measure film thickness using a profilometer were not very successful. Examination via a scanning electron microscopy (SEM) revealed that the metallic films were wrinkled, presumably due to stresses in the coating and poor adhesion to the substrate. In the regions where the profilometer stylus had contacted the film, the film was scraped from the substrate surface. It was possible to determine the thickness of the coatings in these disturbed areas with SEM imaging. Figure 18 shows LDT deposition at 10° increments between 60° and 110° from thrust axis. The uncertainty in the thickness measurements is on the order of 10%. Currently, there is no firm explanation for the apparent enhanced deposition observed at the 80° position. It is conceivable that erosion from the edges of grid holes could lead to a complex angular deposition distribution[20]. X-ray dispersive spectroscopy (XDS) of the metal films revealed their composition to be molybdenum metal (no evidence of tantalum contamination), with a significant amount of xenon detected. The source of the xenon is either from background xenon within the LDT vacuum chamber or, possibly, impingement of low-energy (<10 eV) charge-exchange xenon ions. The witness monitors at 50° and 40° were found to be eroded 1.7 and 7.7 μm. Energetic xenon ion sputtering causes the erosion of these witness monitors. The witness monitors at larger angles may also be impinged by xenon ions capable of sputtering material, but not at a sufficient flux to prevent deposition of sputtered molybdenum.




Figure 18. Molybdenum Accumulation for NSTAR LDT Witness Monitors Versus Angle from Thruster Axis

4.1.2 Flight Correlation—On DS1, the line-of-sight contamination monitors are located 75 cm from the thruster centerline, 85° off thrust axis. Even though the grid is an extended source, the deposition thickness is roughly inversely proportional to square of distance from grids. A contamination witness located at the DS1 QCM0 position would have accumulated approximately 1300 Å during LDT. Rates of contamination accumulation during IPS thrusting can be conveniently expressed in terms of Angstroms of molybdenum per 1000 hour (1 khr) of operation. The expected deposition rate for an LDT witness monitor in the equivalent position of the line-of-sight IDS contamination monitor is 160 Å/khr.

4.2 Flight Contamination Results

The quartz crystal microbalance (QCM) sensors mounted in the IDS Remote Sensors Unit have produced useful data for assessing the contamination environments on DS1. The IDS QCM sensors are 10-MHz fundamental-frequency devices; hence, the frequency-to-area-mass-density conversion is 4.43 ng/cm²-Hz[16]. QCM beat frequencies are sensitive to changes in temperature and solar illumination of the sense crystal. To extract low-level contamination information, QCM data often must be corrected for temperature and solar illumination. The magnitude of the IPS-induced contamination for the line-of-sight (QCM0) sensor is such that these corrections are not necessary. The non-line-of-sight (QCM1) sensor, though, had significantly less accumulation; therefore, its data should be corrected prior to precise quantitative interpretation. The data, as presented in this report, have not been corrected for sense crystal temperature or solar illumination. The preliminary results are discussed chronologically in this section.

4.2.1 Launch Operations—The final pre-flight functional test of the IDS prior to launch was conducted on DOY 293-1998. Data from the QCMs provide the pre-launch baseline for assessing launch-phase contamination in the vicinity of the DS1 to launch-vehicle interface. The pre-launch readings were obtained at 16 °C and are 2475 Hz and 2085 Hz for QCM0 and QCM1, respectively. Following launch, DS1 was oriented with the Sun vector aligned with the spacecraft X-axis. In this orientation, QCM0 is illuminated with a Sun angle of approximately 46°, whereas QCM1 is in the shadow of the DS1 propulsion module. The IDS was not activated until 1998-298 at 2201 hour (approximately 34 hours after launch). The initialization of IDS included a special activity (“DFrost”) intended to bake-off volatile contamination from the QCMs and calorimeters by heating the sensors to +75 °C. Very little change (<50 Hz) was observed in the beat frequency of either QCM as a result of the initial post-launch DFrost.

The frequencies and temperatures for QCM0 and QCM1 just prior to DFrost were 2260 Hz (at +30 °C) and 2272 Hz (at +16 °C) respectively. Since QCM0 was exposed to the sun after launch, it is suspected that most of the contaminants accumulated on it were evaporated prior to IDS initialization. The beat frequency for QCM1 increased by 187 Hz from pre-launch to IDS initialization, yielding an estimated 0.8 μg/cm² (80 Å) accumulation for launch-phase contamination. This accumulation was not affected by the DFrost activity, but was removed when DS1 rotated to expose the NSTAR ion engine to the Sun (NSTAR Decontamination Maneuver). Figure 19 shows the early mission response of QCM1 to the DS1 orientation with respect to the Sun shown in Figure 20. Note the substantial frequency and temperature changes near DOY 304-1998 associated with the NSTAR Decontamination Maneuver. There is an additional turn on DOY 305-1998 that further effects the QCM1 frequency and temperature. On DOY 306-1998, DS1 returned to the nominal Sun on X-axis orientation. Using the frequency reading at this time, it appears that about a 165-Hz decrease occurred as a result of this solar-stimulated bakeout. Based on this interpretation of QCM1 data, it appears that the RSU surfaces were contaminated during the DS1 launch phase with approximately 80 Å of low-volatility organic material, most of which was removed upon exposure to the Sun.

4.2.2 IPS Operations—The QCM data for the IPS operations of the first year of flight for DS1 are illustrated in Figures 21a through 21g. The figures are arranged so the response of the line-of-sight (QCM0) and non-line-of-sight (QCM1) sensors can be compared side-by-side. The four pairs of figures represent time intervals during which IPS operations of substantial duration occurred. Data for minor

thrusting events, such as the brief “S-Peak” test on DOY 022-199 and the trajectory correction maneuvers prior to the Asteroid Braille encounter, do not show significant accumulations on either QCM. Similarly, data for the long, non-thrusting intervals are not shown because no accumulation occurred on either QCM in these periods.




Figure 19. QCM1 (Non-line-of-sight) Early Mission Response




Figure 20. Early Mission DS1 Sun Orientation

The first period of extended IPS operations occurred from DOY 328-1998 to DOY 005-1999. The line-of-sight sensor (QCM0) response is shown in Figures 21a and 21c; the shadowed-sensor (QCM1) response is seen in Figures 21b and 21d. The initial IPS operations consisted of 10 days thrusting with the thrust vector essentially Earth-pointed. During these initial operations, the NSTAR engine was first operated at low-thrust (mission levels 6 to 27) for five days. During this period, QCM0 frequency increased by 123 Hz, while QCM1 increased by 25 Hz. To determine the deposition rates for mission level 27 (ML27), least squares fits of the frequency data for the 117-hour interval starting on DOY 329-1998 and ending on DOY 334 were performed. The resulting slope in units of Hz/day was converted to Å(Mo)/kHr by multiplying by 1.804.

Deep Space 1 Technology Validation Report—*Ion Propulsion Subsystem Environmental Effects on Deep Space 1: Initial Results from the IPS Diagnostics Subsystem*




Figure 21a. QCM0 Data for 1998-317 through 1998-365




Figure 21b. QCM1 Data for 1998-317 through 1998-365




Figure 21c. QCM0 Data for 1999-001 through 1999-012




Figure 21d. QCM1 Data for 1999-001 through 1999-012




Figure 21e. QCM0 Data for 1999-074 through 1999-124

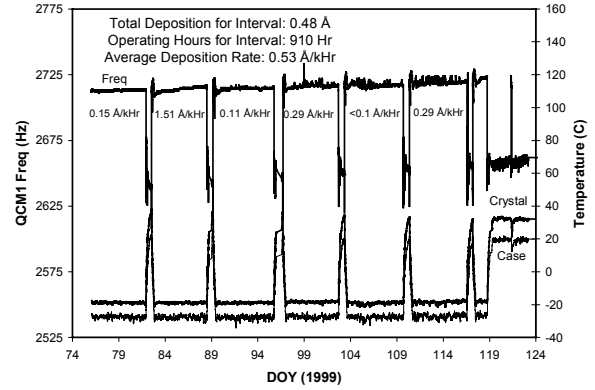


Figure 21f. QCM1 Data for 1999-074 through 1999-124




Figure 21g. QCM0 Data for 1999-214 through 1999-300




Figure 21h. QCM1 Data for 1999-214 through 1999-300

QCM0 data for the remaining thrusting of the initial period shows some interesting features. On DOY 338-1998, DS1 performed a turn to orient the thrust vector from Earth-pointed to the desired mission trajectory thrust attitude. In this turn, the solar illumination of the thruster increased from grazing (80° off of the thrust axis) to about 0.77 suns (40° off of the thrust axis). Note that the molybdenum deposition rate for QCM0 at ML83 prior to DOY 338-1998 was 88 Å/kHr, whereas after the turn, the deposition rate increased to 197 Å/kHr. The accumulation rate of QCM1 almost doubles after the turn. It is not yet known whether this rate change is due to thermal effects on the NSTAR ion engine grids. There have been no reports of change in mass sensitivity with varying Sun angle on QCMs; therefore, it is unlikely that the rate change is an instrument artifact.

Following the turn to thrust attitude, DS1 continued thrusting until DOY 342-1998. Other technology activities, including initial turn-on of the Plasma Experiment for Planetary Exploration (PEPE) instrument were performed. On DOY 346-1998, IPS was restarted at low-thrust level (ML6) to assess the effects on the PEPE instrument. The on-board sequence raised the IPS thrust level to ML85 after 15 minutes. The available power for IPS thrusting was overestimated, resulting in a DS1 “safe-mode” transition. IPS thrusting resumed on DOY 348-1998 after DS1 spent two days in safe mode.

The first IPS thrust segment ended with two weeks of essentially continuous thrusting, with the thrust levels gradually decreasing from ML78 on DOY 352-1998 to ML 72 on DOY 005-1999. During this interval, the DS1 on-board navigation software would update the thrust vector and level at 12-hour intervals. The IPS thruster was turned off at 1600 hours on DOY 005-1999. The deposition on QCM0 steadily increased over this interval, except for a brief interval on DOY 356-1998 where DS1 re-oriented to place the Sun on the X-axis for approximately 3 hours. QCM1 also showed consistent frequency increase, although at an order-of-magnitude lower than that for QCM0. The thrust segment continued into early 1999, with steady accumulation by both QCMs witnessed in Figures 21c and 21d. Subsequent to engine turn-off on DOY 005-1999, DS1 performed maneuvers to characterize stray-light into the MICAS imager. The effect of minor Sun-angle changes caused by attitude control system dead-banding on QCM1 (100 Hz oscillation) is quite evident for DOY 009-1999 through DOY 012-1999.

The next major IPS thrust interval was the C1A and C1B activities performed from DOY 075-1999 until DOY 117-1999. This thrusting was performed with weekly optical navigation (OpNav) activities and high-rate telemetry downlink intervals. The thrusting duty cycle was typically greater than 90% during this interval. The OpNav/downlink events are readily identified in Figures 21d and 21e by 100-

Hz frequency dips in both QCM0 and QCM1 as well as 60 °C temperature increases for both sensors. The deposition rates for QCM0 are labeled in Figure 21d with time-averaged thruster mission levels for each thrust segment. During the C1A and C1B activities, the on-board navigator commanded the desired IPS mission level. The DS1 power management software would monitor battery state-of-charge and perform thrust reduction as required. For this period, the non-line-of-sight sensor (QCM1) accumulated only about 1% of the amount of molybdenum collected by QCM0. This value is consistent with pre-flight estimates for production and collection of ionized molybdenum from the thruster plume.

Subsequent to the Asteroid Braille encounter on DOY 210-1999, IPS operated for an interval of almost 12 weeks. As the DS1-Sun distance decreased, the mission level gradually increased during the C2A and C2B segments. The deposition rates of both QCMs also increased during this period, as seen in Figures 21g and 21h. The brief, periodic spikes in the QCM frequency data occur at each of the weekly OpNav and downlink sessions, again caused by Sun-angle changes. The accumulation of molybdenum on the shadowed QCM is about 5% of that witnessed by the line-of-sight sensor.

The four thrusting segments shown in Figures 21a through 21h account for more than 95% of the IPS operating time for the first year of the mission. Of the 250 Å of molybdenum collected on the line-of-sight QCM in the first year of operation, almost 95% of the accumulation are shown in these figures. The shadowed QCM collected the equivalent mass of a 25-Å thick deposit of molybdenum in the first year. It is possible that a portion of the deposited mass on the shadowed QCM is not molybdenum, perhaps from general spacecraft outgassing contamination. For the thrusting conditions thus far, the shadowed QCM has accumulated approximately 10% of molybdenum deposited on the line-of-sight sensor. The source of this non-line-of-sight contaminant is attributed to ionized molybdenum, moving along trajectories effected by electrostatic potentials associated with the thruster plume and spacecraft surfaces. Since the DS1 solar arrays do not extend into the line-of-sight zone, are well removed from the thruster (>2 m), and are negatively grounded, the amount of molybdenum deposited on the SCARLET concentrator lenses is expected to be very small.

4.2.3 Deposition Dependence on IPS Mission Level—The deposition rates for QCM0 and QCM1 at various NSTAR mission levels are summarized in Figure 22. The effective deposition rate for the full power LDT is indicated in the right-hand side of the figure. Due to the IPS operations profile, there is no data available for mission levels 50 through 70. As indicated before, the line-of-sight QCM0 accumulates molybdenum at a substantially higher rate than

the shadowed QCM1. The line-of-sight sensor deposition rate appears roughly proportional to the square of the mission level, whereas the non-line-of-sight rate seems more strongly affected by mission level. The rate of production of ionized molybdenum is expected to increase dramatically with mission level for the following reasons[21]:

- More sputtered molybdenum atoms are produced at higher mission levels due to increased impingement by charge-exchange xenon.
- Higher electron temperatures are observed at higher mission levels, increasing the rate for electron-impact ionization of neutral molybdenum.
- More beam ions are produced by the engine, increasing the rate for charge-exchange ionization of molybdenum atoms.




Figure 22. Mo Deposition Rates Versus Mission Level (QCM0 is the line-of-sight sensor)

The molybdenum collection rate by the non-line-of-sight sensor normalized to that of the line-of-sight QCM is shown in Figure 23. The ratio appears to increase strongly with mission level. The early mission data points highlighted on the plot correspond to initial IPS operations at low and high mission levels that show an enhanced collection rate by the non-line-of-sight sensor. It is possible that this enhancement is due to contamination from spacecraft outgassing, since the ion engine heats the propulsion module assembly during operation. It is also possible that spacecraft outgassing contributed to the trend at high mission levels (> 70), since the early mission profile consisted of gradually decreasing thrust. Unfortunately, this ambiguity may not be directly resolved in the future because the Sun distance for DS1 will remain above 1.3 AU for the remainder of the mission, precluding IPS operations at mission levels greater than 70. Correlation of these rates with certain IPS telemetry, such as the accelerator grid impingement current, may improve the understanding of the mission-level dependence.

4.2.4 Thermo-optical Property Changes—The IDS contamination monitors include line-of-sight and non-line-

of-sight calorimeters. Under ideal conditions for analysis, calorimeters should have a 2π -steradian field-of-view to space. Of course, this condition is clearly not possible for the line-of-sight calorimeter. The requirement for Sun-viewing and the desire to correlate mass deposition with thermo-optical property changes drove the configuration of the non-line-of-sight calorimeter to the present state. The presence of the DS1 spacecraft (with IPS thruster) in the field-of-view of the calorimeters has substantially complicated the data analysis for these sensors. Some semi-quantitative analysis is possible for the line-of-sight calorimeter. The temperature of this calorimeter increased dramatically in the early part of the mission, as seen in Figure 24.




Figure 23. Ratio of Non-Line-of-Sight to Line-of-Sight Deposition Rates as a Function of Mission Level




Figure 24. Response of IDS Line-of-Sight Calorimeter During Initial IPS Operations

The active calorimeter element (the disk) increases in temperature by more than 50 °C within several days of high mission-level operation of the IPS thruster. The solar illumination of the calorimeter, illustrated in Figure 25, remained constant at about 93 mW/cm² from DOY 323-1998 through 338-1998. DS1 turned to the trajectory thrust

attitude on DOY 338-1998, reducing the solar input to approximately 75 mW/cm². The solar input gradually decreased to 70 mW/cm² until DOY 343-1998, when DS1 again changed attitude. The relatively constant period of solar illumination between DOY 322-1998 to 339-1998 provides the opportunity to simply estimate changes in the thermo-optical properties of the line-of-sight calorimeter.




Figure 25. Solar Irradiance History for the Line-of-Sight Calorimeter During Initial IPS Operations

The change in thermo-optical properties for the initial four days of high-power operation is determined from the values in Table 3.

Table 3. Selected Parameters for Estimating the Change in Thermo-optical Properties

Quantity	335-1998	339-1998
Insolation (mW/cm ²)	93.5	93.0
T _{disk} (°C)	11.1	45.2
T _{cup} (°C)	31.7	48.2
Q _{disk-cup} (mW)	31	3
T _{sky} (K)	243	243

It is first necessary to estimate the effective sky temperature due to the radiative heat load from the spacecraft and thruster into the calorimeter disk. The initial value for the ratio of solar absorptance to hemispherical emittance (α/ϵ) is taken as the pre-flight value, $0.1 = 0.08/0.8$, since little contamination was encountered during the launch phase. The pre-flight measured conductive heat leak between the disk and cup is 1.5×10^{-3} W/cm². At equilibrium, the radiative heat loss from the disk is equal to the solar-heat input and heat leak from the cup.

$$Q_{rad} = Q_{sun} + Q_{disk-cup} = \epsilon\sigma A(T_{disk}^4 - T_{sky}^4)$$

Using the values for DOY 335-1998, effective sky temperature is estimated to be 243 K (−30 °C). This value may seem high, but the NSTAR thruster may reach

temperatures of 500 K at in operation. The effective sky temperature is assumed to remain constant for DOY 339-1998. Neglecting the minor heat loss between the disk and cup, the estimated value for α/ϵ increases to 0.4. This is a significant change in radiator properties (to typical design end-of-life), with an estimated molybdenum accumulation of about 10 to 15 Å.

5.0 IPS PLASMA WAVE & EMI CHARACTERISTICS

5.1 Plasma Wave Electric-field Measurements

5.1.1 Ground Test—An IPS compatibility test (ICT) with the DS1 spacecraft was performed in the JPL 25-foot space simulator facility in February 1998. During the ICT, the IPS was briefly operated at TH0 (ML6), TH7–8 (ML55–ML62), and TH14 (ML104) thrust levels. The IDS Engineering Model, which included the flight Plasma Wave Antenna (PWA) pre-amplifier and Plasma Wave Spectrometer (PWS) board from TRW, was used in the ICT. The IDS used a rigid, non-flight 2-m tip-to-tip wire antenna to monitor electric field signals. A flight-like search coil was used to collect AC magnetic field data.

At the time of the DS1 ICT, the IDS software manager was not on-board; therefore, DS1/IDS command and data communications were invoked by primitive commands to the DS1 MIL-STD-1553B bus controller hardware. IDS could not transmit time-domain data in the “burst” mode because no processing of the IDS bus traffic was performed by the DS1 flight computer at this time. Data from IDS were captured by an external MIL-STD-1553B bus monitor. Therefore, the DS1 test conductor only executed IDS configuration or gain commands during periods of low spacecraft activity. No IDS commanding was performed during IPS thrust operations. The IDS team prepared several PWS gain commands in preparation for the ICT. For the initial ML6 operations, the PWS gain was set at a relatively low level. Upon examination of the PWS data, the PWS gain was set to a high level for the remainder of the ICT.

PWS electric-field data obtained during the DS1 ICT is shown in Figure 26. A few features are readily noted in the power spectra. A large peak appears in the 1-MHz to 15-MHz region, attributed to IPS electron-plasma frequency noise. A lesser peak is seen in the 200-Hz to 4-kHz region; the source of this signal is not yet understood. The amplitude of the PWS signal is less than 0.1 V_{p-p}/m, except near 15 MHz, where the signal approaches 0.3 V_{p-p}/m. Note that there is little signal observed in the 10-kHz to 300-kHz frequency region during the ICT.

5.1.2 Flight Measurements—For purposes of comparison with ground measurements made during the DS1 ICT, data from a brief IPS activity on DS1 to assess power production from the SCARLET solar arrays is presented. This DS1 test, referred to as “S-Peak,” operated the IPS for a relatively

brief interval (less than 40 minutes total). The IPS is always started with high-cathode flow rates; the characteristic time to reach steady-state-flow conditions is generally several hours. Therefore, this brief S-Peak test most closely resembles the IPS conditions during the ICT. Due to the spacecraft-to-Sun range, though, DS1 was not able to achieve the ML104 maximum level witnessed in the ICT.




Figure 26. Plasma Wave Spectrum for ICT Thrust Levels




Figure 27. Plasma Wave Spectrum for S-Peak Thrust Levels

PWS electric-field data obtained during the “S-Peak” test is shown in Figure 27. Since time-domain data collection was enabled to capture high-amplitude events during the S-Peak test, the PWS gain settings were lower than that for the DS1 ICT. The PWS noise “floor” for S-Peak is approximately 0.01 V_{p-p}/m . The high-frequency feature between 1 MHz and 15 MHz is about 10 dB higher in amplitude in the flight S-Peak than what was observed in the ground-based ICT. Unlike the ICT, essentially no signal amplitude is observed between 200 Hz to 4 kHz during S-Peak. A substantial signal is observed in the 10-kHz to 300-kHz frequency region in the S-Peak data (in contrast to the minimal signal observed in this frequency regime during ICT). Both the

ICT and S-Peak data sets appear to show an amplitude “dip” between 300 kHz and 2 MHz.

Characteristic plasma wave signal measurements under IPS steady-state thrust conditions were obtained during IPS Acceptance Tests IAT1 and IAT2. The results for IAT1 are shown in Figure 28. The plot-symbol size approximates the amplitude-error bars at high signal levels. The PWS signal might be expected to correlate with the thrust level for the IPS. The data in Figure 28 clearly shows no straight-forward correlation between plasma-noise amplitude and IPS-thrust level. Note that the highest thrust level (TH12, ML90) has a plasma-wave spectrum almost the same as that for TH3 (ML27). The highest plasma noise in IAT1 is observed for TH11 (ML83). Maximum signal levels, at 40 kHz and from 2 MHz to 15 MHz are approximately 0.5 V_{p-p}/m , similar to amplitudes observed in the S-Peak data. The behavior in the low-frequency region (below 10 kHz) with thrust level is not well understood, but could be due to inter-modulation between switching power-supply modules within the IPS power-processing unit.




Figure 28. Plasma Wave Spectra for IAT1 Mission Levels

Plasma-wave-noise measurements obtained during the lower thrust level IAT2 are shown in Figure 29. Note that the lowest thrust level (TH0, ML6) has a noise spectrum almost as high as that for TH4 (ML34). The spacecraft noise level just prior to initiation of IAT2 is plotted as a solid black line in the Figure. The spacecraft noise includes a signal from an unknown source in the 2-kHz to 7-kHz region. This signal appears to be attenuated by thruster operations at ML13 through ML26. Maximum signal levels, at 40 kHz and 2 MHz to 15 MHz, approach 1V/m. Again a characteristic “dip” in the spectrum is observed in the 300-kHz to 1-MHz frequency region.

The plasma noise from the IPS occasionally changes dramatically during thrust-level transitions. Upon transition to a higher thrust level, the IPS is designed to first increase the xenon flow, then increase the ion-beam current and

other IPS electrical parameters. Increased xenon flow at a fixed beam current, will increase the production of charge-exchange xenon. This charge-exchange xenon plasma behaves as an electrically conducting medium for the plasma noise. A dramatic example of this behavior is illustrated in Figure 30. The amplitude of the plasma noise in the 22-kHz band increases by 1000-fold during the 2-minute transition from ML20 to ML27. Note that the steady-state plasma wave signatures for these two thrust levels are within a factor of two of each other.




Figure 29. Plasma Wave Spectrum for IAT2 Mission Levels




Figure 30. Plasma Wave Spectrogram for IPS Transition from ML20 to ML27

The transition between IPS ML83 to ML90 is shown in Figure 31. In this case, the plasma noise decreases dramatically in the lower frequency region (<10 kHz). This phenomenon has been repeated in ground test by reducing neutralizer flow or discharge current. In the ground test, it is possible for a secondary plasma sheath associated with the chamber walls to envelope a portion of the antenna. In flight, the higher noise level at ML83 might be due to the amount of residual xenon available for producing a noisy plasma discharge within the neutralizer. Further experimentation in flight will not occur until after completion of the extended science mission because reduced xenon flow represents an erosion risk to the cathodes. (A

common plenum tank controls both the NSTAR IPS neutralizer and discharge cathodes; therefore, the erosion risk exists for both devices.)




Figure 31. Plasma Wave Spectrogram for IPS Transition from ML83 to ML90

5.2 AC Magnetic Fields (EMI)

5.2.1 Ground Test

In addition to the electric-field measurements, the IDS made simultaneous measurements of AC magnetic fields during the DS1 IPS compatibility test (ICT). In spite of setting the gain to the maximum level after the TH0 (ML6) initial firing of the IPS, no signals above the noise floor were recorded during the test. Prior to and subsequent to IDS delivery to the ICT, the IDS engineering model search coil easily detected AC magnetic field stimuli applied with a small excitation coil. The absence of AC magnetic signature in the ICT ground test is very surprising, given the amplitudes observed in flight.

Measurements were made with engineering model search coil in NSTAR characterization tests CT31 and CT36, capturing signals with a fast digital oscilloscope. As seen in Figure 32, the search coil shows a weak response to transient events, such as the IPS engine start, but does not show much electromagnetic interference (EMI) noise with steady-state engine operations. Whether the lack of strong AC magnetic signals is due to chamber effects or EMI-shielding or grounding considerations is under debate.




Figure 32. Response of the Search Coil Magnetometer to IPS Start During Ground Testing

5.2.2 *Flight Measurements*—AC magnetic-field data recorded by the IDS engineering search coil (SCM0) during IAT1 is shown in Figure 33. Some of the characteristic trends observed in the electric-field data (Figure 30) are also seen for the magnetic (B-fields). The highest amplitude B-fields are found at ML83 in the 1-kHz to 5-kHz region. The peak amplitude for ML90 is 10 dB below that of ML83, as found in the E-field spectra. The lowest B-fields in IAT1 are found at ML27 and ML48, which differs from the E-field measurements where ML90 was the least-noisy operating point. The lower-frequency signals (50 Hz to 200 Hz) appear to have less variation with operating level and are not consistent with the order witnessed in the 1-kHz to 5-kHz region. Until the IAT2 test was performed, the nature of the low-frequency magnetic field signals were not understood.




Figure 33. AC Magnetic Field Spectra for IAT1 Mission Levels

Data obtained during the DS1 IAT2 activity is shown in Figure 34. The figure shows the relative contribution to a known non-IPS source of EMI on the DS1 spacecraft: the engine gimbal assembly (EGA) stepper-motors for performing thrust vector control of the IPS engine. IAT2 included special EGA motion patterns for magnetic field and charge-exchange plume mapping experiments (this data is still under analyses). The attitude control system software maintained DS1 pointing using only the reaction control subsystem (RCS) hydrazine thrusters during this period of IAT2. As a result, the DS1 search coils could distinguish between EMI produced by the EGAs and the IPS during ion engine operations. Note that the EGA noise amplitudes are comparable to IPS noise, though at much lower frequency (< 400 Hz).

5.3 Plasma Wave Transient Signals

5.3.1 *Ground Test*—As indicated in section 5.1, the DS1 flight software to control the IDS was not available during

the ICT. Time-domain data from the plasma wave antenna and search coil sensors could not be captured during this integrated ground test of DS1 and IPS. Time-domain waveform data from plasma wave antennas were recorded during NSTAR developmental and characterization tests using flight-like sensors and laboratory digital oscilloscopes.




Figure 34. AC Magnetic Field Spectra for IAT2 Mission Levels

Examples of a typical high-amplitude, IPS-generated event are shown in Figure 35 and Figure 36. This event occurs during discharge ignition during IPS start-up. An actively amplified monopole antenna detected the data in Figure 35. The amplitude of this event is 8 V_{p-p}/m. Data shown in Figure 36 was simultaneously recorded with a 2-m tip-to-tip dipole antenna with an engineering model IDS PWA pre-amplifier. Notice that amplitude recorded by the dipole antenna is only about 2 V_{p-p}/m, about a factor of 4 less than the monopole signal.




Figure 35. IPS Ignition in CT36 Ground Test (monopole)

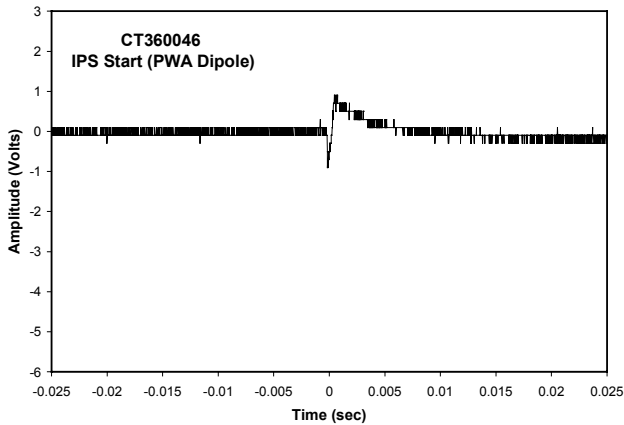


Figure 36. IPS Ignition in CT36 Ground Test (PWA dipole)




Figure 38. B-Field Transient Signal for Flight IPS Ignition

The IDS recorded several IPS-ignition events in flight. Data for a typical IPS ignition is shown in Figure 37 below. The peak signal at $t=0$ seconds is approximately 1 V/m, consistent with the level observed in PWA dipole measurement from the CT36 ground test. After the ignition event, the noise from the IPS plasma is clearly visible in the IDS PWA data. Simultaneous magnetic field data for IPS ignition from the IDS search-coil magnetometer is displayed in Figure 38. Peak field strengths of about 50,000 nT are observed for IPS-discharge ignition.




Figure 37. E-field Transient Signal for Flight IPS Ignition




Figure 39. E-field Signature for IPS Recycle at $t = -0.45$ (The large signal near $t=0$ is due to hydrazine thrusters firing.)




Figure 40. B-field Signature for IPS Recycle at $t = -0.45$ (The large signals from $t = -0.3$ to 0.45 are from gimbal actuators. The transient spike at $t=0$ is from the RCS thruster valve.)

The IPS can also produce high-amplitude transient-field events when a momentary ionization arc between the grids induces a “recycle” event. The NSTAR power processor unit will disable the ion beam power supplies within a few microseconds of a fault condition in the output. Within a second of disabling the beam supplies, the power processor gradually restores the beam supplies to the thrust level. Examples of the E- and B-field transients for a recycle event are shown in Figures 39 and 40, respectively.

Notice that the IPS stops at $t = -0.45$ seconds and an RCS thruster firing occurs at $t = 0$. The low frequency magnetic oscillations between $t = -0.3$ and $t = 0.45$ are due to the engine gimbal assembly motors.

The DS1 reaction control system (RCS) thrusters are responsible for some of the largest amplitude-transient signals observed by the IDS. As shown in Figures 39, 40, 41, and 42, the RCS-produced signals are substantial. Electric-field amplitudes in excess of $2 \text{ V}_{p-p}/\text{m}$ are typically observed for the RCS thruster firings. The origin of the high-amplitude E-field signal is not fully understood; however, a strong candidate is the ability of low-density gas flows to discharge electrically charged surfaces. The plasma wave antenna will become moderately charged due to the photoelectric effect. Some variation of the E-field amplitude has been observed with changes in Sun angle on DS1, supporting the possibility that charge dissipation is responsible for the signals. The magnetic field signals in Figure 42 are attributed to the solenoid valve-drive pulses. The various thruster firing combinations on DS1 yield unique, but reproducible, magnetic-field signatures. The magnetic field signature typically begins approximately 15 msec prior to the electric field signal in RCS thruster firings.

On several occasions, strong E-field transient events have been recorded by the IDS without RCS or IPS operations. These E-field signals do not have a simultaneous magnetic signature, suggesting a momentary plasma discharge. Such events have been attributed to hypervelocity impacts and have been observed in prior space missions (for example, Voyager). Figures 43 and 44 provide an example of such an event on DS1.




Figure 41. E-field Signature for RCS Thrusters Firing at $t=0$




Figure 42. B-field Signature for RCS Thrusters Firing at $t=0$




Figure 43. E-field Signature for Particle Impact at $t=0$




Figure 44. B-field Recording for Particle Impact at $t=0$

6.0 IPS DC-MAGNETIC FIELDS

6.1 Ground Magnetic Field Mapping

The NSTAR ion engine includes strong permanent magnets arranged in a “ring-cusp” geometry to enhance the ionization efficiency within the discharge chamber[7]. These rare-earth permanent magnets are fabricated from samarium-cobalt ($\text{Sm}_2\text{Co}_{17}$) and have been thermally conditioned to improve their long-term stability. An important issue regarding the IPS permanent magnets is the stability of the fields during the lifetime of a science mission. The magnets are known to exhibit temperature-dependent changes in field strength; this dependence can be accurately determined prior to launch. The long-term stability of the temperature-compensated magnetic field characteristics is a critical factor for determining the compatibility of IPS with magnetic field science measurements during a mission.

A simple finite-element magnetic field model for the NSTAR ion engine was constructed using the student version of Q-Field (Tera Analysis). The configuration of the magnets permitted a simple, axial-symmetric model to be constructed. The location and pole orientation for the magnets were determined from NSTAR assembly drawings. The magnetic properties of the $\text{Sm}_2\text{Co}_{17}$ were obtained from the supplier literature. Figure 45 illustrates the magnetic flux density with a color scale and magnetic field lines at contour intervals of 500,000 nT. An outline for the NSTAR ion-engine shell is provided in the figure for clarity. There is a large external field lobe opposite from the ion-beam direction. The internal “ring-cusp” field lines are evident within the discharge chamber region. The upper bound for the field magnitudes for the IDS FGM sensors is approximately 11,000 nT for the inboard sensor and 3,200 nT for the outboard sensor.

The initial assessment of IPS magnetic fields and the long-term stability was performed in conjunction with the NSTAR 8000-hour life demonstration test (LDT) performed with an engineering model thruster (EMT#2). Prior to the start of the LDT, EMT#2 was characterized in the JPL Magnetic Mapping facility. As expected, very strong magnetic fields were observed in the mapping operation. A polar plot of the IPS magnetic is shown in Figure 46. This plot is overlaid upon a cross-sectional view of the IPS thruster. (Note that the orientation of the engine is reversed in Figures 45 and 46). The peak field, at a 1-meter distance from the approximate center of the IPS, was found to be 12,000 nT along the thruster centerline. Smaller field lobes were found roughly perpendicular to the thrust axis. This external-field geometry is consistent with the configuration and orientation of the magnets within the thruster assembly. The slight tilt of the lobes perpendicular to the thruster centerline is due to an offset of the engine magnetic “center” from the axis of the rotation table in the magnetic mapping

facility. Based on the EMT measurements, the predicted field magnitude for the IDS FGM sensors was about 7000 nT for the inboard sensor and 2800 nT for the outboard sensor. The variance from the magnetic model is due primarily to the effects of the thermal conditioning on the $\text{Sm}_2\text{Co}_{17}$ magnetics.




Figure 45. Magnetic Field Model for the NSTAR Ion Engine




Figure 46. Pre-LDT Magnetic Map of EMT#2 Thruster

Subsequent to the completion of the 8000-hour LDT, the EMT#2 ion thruster was returned to the JPL Magnetic mapping facility. Since permanent fixtures for precisely positioning the IPS engine and magnetometer sensors within the magnetic mapping were not available, the mapping configuration was reconstructed based on photographic documentation of the pre-LDT set-up. The post-LDT mapping data were found to repeat the original results (within the ability to accurately re-create the pre-LDT mapping configuration). The estimated limit of magnetic

strength degradation for the 8000-hour test is less than 5%. The thermally-conditioned permanent $\text{Sm}_2\text{Co}_{17}$ magnets used in the NSTAR ion engine demonstrate stable magnetic characteristics after long-term operation at full power.

6.2 Flight Measurements

The following describes some long-term investigations on the DS1 FGM data. There is a list of interesting questions concerning the behavior of the ion engine permanent magnetics with respect to the FGM data.

- (1) Does the temperature play a significant role in the magnetic measurements?
- (2) If there is a temperature dependency, is it possible to make a model for temperature correction on the data?
- (3) Do the magnetic moments of the ion engine magnets vary with the time?

The following results are based on data transmitted to the TU-Braunschweig from launch to DOY 077-1999. Therefore, only the first six months of the mission is covered by this analysis. For the day of the encounter of DS1 at Braille, limited data are available (DOY 209/210-1999).

When the data are shown versus time, the x-axis is shown in units of day of the year 1999. Thus the days in 1998 are handled as “negative days.” The magnetic and temperature data on the y-axis is the average of the specific data over the period of the assigned day (24-hour average).

6.2.1 Investigation of the Mean Residual Field—The plots in Figure 47 show the 24-hour averaged FGM data of the

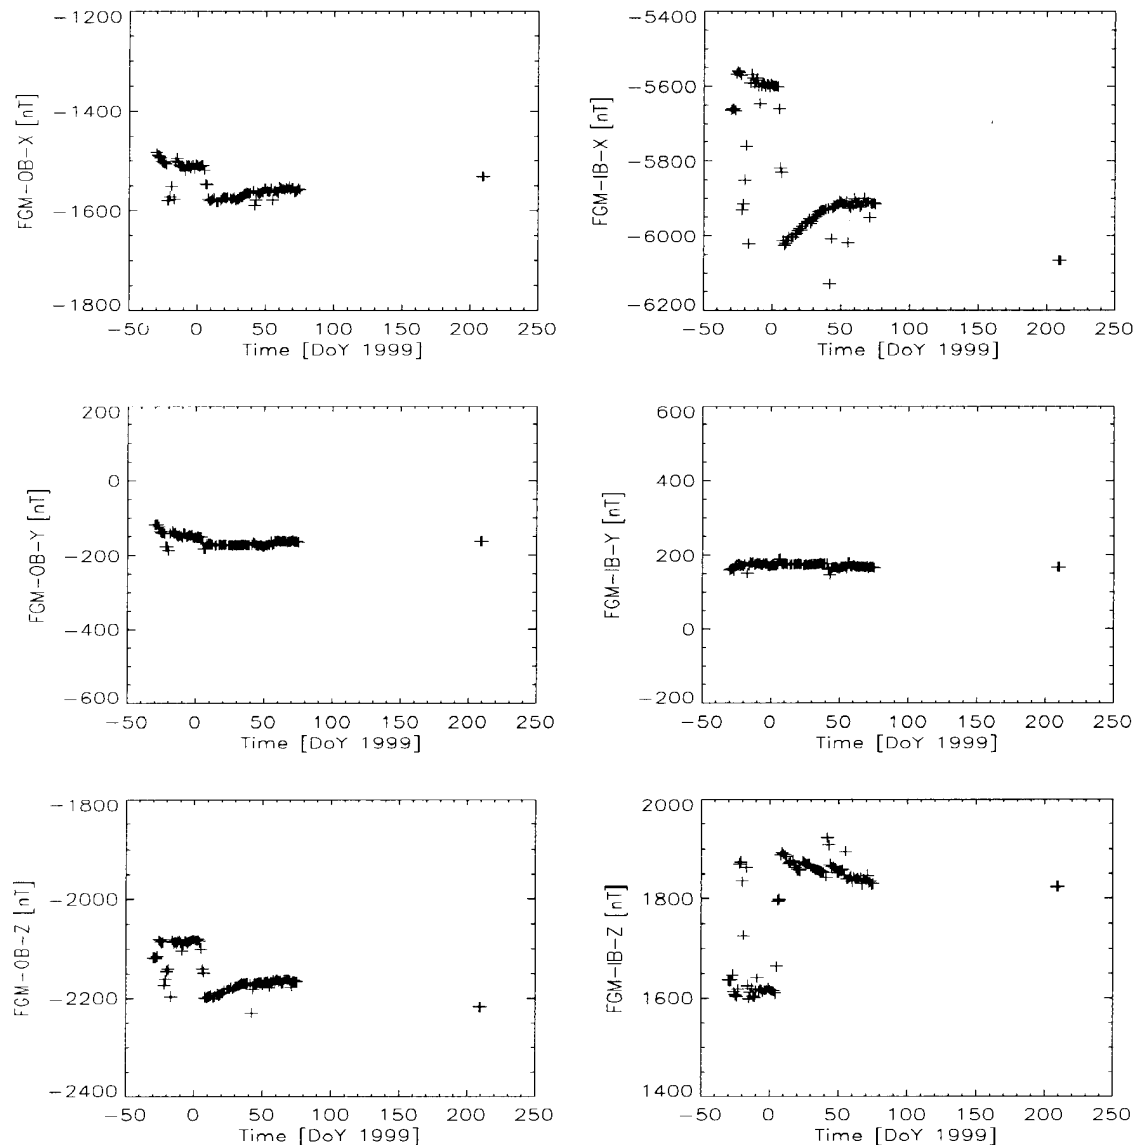


Figure 47. The 24-hour Averaged, Calibrated FGM Data in DS1 Coordinates

outboard (left column) magnetometer and the inboard (right column) magnetometer. The data are calibrated and displayed in spacecraft coordinates.

The ambient field of interplanetary space is in the order of a few nanotesla; the offsets of the magnetometers are also in the order of a few nanotesla. Therefore, it is quite obvious that the resulting huge magnetometer readings are caused by the spacecraft. The data show a strong variation over the time, especially in the x and z components. The inboard sensor (FGM-IB) shows larger absolute values and higher variations. These field variations are caused by the IPS permanent magnets. It is a known fact that the magnetic moment of a probe is strongly temperature dependent. Therefore, the next step is to look at the various temperature sensors on board DS1.

Figure 48 shows the data of four temperature sensors:

- T_INT refers to the internal ion-engine temperature sensor (IPS_THR_TMP).
- T_EXT refers to the external ion-engine temperature sensor (IPSTHRMSKTMP).

- T_FGM_IB refers to the inboard magnetometer temperature sensor.
- T_FGM_OB refers to the outboard magnetometer temperature sensor.

All the sensors show nearly the same structure; however, the sensors show different absolute values and different amounts of variation. At the beginning of the mission high temperatures are indicated. This corresponds to the operating ion engine. The engine was switched off on January 5, 1999. The sudden temperature decrease is easily seen in the data. The operating ion engine causes a higher temperature on the outboard sensor than on the inboard sensor. This might be due to the fact that the outboard sensor is placed a little bit nearer to the ion beam regime than the inboard sensor. The temperature measured inside the propulsion system decreases by about 150 °C when the engine is deactivated. In the following time, the system seems to be heated up exponentially. This is probably caused by gradual change in solar flux on the engine.




Figure 48. The 24-hour Averaged Temperature Data from Thruster and FGM Sensors

Deep Space 1 Technology Validation Report—*Ion Propulsion Subsystem Environmental Effects on Deep Space 1: Initial Results from the IPS Diagnostics Subsystem*

The comparison between the FGM magnetic field data and the measured temperatures suggests a linear model of the temperature dependence of the magnetic moments of the ion engine permanent magnets.

The best fit of such a model is shown in Figure 49. The intercept and slopes for the best fit for each component are printed above each plot. All components of the measured

magnetic field data show a linear dependency of the internal temperature T_{INT} . The x and z components show huge temperature variations. This is due to the geometric orientation of the magnetometer on the boom relative to the engine magnets. At 0 °C, the inboard FGM is in a 6315-nT field, whereas the outboard FGM is in a 2710-nT field.

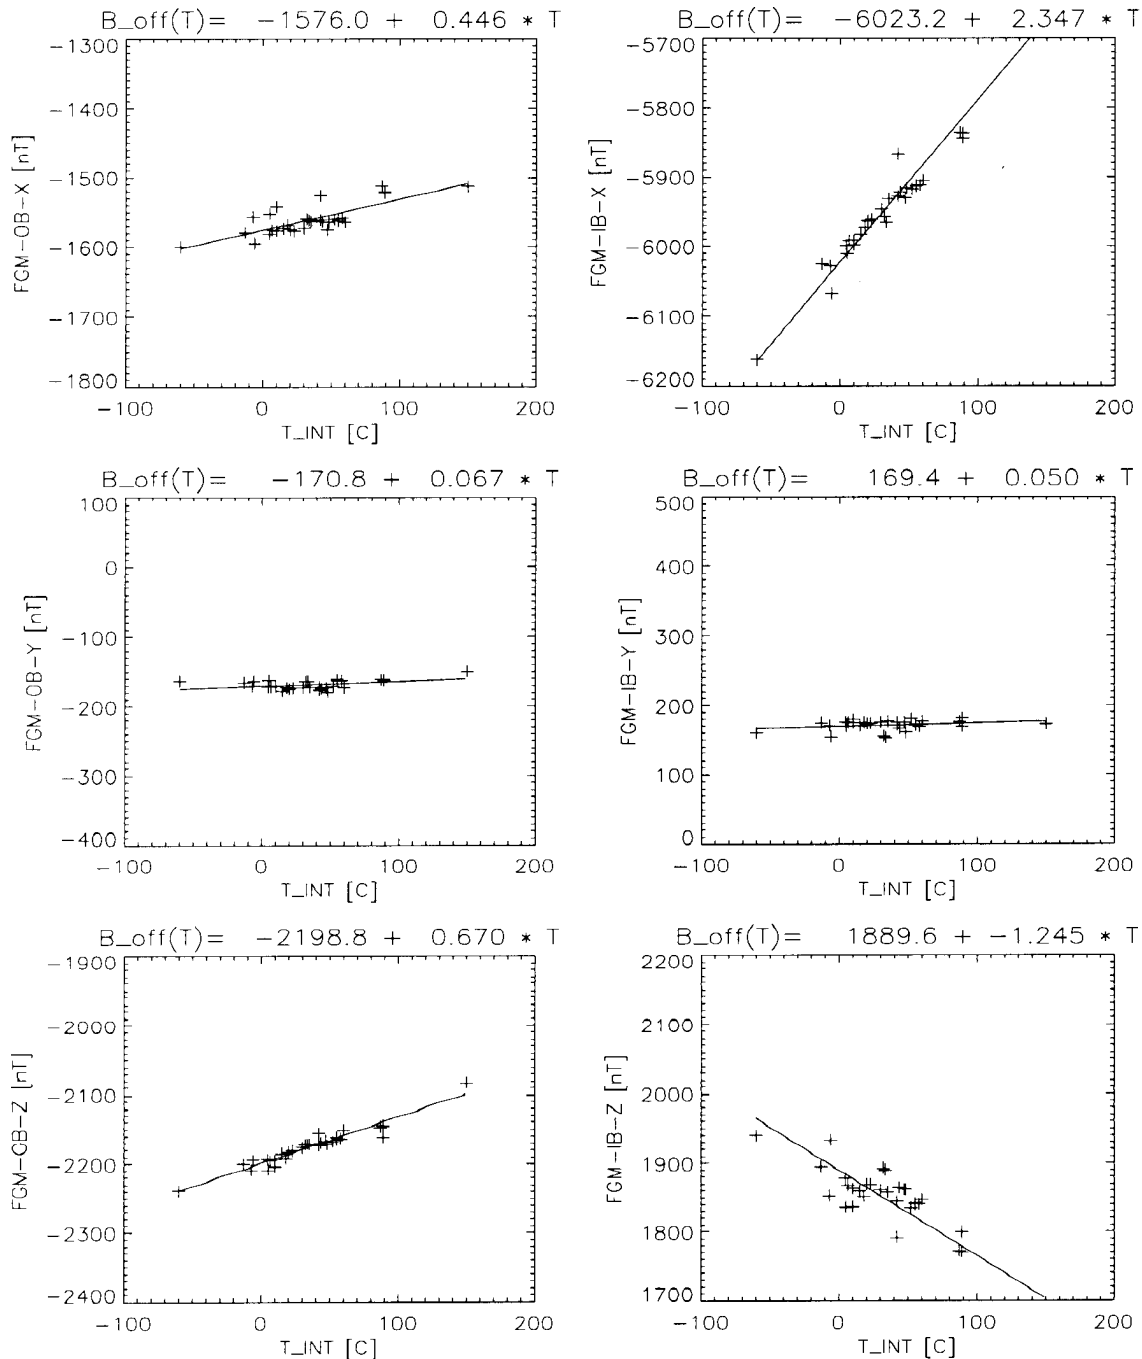


Figure 49. The Linear Temperature Model of the IPS Thruster Magnetic Field

The application of this model to the data leads to linear-temperature-corrected magnetic-field data, shown in Figure 50. The strong temperature dependence is diminished and the resulting residual field is suppressed. However, the model is not completely perfect. Especially on the x and z components of the FGM-IB magnetometer, which is located near the magnet, some linear (in the time domain, not in the temperature domain!) trend remains. This could be caused

by a temporal variation of the magnets themselves. Further investigation is required to resolve the magnetic stability.

A further cross check of the temperature model is given by investigation of the engine-temperature-corrected data versus the FGM sensor temperatures. These data are shown in Figure 51.

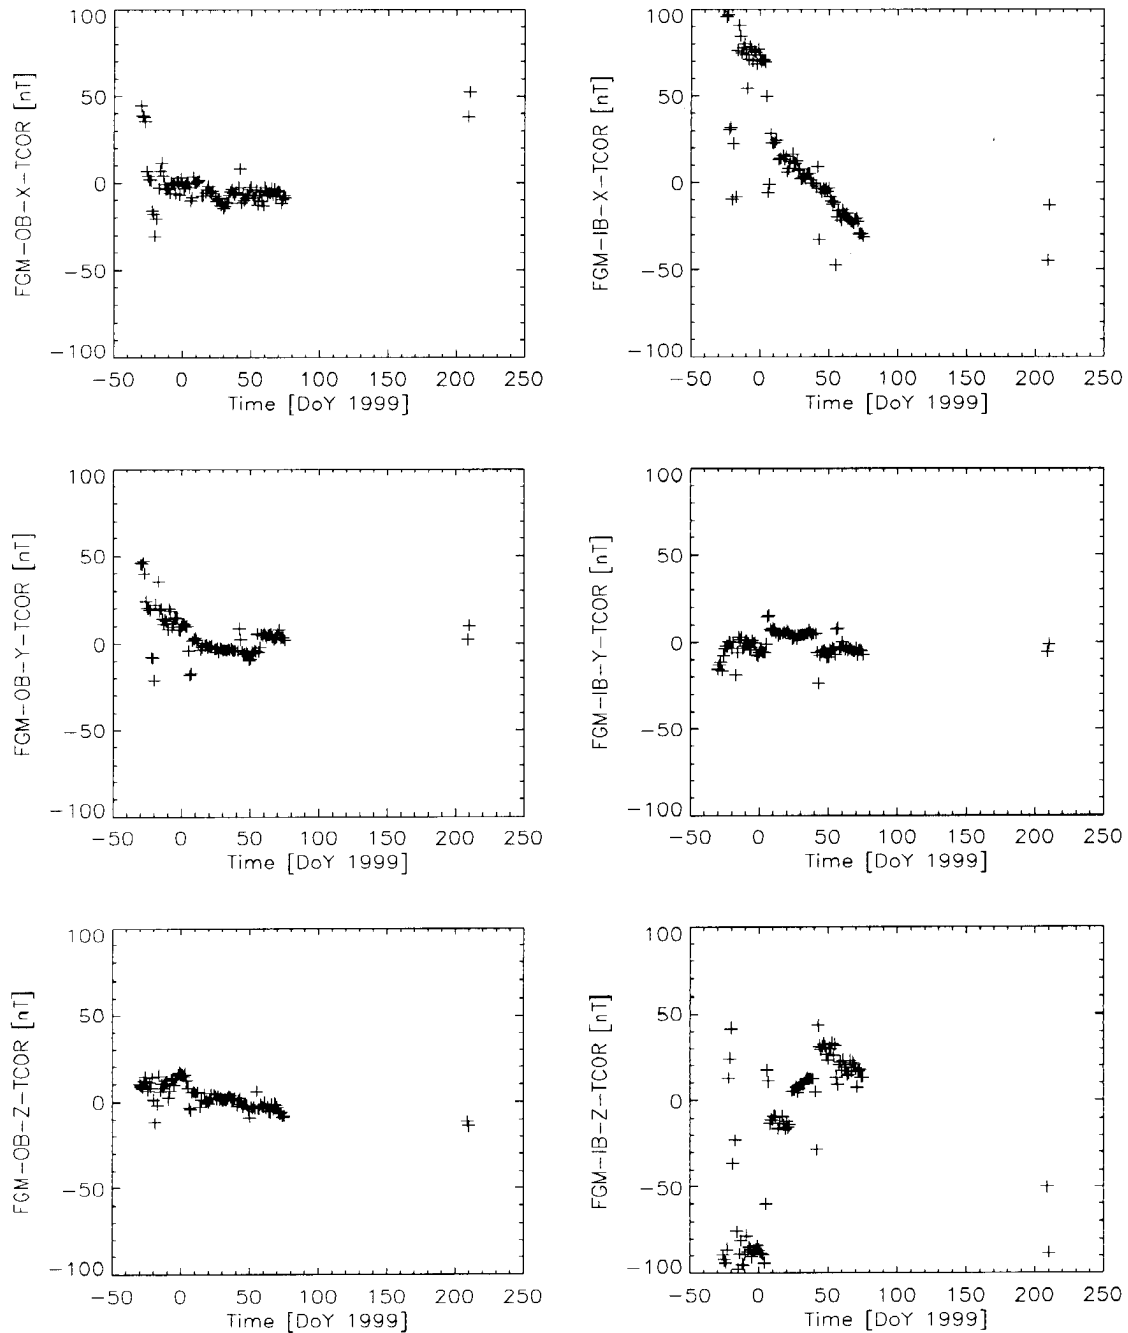


Figure 50. Residual Magnetic Field Data After Temperature Correction

Deep Space 1 Technology Validation Report—*Ion Propulsion Subsystem Environmental Effects on Deep Space 1: Initial Results from the IPS Diagnostics Subsystem*

The plots show that there is almost no temperature dependence to be seen. The data are straying nearly

randomly. This means that the temperature model does include the temperature-caused effects sufficiently.

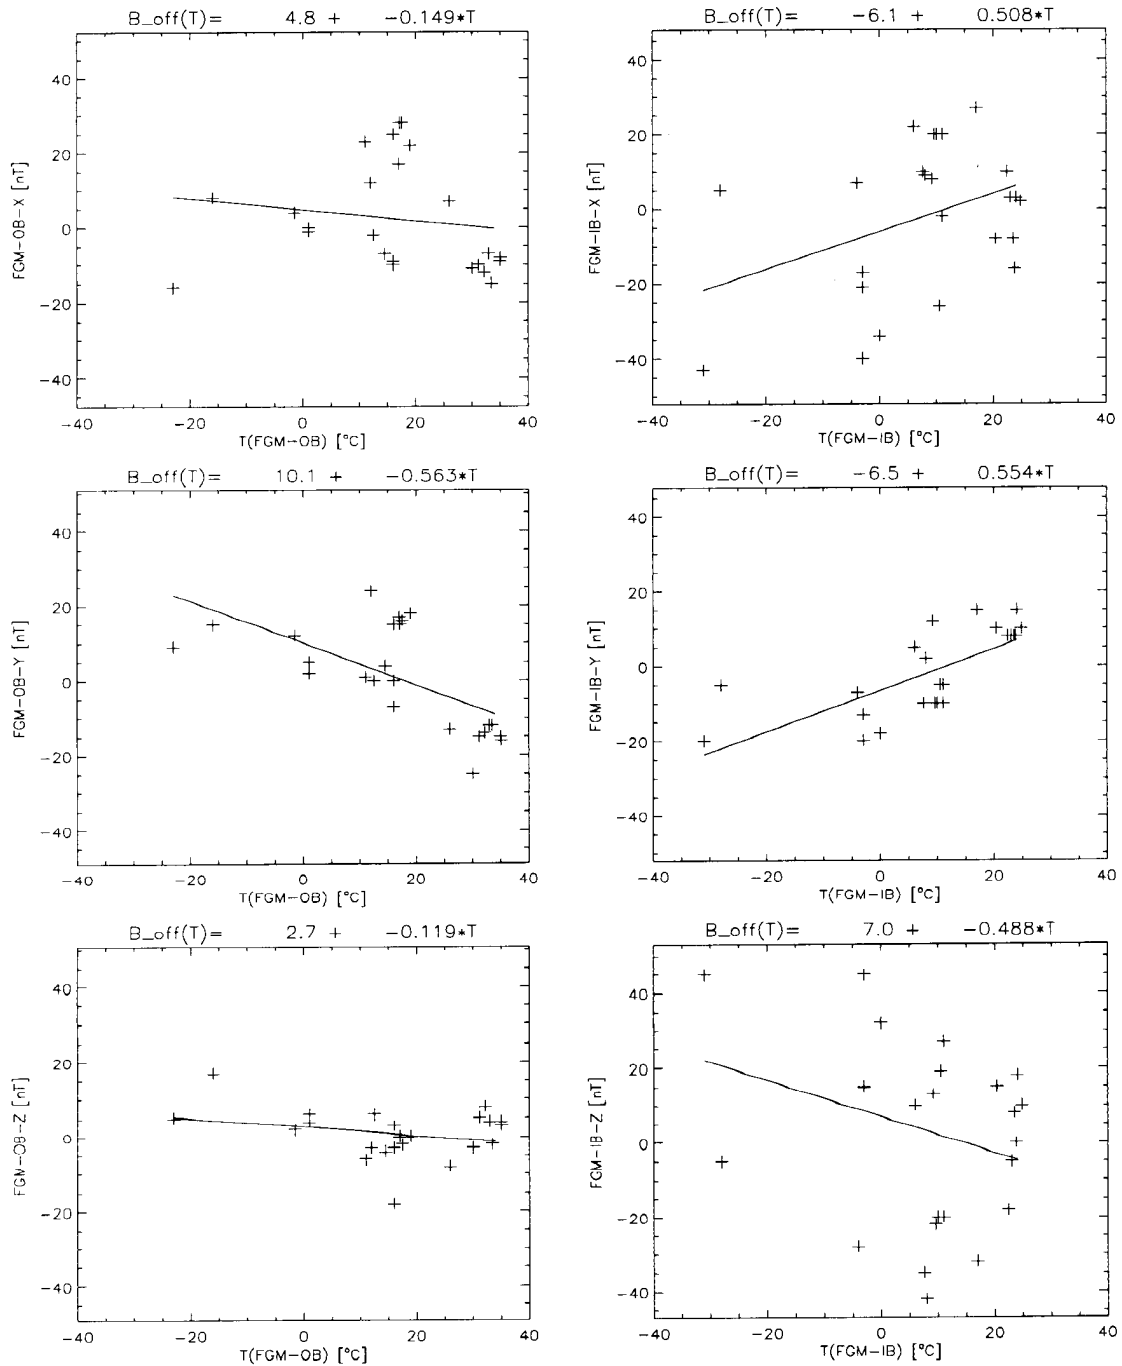


Figure 51. Temperature-corrected FGM Field Data versus FGM Temperature

7.0 CONCLUSIONS

DS1 provided an excellent opportunity for in-depth investigation of interactions of an ion propulsion system with an interplanetary spacecraft in flight. The NSTAR project recognized the importance of characterizing the local environment due to IPS operations and chose to fly a diverse set of instrumentation. The sensors were selected to capture the range of expected signals from the IPS. Hence, the sensor sensitivity and response characteristics are generally less than what is found in space-science instrumentation. Notable exceptions to the above statement are the flux-gate magnetometers provided by the Technical University of Braunschweig. The FGMs have performed exceptionally well throughout the mission and may have detected a weak (2 nT) magnetic signature during the flyby of Asteroid Braille[23]. The IDS has succeeded in collecting the data required to characterize the local environment and effects induced by the IPS operating on DS1.

Analysis of the IDS measurements of ion energies and densities and electron temperatures have validated sophisticated numerical-simulation models of the charge-exchange plasma produced by the IPS. Although a wider Langmuir-probe voltage-sweep range would have permitted independent electron density determination, the Langmuir probe performance was sufficient to obtain electron temperature data. Ion-current measurements from the retarding-potential analyzer allowed the charge-exchange density to be determined. The IPS charge-exchange plasma induced a shift of the DS1 spacecraft potential by -6 V to -10 V with respect to ambient “space ground.” In ground testing, the spacecraft was tied to Earth ground, as were the walls of the vacuum facility. A peak plasma potential of 5 to 7 V was observed in ground test, whereas the peak plasma potential exceeded 15 V with respect to spacecraft ground in flight. In terms of effects of the charge exchange on spacecraft subsystems, no degradation of the spacecraft power system due to parasitic current collection by the SCARLET solar arrays was observed during IPS operations. PEPE detected a substantial flux of charge-exchange xenon ions without effecting measurements of solar wind protons. The effects of IPS operations on PEPE solar wind electrons are still being evaluated.

The IDS contamination monitors returned high-quality measurements of deposition of IPS grid-erosion products during the DS1 mission. The line-of-sight quartz crystal microbalance accumulated 25 nm of molybdenum after 3500 hours of IPS operation. The line-of-sight accumulation is consistent with deposition observed during the 8000-hour LDT. The amount of non-line-of-sight molybdenum accumulation (2.5 nm after 3500 hours) is higher than the pre-flight prediction of <0.5 nm. Non-line-of-sight deposition is due to surface accumulation of molybdenum ions, whose trajectories are deflected by local electrostatic

fields on DS1. The pre-flight estimate of molybdenum ion production did not include the possibility of charge-exchange between neutral molybdenum and beam ions. Subsequent communication[22] revealed the Mo-Xe⁺ cross-section is surprisingly large; this channel dominates in the formation of molybdenum ions. Non-line-of-sight contamination measurements in ground test are not feasible due to interference from material sputtered from chamber walls. Flight measurements are the only reliable source for assessing non-line-of-sight deposition from the IPS engine on DS1.

The IDS Plasma Wave Spectrometer characterized the electrostatic wave and electromagnetic noise environments produced by the IPS and other DS1 subsystems. A large volume of both spectral and time-domain data were obtained throughout the DS1 mission, especially during IPS operations. There is not a direct correlation of noise amplitude with IPS operating power. The IPS noise levels are bounded as follows:

- IPS E-field continuous noise: < 1 V/m, < 15 MHz.
- IPS E-field transient: < 2 V/m for < 1 ms.
- IPS B-field continuous noise: < 10 μ T, < 10 kHz.
- IPS B-field transient: < 200 μ T for < 2 ms.

Limits for the major DS1 subsystem noise sources, namely the hydrazine reaction control subsystem (RCS) thrusters and engine gimbal actuators (EGAs), are bounded by:

- RCS thruster E-field transient: < 5 V/m for < 10 ms.
- RCS thruster B-field transient: < 200 μ T for < 40 ms.
- EGA B-field continuous noise: < 10 μ T, at 100 Hz.
- EGA B-field transient: < 100 μ T for < 1 s.

From a spacecraft systems-engineering perspective, the IPS does not produce *peak* electromagnetic or electrostatic noise beyond that of other spacecraft subsystems[15]. Note that, when operating, the IPS produces noise continuously; conversely, the other spacecraft sources are typically transient in nature. A major finding is the IPS does not introduce any interference in spacecraft communications or other subsystem operations.

The presence of high-strength permanent magnets within the IPS is a concern for performing magnetic-science measurements. The opportunity for science measurements is improved if the IPS permanent magnetic field can be accurately characterized and removed as background from the science measurement. For high-sensitivity magnetometers, it is important to locate the sensors as far as possible from the IPS and have accurate knowledge of the geometric orientation and temperature-compensated magnetic fields. Long-term degradation of the magnets can introduce significant errors in this approach to removing the background field. The IDS FGM sensors provided in-flight temperature-compensation data for and demonstrated the

long-term stability of the IPS magnets. The use of dual FGM sensors and principal component-analysis technique led to the possible detection of a weak magnetic signature at Asteroid Braille[23], demonstrating the potential for performing magnetic science even in the presence of large, local magnetic fields.

8.0 ACKNOWLEDGMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government, or the Jet Propulsion Laboratory, California Institute of Technology.

The IDS development effort was primarily funded by the NSTAR Project under the management of John F. Stocky, who provided valuable advice as well a programmatic support. The support of the NSTAR Project Office, especially from Michael Marcucci, Jim Tribbett, and Henrik Gronroos, is deeply appreciated. The ground-test portion of this effort was made possible by the following JPL personnel: John Brophy, Jay Polk, Keith Goodfellow, Pablo Narvaez, and John Anderson. Additional IDS development support was provided by members of the Integrated Space Physics Instrument team, including the following individuals not included in the author list: Kim Leschley, Gerry Murphy, and Gregg Vane. The following individuals at Physical Sciences, Inc. contributed substantially to the DSEU hardware and software development: Prakash Joshi, B. David Green, George Caledonia, Eric Lund, Michael Hinds, and Brian Root. Valuable hardware fabrication and test support was received from Chris Coles at Sentran Corp. Key personnel supporting fabrication and assembly of the IDS at JPL include Dave Rooney, John Bousman, Michael Parks, Michael O'Connell, and Greg Hickey.

The Deep Space 1 spacecraft team provided substantial assistance in the development, integration, and test of the IDS hardware. The authors appreciate the extra effort from the following individuals: Leslie Livesay, Gaylon McSmith, J. Sean Howard, Gary Glass, and Andy Rose. The New Millennium Program Chief Scientist, David Crisp, provided crucial support in obtaining approvals for the TUB hardware contribution. The DS1 Mission Operations and Deep Space Network teams supported the IDS operations phase with much appreciated extra effort from Marc Rayman, Phil Varghese, Bud Ford, Hank Hotz, Robert Gounley, Curt Eggemeyer, Kathy Moyd, and Tom Boreham.

Important hardware contributions to the IDS were made by the Institute for Geophysics and Meteorology of the Technical University of Braunschweig (TUB) and TRW. TUB provided the flux-gate magnetometers with signal processing electronics. TRW provided the plasma-wave spectrometer electronics and the plasma wave pre-amplifier. These contributions made possible the important electric and magnetic fields measurements obtained by IDS and substantially augmented the validation of ion propulsion on DS1.

The work by Guenter Musmann, Falko Kuhnke, Ingo Richter, Carsten Othmer, and Karl-Heinz Glassmeier was financially supported by the German Bundesministerium fuer Bildung und Forschung and the Deutsches Zentrum fuer Luft- und Raumfahrt under contract 50 OO 99037.

9.0 REFERENCES

- [1] Brinza, D., editor, Report on the NASA/USAF Workshop on Environmental Diagnostics for ELITE/STAR, *JPL D-11595*, 1994.
- [2] Carruth, M.R., A review of studies on ion thruster beam and charge-exchange plasmas, *AIAA Paper 82-1994*, 1982.
- [3] Katz, I., V. Davis, J. Wang, and D. Brinza, Electrical potentials in the NSTAR charge-exchange plume, *IEPC Paper 97-042*, 1997.
- [4] Katz, I., Parks, D. E., Mandell, M. J., and Schnuelle, G. W., Parasitic current losses due to solar-electric propulsion generated plasmas, *J. Spacecraft and Rockets*, 19(2), 129, 1982.
- [5] Jongeward, G., Mandell, M. J., Katz, I., Bucholtz, B. Snyder, S. and Wilbur, P., Conductive Nature of Low Frequency (<1 MHz) Electromagnetic Fields Generated by a Hollow Cathode Plasma Contactor, *AIAA Paper 94-3313*, 1994.
- [6] Huba, J. D., NRL Plasma Formulary, Revised 1998, Naval Research Laboratory, *NRL/PU/6790—98-358*, 28, 1998.
- [7] Patterson, M. J., Performance characteristics of ring-cusp thrusters with xenon propellant, *AIAA Paper 86-1392*, 1986.
- [8] Brophy, J. R., Polk, J. E., and Pless, L. C., Test-to-failure of a two-grid, 30-cm-dia. ion accelerator system, *IEPC Paper 93-172*, 1993.
- [9] Christensen, J. A., Freick, K. J., Hamel, D. J., Hart, S. L., Norenberg, K. T., Haag, T. W., Patterson, M. J., Rawlin, V. K., Sovey, J. S. and Anderson, J. R., Design and fabrication of a flight model 2.3 kW ion thruster for the Deep Space 1 Mission, *AIAA Paper 98-3327*, 1998.
- [10] Polk, J. E., Anderson, J. R., Brophy, J. R., Rawlin, V. K., Patterson, M. J., Sovey, J., and Hamley, J., An overview of the results from an 8200 hour wear test of the NSTAR ion thruster, *AIAA Paper 99-2446*, 1999.

Deep Space 1 Technology Validation Report—*Ion Propulsion Subsystem Environmental Effects on Deep Space 1: Initial Results from the IPS Diagnostics Subsystem*

- [11] Wang, J., D. Brinza, and Young, M., Three-dimensional particle simulation modeling of ion propulsion plasma environment for Deep Space 1, submitted to *J. Spacecraft and Rockets*, 2000.
- [12] Davis, V., Katz, I., Mandell, M., Brinza, D., Henry, M., Wang, J. and Young, D., Ion engine generated charge-exchange environment, comparison between NSTAR flight data and numerical simulations, *AIAA Paper 2000-3529*, 2000.
- [13] Arnold, G. S., Brinza, D. E., Joshi, P., Keener, D. N., Space active modular materials experiment, *SPIE Proc. Vol. 3427*, 225, 1998.
- [14] Joshi, P., Malonson, M., Green, B. D., McKay, J., Brinza, D. E. and Arnold, G. S., Space Environment and Effect Monitoring Instrumentation for Small Satellites, *J. Spacecraft and Rockets*, 35(6), 821, 1998.
- [15] Henry, M. D., Brinza, D. E., Mactutis, A. T., McCarty, K. P., Rademacher, J. D., vanZandt, T. R., Johnson, R., Musmann, G. and Kunke, F., NSTAR Diagnostics Package Architecture and Deep Space One Spacecraft Event Detection, *IEEE 2000 Aerospace Conference Paper 11.0502*, 2000.
- [16] Wallace, D. A. and Wallace, S. A., Realistic performance specifications for flight quartz crystal microbalance instruments for contamination measurement on spacecraft, *AIAA Paper 88-2727*, 1988.
- [17] Reichardt, P. J. and Triolo, J. J., Preflight testing of the ATS-1 thermal coatings experiment, *Thermophysics of Spacecraft and Planetary Bodies. Progress in Astronautics and Aeronautics, Vol. 20*, Academic Press, 1967.
- [18] Hurst, E. B. and Thomas, G. Z., Diagnostic system design for the ion auxiliary propulsion system (IAPS) – Flight test of two 8 cm mercury ion, *NASA Technical Memorandum 81702*, 1981.
- [19] Sakabe, S. and Izawa, Y., Simple formula for the cross sections of resonant charge transfer between atoms and their positive ions at low impact velocity, *Phys Rev A*, 45(3), 2086, 1992.
- [20] Polk, J. E., Brophy, J. R. and Wang, J., Spatial and temporal distribution of ion engine accelerator grid erosion, *AIAA Paper 95-2924*, 1995.
- [21] Brinza, D., J. Wang, J. Polk, and M. Henry, Deep Space One Measurements of Ion Propulsion Contamination, submitted to *J. Spacecraft and Rockets*, 2000.
- [22] Private communication with Rainer Dressler, AFRL/VSBS.
- [23] Richter, I., Othmer, C., Kuhnke, F., Glassmeier, K.-H., Brinza, D. and Tsurutani, B., Magnetometer observations during the fly-by of the Deep Space 1 spacecraft at the Asteroid Braille, *GP31A-09*, 2000 American Geophysical Union Spring Meeting, Washington D.C, 2000.

Appendix A. List of Telemetry Channels and Names

There is a fairly long list of IPS and spacecraft telemetry channels that indicate the state of the ion engine, spacecraft attitude, etc. The IDS instrument data, though, is not channelized, the data is contained in APIDs 3 and 4. We developed specific post-processing software to decommutate IDS sensor data and apply engineering unit

conversion factors. The DS1 channelized data is essentially ancillary data required to interpret the IDS data. The following list of channelized data is only part of the picture for our NSTAR validation activity. (David Brinza and Michael Henry, 10/15/99.)

Channel	Mnemonic	V-0135	XBEAMCUR	P-3202	XCA_plat_htr
P-3149	pdufet_dseu	V-0136	XBEAMVOL	P-3203	XE_tank_htr
P-2064	non_bus2_i	V-0137	XDISCURR	P-3204	XE_line1_htr
V-4068	ace_dseu1_t	V-0138	XDISVOLT	P-3205	XE_pl_ln_htr
V-4069	ace_dseu2_t	V-0139	XDHTRCUR	P-3206	XE_pl_t1_htr
A-4005	ace_rsu_t	V-0140	XDHTRVOL	P-3207	XE_pl_t2_htr
V-0430	DSEUdig_bd_t	V-0141	XLINECUR	P-3208	DCIU_htr
V-0450	RSUpreamp_t	V-0142	XLINEVOL	P-3209	PPU_htr
V-0447	DSEUmgr_mode	V-0143	XNTRCURR	P-3210	HPCU_htr
V-0436	DSEUsens_md	V-0144	XNTRVOLT	P-3211	N2H4_tnk_htr
V-0439	FMPsta_word1	V-0145	XNHTRCUR	P-3212	N2H4_svbm_ht
V-0445	DSEUh10DACIp	V-0146	XNHTRVOL	P-3213	N2H4_boon_ht
V-0498	SCAN_period	V-0147	XNTRCOMN	P-3214	N2H4_ln1_htr
V-0500	SCAN_skip	V-0188	XPAMSRD1	P-3215	N2H4_ln2A_ht
V-0501	FMP_period	V-0190	XPAMSRD2	P-3216	N2H4_ln2B_ht
V-0503	FMP_skip	V-2512	EGA1_pos	P-3217	N2H4_tc1_htr
V-0504	BURST_period	V-2522	EGA2_pos	P-3218	N2H4_tc2_htr
V-0506	BURST_skip	V-3402	XMAINFLOW	P-3219	IEM_SRU_htr
V-3025	dseuSCANdata	V-3403	XCATFLOW	P-3220	IPS_act1_htr
V-3026	dseuFMPdata	V-3404	XNEUFLOW	P-3221	IPS_act2_htr
V-3027	dseuBURSTdata	V-4063	ips_thr_tmp	P-3222	XPA_htr
D-0053	buf_pkt_03	V-4064	IpsThrMskTmp	P-3223	KAPA_htr
D-0054	sent_pkt_03	A-1401	acmSunBody0	P-3224	SADM_py_htr
D-0069	buf_pkt_04	A-1402	acmSunBody1	P-3225	SADM_ny_htr
D-0070	sent_pkt_04	A-1403	acmSunBody2	P-3226	Battery_htr
F-0380	PktOverflow	A-1711	sada_angle_0	P-3227	DSEU1_htr
F-0381	PktMsgCount	A-1712	sada_angle_1	P-3228	DSEU2_htr
D-0001	spc_used_tot	P-2040	sa1_i	P-3229	RSU_htr
B-0011	bmDSEUgdc dct	P-2050	sa2_i	V-0120	XPOWRLVL
B-0012	bmDSEUbdcdct	P-3006	pps_100w	V-3105	XTHRSSTS
V-0133	XACCLCUR	P-3200	XFS_shf1_htr		
V-0134	XACCLVOL	P-3201	XFS_shf2_htr		

APIDs 3 and 4


Appendix B. Date of Turn-on/off and Frequency of Data Capture

The IDS was first activated on day after launch on DOY 1998-298T22:05:43. Since the initial activation, IDS has been operated continuously, except during spacecraft safe-mode operations. Data collection from the IDS has occurred during all IPS operations (see IPS Appendix B) with data rates established by negotiation with the DS1 mission operations team. The IDS team supported development of IPS Acceptance Test sequences performed during the

mission to insure capture of critical data during these IPS validation activities. During DS1 cruise, higher data rates were supported for IPS ignition events, with reduced data rates during long-duration thrust segments. IDS was also operated during the Asteroid Braille fly-by, with IDS data rate and burst-mode commands integrated in the fly-by sequence. IDS data collection is anticipated to occur until the DS1 end-of-mission.

Appendix C. List of Acronyms and Abbreviations

AC	Alternating Current	LP0	Langmuir Probe (spherical)
ACS	Attitude Control System	LP1	Langmuir Probe (planar ring)
AU	Astronomical Unit	ML#	Mission Level (#)
CAL0	Calorimeter (Line-of-sight)	MLI	Multi-layer Insulation
CAL1	Calorimeter (Non-line-of-sight)	NDP	NSTAR Diagnostics Package
CEX	Charge-exchange Xenon	NSTAR	NASA SEP Technology Applications Readiness
DC	Direct Current	PCB	Printed Circuit Board
DOY	Day-of-year	PIC	Particle-in-cell
DS1	Deep Space One	PPU	Power Processor Unit
DSEU	Diagnostics Sensors Electronics Unit	PWA	Plasma Wave Antenna
EGA	Engine Gimbal Assembly	PWS	Plasma Wave Spectrometer
EMC	Electromagnetic Compatibility	QCM0	Quartz Crystal Microbalance (Line-of-sight)
EMI	Electromagnetic Interference	QCM1	Quartz Crystal Microbalance (non-line-of-sight)
EMT	Engineering Model Thruster	RCS	Reactive Control Subsystem
eEV	Electron Volt	RF	Radio Frequency
EWB	Environment Work Bench	RPA	Retarding Potential Analyzer
FGM_IB	Flux-Gate Magnetometer (Inboard)	RSU	Remote Sensors Unit
FGM_OB	Flux-Gate Magnetometer (Outboard)	SCM0	Search Coil Magnetometer (Miniature)
FMP	Fields Measurement Processor	SCM1	Search Coil Magnetometer (Science, Inactive)
IAT1	IPS Acceptance Test #1	SEM	Scanning Electron Microscopy
IAT2	IPS Acceptance Test #2	SEP	Solar Electric Propulsion
IDS	IPS Diagnostics Subsystem	SMA	Shape-memory Alloy
IPS	Ion Propulsion Subsystem	TUB	Technical University of Braunschweig
LDT	Life Demonstration Test	VDC	Volts Direct Current
LOS	Line-on-sight		



Ka-Band Solid-State Power Amplifier (KAPA) DS1 Technology Validation Report

Martin I. Herman, Luis R. Amaro, Chien-Chung Chen,
Gerald S. Gaughen, William A. Hatch, James S. Howard,
Andrew Makovsky, Kermit I. Pederson, Steven M. Petree,
Rocco P. Scaramastra, F. H. Taylor, Joseph D. Vacchione, Sam Valas
*Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California 91109*

Sam Valenti
*Lockheed Martin Corporation
Communications and Power Center
100 Campus Drive
Newton, Pennsylvania 18940*



Table of Contents

<u>Section</u>	<u>Page</u>
Ka-Band Solid-State Power Amplifier (KAPA) Fact Sheet	iii
Ka-Band Solid-State Power Amplifier (KAPA) DS1 Technology Validation Report	1
Abstract	1
1.0 Introduction	1
2.0 KAPA Description and Flight Qualification	1
3.0 KAPA Flight Validation	2
3.1 Validation Criteria	3
3.2 Validation Evaluation/Summary	3
4.0 Ka-band Technology	3
5.0 Summary and Conclusion	4
6.0 Acknowledgment	4
7.0 References	4
Appendix A. List of Telemetry Channels and Names	5
Appendix B. Date of Turn-on/off and Frequency of Data Capture	6

Figures

<u>Figure</u>	<u>Page</u>
Figure 1. DS1 Telecom Subsystem Block Diagram.....	2
Figure 2. Interior View of the +X, +Y Panel.....	3

Ka-Band Solid-State Power Amplifier (KAPA) Fact Sheet



Lockheed Martin’s Communications and Power Center (CPC) offers a complete line of communications products that includes a Ka-band Solid-State Power Amplifier (SSPA). Our standard Ka-band SSPA has greater than 2.5 W of output power and an overall efficiency of 14%.

The “plug-in” module approach allows the combination of multiple modules to obtain power output as high as 20 W. The SSPA is integrated with a high-efficiency electronic power conditioner (EPC) and consists of a three-stage radio frequency (RF) driver module and a three-stage RF output module. Input and output WR28 waveguide isolators are used for low voltage standing wave ratio (VSWR) and output module protection.

The RF output module combines three stages of amplification: stage one represents the basic building block of the entire output module. The RF output module (Figure 1) consists of a fully metalized diamond substrate that acts as a heat-dissipation path and carrier to the RF.

The block diagram of the SSPA is shown in Figure 2 with electrical performance at -14°C , $+23^{\circ}\text{C}$ and $+40^{\circ}\text{C}$. Figure 3 shows the key performance parameters versus temperature and frequency. Temperature compensation and telemetry circuitry are incorporated in the SSPA design, allowing for complete system integration. The SSPA is fully space qualified, and the physical design and layout have successfully passed hundreds of non-operational thermal cycles ranging from -55°C to $+125^{\circ}\text{C}$.




Figure 2. SSPA Block Diagram




Figure 1. RF Module




Figure 3. Key Performance Parameters Versus Temperature and Frequency

For further information on the SSPAs and the full CPC product line call: 215-497-1559, or Fax: 215-497-1564. Contact our web site at <http://www.payloads.com>. Lockheed Martin Communications and Power Center 100 Campus Drive Newtown, PA 18940

Ka-Band Solid-State Power Amplifier (KAPA) DS1 Technology Validation Report

*Martin I. Herman, Luis R. Amaro, Chien-Chung Chen, Gerald S. Gaughen, William A. Hatch
James S. Howard, Andrew Makovsky, Kermit I. Pederson, Steven M. Petree
Rocco P. Scaramastra, F.H. Taylor, Joseph D. Vacchione, Sam Valas*
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California

Sam Valenti
Lockheed Martin Corporation, Communications and Power Center, Newton, Pennsylvania

ABSTRACT

Communication subsystems for future missions must be low-mass and enable equivalent if not greater data return to the scientific community over the current X-band (8.4 GHz) links. One potential solution is to increase the downlink frequency to Ka-band (32 GHz). A major component required is the development of a power amplifier that can boost a transponder's exciter power from 0.5 mW_{rf} to more than 2W_{rf}. This paper describes the basic characteristics of a Lockheed Martin Engineering Test Module Ka-band solid state power amplifier (SSPA) that was provided to the New Millennium Program for flight validation on the Deep Space 1 (DS1) mission. Initial in-flight data shows that the unit has been functioning nominally during the past year (1680 hours of operation accumulated). In addition, the unit has been power cycled 28 times and has gone through multiple-thermal cycles (due to the trajectory combined with autonomous spacecraft maneuvers for optical navigation measurements).

1.0 INTRODUCTION

The Ka-band Solid-State Power Amplifier (KAPA) is one of eight Level-1 technology validation objectives of the New Millennium Deep Space 1 (DS1) mission. The principal goal of the New Millennium Program (NMP) is to validate selected high-risk, high-benefit technologies to reduce the risks and costs future missions would experience in their use. With successful flight validation of the technology, the risk of using them is substantially reduced. Knowledge gained from incorporating the new capability into the spacecraft, ground system, and mission design sets a beneficial precedent for future missions.

KAPA was developed by Lockheed Martin Communications and Power Center under their own internal IR&D funding. An Engineering Test Module Unit was delivered to DS1 and integrated into the Telecommunication subsystem.

This unit has successfully demonstrated the highest-power solid-state Ka-band amplifier ever used for deep space communications. With future improvements in ground facilities and spacecraft hardware, Ka-band holds a potential four-fold increase in data rate in comparison with X-band. This is extremely important since a faster data rate reduces ground resources/mission operation support and project cost. Another benefit of going to Ka-band is the availability of greater bandwidth. Both NASA and commercial programs recognize this and are developing the technology to move beyond microwave bands, which are becoming crowded due to PCS and other emerging information technology ventures.

2.0 KAPA DESCRIPTION AND FLIGHT QUALIFICATION

KAPA's mass was 0.66 kg (this includes input/output isolators, power supply, telemetry circuitry, and RF electronics), with a RF output power of 2.2 W and a gain of 36 dB.

The unit was qualified to DS1 requirements that include:

- Random Vibration:

20 Hz	0.0322 G ² /Hz
50–500 Hz	0.2 G ² /Hz
2000 Hz	0.0126 G ² /Hz
Overall	13 G _{rms}
- Thermal Vacuum cycling from –14° C to +40° C
- Full EMC testing to MIL SPEC 461

Unique features include built-in input/output isolators and engineering telemetry monitors (two-gate currents, output drain voltage, and internal unit temperature). Due to the short development time for this unit, Lockheed Martin did not hermetically seal it. After delivery, some accelerated testing on other similar power devices has shown no major degradation after an initial burn-in. After 250 hours of ground operation (in both vacuum and atmosphere), the flight unit did not show any operational degradation. Caution was exercised

to prevent the unit's operation for too long in open atmosphere or letting the unit's temperature drop below the dew point.

The key technology for KAPA is the use of 0.25 micron GaAs Pseudomorphic High Electron Mobility Transistors (PHEMT). The efficiency could have been optimized further with the use of 0.15 micron devices; however, time and resources defined what the final product would be in this fast-paced program.

3.0 KAPA FLIGHT VALIDATION

The telecommunication subsystem for DS1 was single string, as mandated by the project. Figure 1 is the DS1 telecom subsystem block diagram. The primary communication link is on Channel 19 at X-band for both uplink and downlink (7.168 GHz and 8.421 GHz, respectively).

As part of the technology demonstration, we have an auxiliary Ka-band downlink (32.155 GHz). The heart of the Ka-band downlink is the KAPA itself. Figure 2 is an

interior view of the +X, +Y panel of the DS1 spacecraft, where a major portion of the active telecom subsystem electronics resides. Key components include the KAPA, a Detector Amplifier Module, and an SDST. KAPA dimensions are approximately 14.2 cm × 15.2 cm. The full Telecommunication subsystem was described in [1].

On December 9, 1998, the KAPA was first powered on in-flight (launch of DS1 was on October 24, 1998). Flight operation of the unit has been nominal. As of November 22, 1999, the unit has been power cycled more than 28 times and has logged more than 1680 hours of operating time (over a variety of temperature ranges). In the event that the Ka-band operation was not nominal, it was the responsibility of the flight team to ensure that enough data was available to determine what the anomaly could have been. This was accomplished both internally and externally to the KAPA itself. Internal to the unit—temperature sensor, gate-currents and gate-voltage telemetry are passed to the C&DH subsystem. External to the unit, RF power detectors monitor KAPA's input and output RF power. This ensures that the RF drive from the SDST—or any intervening component—is not responsible for any potential performance degradation.

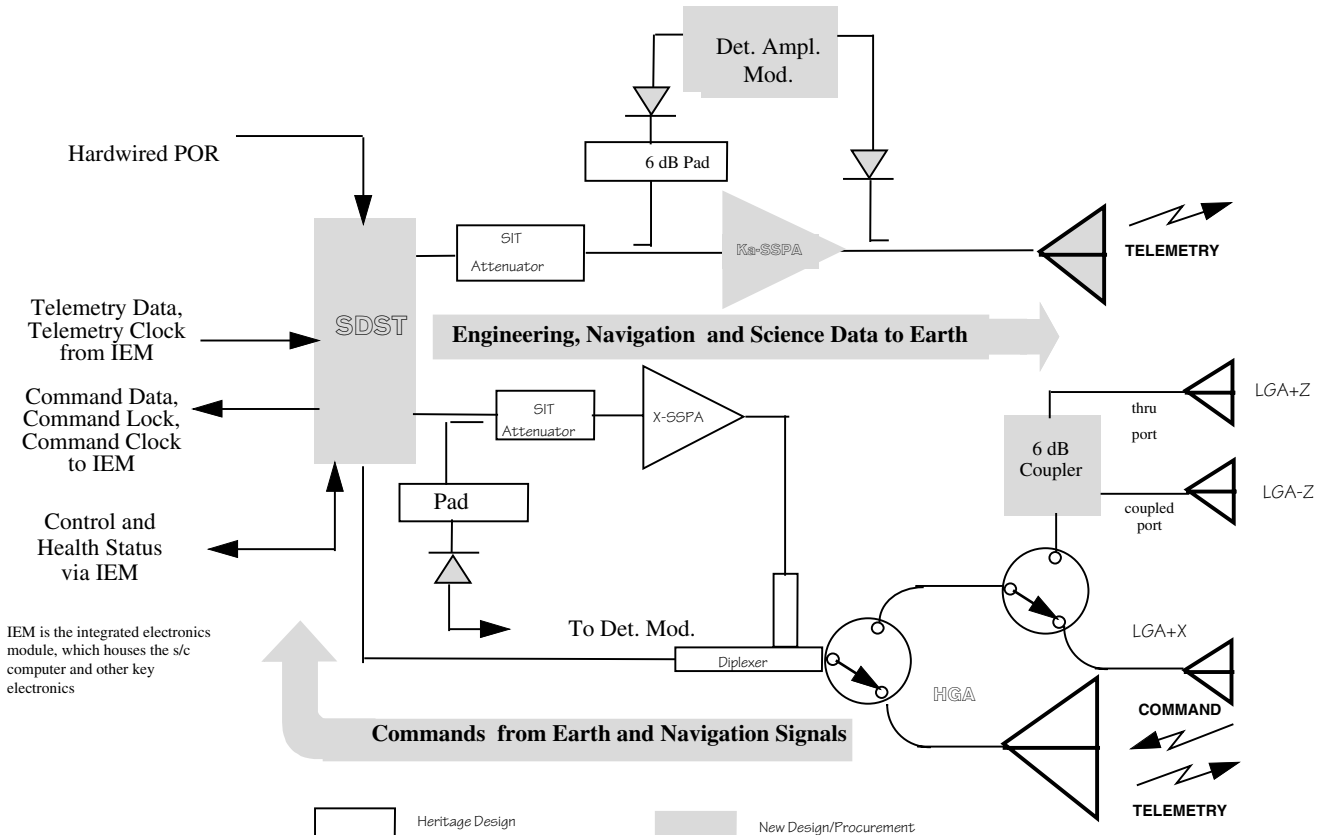


Figure 1. DS1 Telecom Subsystem Block Diagram

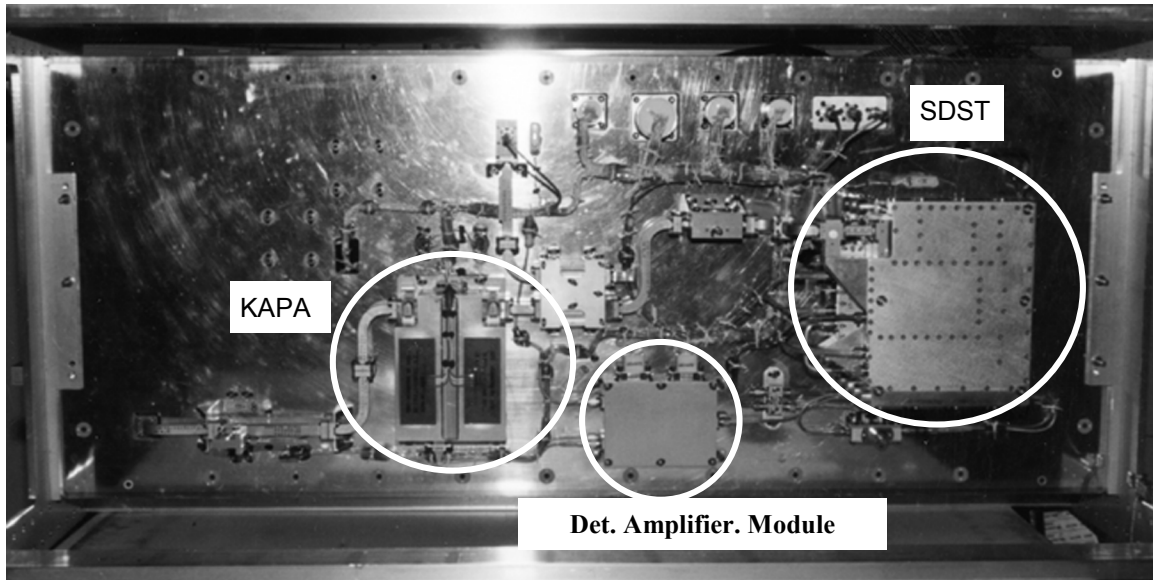


Figure 2. Interior View of the +X, +Y Panel

3.1 Validation Criteria

Pre-Flight:
Development of a 2.5 W _{RF} SSPA which has ~36 dB gain and provides critical engineering telemetry (gate current, drain voltage, and unit temperature) for unit-performance evaluation during flight.
Post Flight:
Launch to L+25 day Due to mission-pointing constraints for the Miniature Imaging Camera and Spectrometer (MICAS), a Ka-band communication link is not allowed during this period. > L+25 day The ability to have a Ka-band communication link is a major validation step.

3.2 Validation Evaluation/Summary

Pre-Launch		
Parameter	Achieved	Benchmark (MGS Mission)
Mass*	0.660 kg	>0.600 kg (and does not have input isolator)
RF Output Power	2.2 W	1 W
Efficiency*	13%	8.7%
Gain	36.4 dB	15 dB

Post-Launch (>L+25)	
Parameter	Achieved
RF Output Power	2.16 W
Efficiency*	12%
Gain	36 dB
*including DC-DC converter	

The unit's overall performance has been nominal. Analysis of data has been of the first order. Mainly, the Telecom Mission Operator plots data looking for any potential hazardous trends and ensures that it is within nominal operating range. A potential output-power step change was observed within the measurement-error range (0.5 dB); however, no visible trend is now apparent.

4.0 KA-BAND TECHNOLOGY

The desire to increase the data volume of future systems can be accomplished by going from X-band (8.4 GHz) to Ka-band (32 GHz). Theoretically, there is a 16-fold advantage. When one takes into account the realities of weather, spacecraft pointing, etc., the potential advantage is predicted to be a factor of 4. The KAPA is a major component required in achieving this important goal. The question now may arise, What does it take to have a Ka-band link? The downlink telemetry is modulated onto a subcarrier and then up-converted to Ka-band in the Small Deep Space Transponder (SDST). From the SDST, the signal may be coupled off to detectors or go directly to the power amplifier to increase signal strength. The KAPA provides this critical function (see Figure 1). From the amplifier, the signal can be routed through

couplers and/or switches to the antenna, where it is radiated into free space.

Collecting all the facts presented thus far:

- Ka-band may enable greater science-data return.
- DS1 has validated operation of the highest-power solid-state Ka-band amplifier for Deep Space Communications.

This begs the question, Would we achieve a potential advantage for Ka-band communications? Initial results from [2] indicate that, based on scaled calculations from DS1 flight data, future systems could achieve the four-fold improvement.

5.0 SUMMARY AND CONCLUSION

DS1 has successfully demonstrated in-flight the operation of a Ka-band (32 GHz) Solid-State Power Amplifier (KAPA), which was an Engineering Test Module Unit provided by Lockheed Martin Communication and Power Center (using their own IR&D funding). This technology, in turn, has enabled further validation of Ka-band's potential advantage over X-band for deep space communications.

6.0 ACKNOWLEDGMENT

KAPA was developed by Lockheed Martin Communications and Power Center under their own internal IR&D funding. An Engineering Test Module Unit was delivered to DS1 and integrated into the Telecommunication subsystem.

Many people have contributed to the success of KAPA and its technology validation activities. The authors would like to acknowledge the following:

From Lockheed Martin CPC:

Mark Karnacewicz
Sandy Conway
Mitch Hirokawa
Bob Novack
Todd Rena
Lew Sponar
William Taft III
Ken Vaughn
Larry Newman

Operation personnel from the DS1 Flight and Mission Support team have labored extensively over the planning and execution of the tech validation activities. The authors will like to acknowledge the participation of the following individuals:

Pam Chadbourne
Kathy Moyd
Marc Rayman
Rob Smith
Ben Toyoshima

Without their support the technology validation and characterization effort of the KAPA would not have been successful.

The work described herein was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

7.0 REFERENCES

- [1] Herman, M. I. et. al., "Deep Space One Telecommunication Development," AIAA/Utah State University Conference on Small Satellites, Logan, Utah, 1998.
- [2] Chen, C.-C., et. al., "DS1 Technology Validation Report for Small Deep Space Transponder," Deep Space 1 Technology Validation Symposium, Pasadena, CA, February 8–9, 2000.

Appendix A. List of Telemetry Channels and Names

Table A1 is a list of all of the telemetry channels that the "monitor" channels in this work. (Jim Taylor, 10/20/99.) KAPA team collects and uses. Note the importance of

Table A1. Channels and Mnemonics

Channel	Mnemonic
T-3252	sdst_evnt_ct
T-3116	aux_osc_temp
T-3124	vco_tmp
T-4003	KAPA_ext_tmp
P-2061	ess_bus_v
A-1637	bbc_CtrlErr0
A-1621	bbc_CtrlErr1
A-1625	bbc_CtrlErr2
T-3502	kapa_dc_pwr
T-3503	kapa_rf_gain
T-3188	ka_tlm_mod
T-3136	katlm_coder
T-3180	ksubcar_freq
T-3101	ka_ranging
T-3105	ka_Exciter
P-3126	KAPA_on_off
M-0130	MCD1 SNR
M-0731	AA5 SS1 SNR
M-0723	AA5 PCN0
M-0727	AA5 PC
M-0725	AA5 SNT
M-0737	AA5 SPE
M-0304	ANT A EL ANG
M-0305	ANT A AZ ANG
M-0308	A CNSCN
M-0309	A CNSCN LOOP


Appendix B. Date of Turn-on/off and Frequency of Data Capture

The KAPA was first turned ON as part of a telecom technology validation activity on 98-343/08:15. It was left ON for over 30 hours, being commanded OFF at 98-344/14:41. Both times are per the ACE log.

The Ka-band downlink has been ON and OFF many times since then. Table B1 is a listing, as obtained from a telemetry query of the KAPA state itself, through 99-280. (Jim Taylor, 10/29/99.)

Table B1. Channels and Mnemonics

Time	KAPA State	Time	KAPA State
1999-011T01:30:38.465	ON	1999-068T22:00:17.339	ON
1999-018T20:12:56.500	OFF	1999-075T06:40:38.551	OFF
1999-018T23:40:06.500	ON	1999-082T02:55:29.682	ON
1999-020T19:14:12.305	OFF	1999-082T12:44:13.950	OFF
1999-020T23:40:12.305	ON	1999-088T20:30:12.579	ON
1999-022T20:04:17.765	OFF	1999-089T04:40:32.570	OFF
1999-022T23:40:07.765	ON	1999-095T23:25:32.828	ON
1999-026T21:44:04.176	OFF	1999-096T09:00:12.904	OFF
1999-026T23:00:24.176	ON	1999-102T22:40:33.763	ON
1999-031T23:23:34.676	OFF	1999-103T05:50:12.513	OFF
1999-032T23:00:22.289	ON	1999-109T22:55:32.914	ON
1999-041T21:11:11.398	OFF	1999-110T05:25:12.943	OFF
1999-043T20:17:46.344	ON	1999-116T20:55:13.383	ON
1999-053T23:20:08.293	OFF	1999-117T04:15:33.301	OFF
1999-054T04:30:44.219	ON	1999-166T20:30:11.570	ON
1999-054T19:01:56.387	OFF	1999-175T12:01:01.508	OFF
1999-055T19:56:42.246	ON	1999-209T14:50:39.103	ON
1999-057T17:55:15.735	OFF	1999-209T19:28:59.126	OFF
1999-058T00:30:15.981	ON	1999-221T19:45:36.647	ON
1999-060T10:20:16.226	OFF	1999-222T03:14:36.659	OFF
1999-060T16:00:16.103	ON	1999-242T20:45:24.536	ON
1999-061T14:55:16.145	OFF	1999-243T06:29:37.526	OFF
1999-061T23:00:17.038	ON	1999-256T18:45:36.736	ON
1999-064T09:55:17.242	OFF	1999-256T23:46:11.708	OFF
1999-064T16:00:17.257	ON	1999-277T21:08:37.596	ON
1999-067T09:55:17.371	OFF	1999-278T05:34:24.616	OFF
1999-067T15:05:17.295	ON		
1999-068T14:50:17.326	OFF		



Low Power Electronics DS1 Technology Validation Report Flight Qualifying Sub-0.25- μm Fully Depleted SOI CMOS Technology

Richard P. D'Onofrio — donofrio@ll.mit.edu

Dr. Craig L. Keast — keast@ll.mit.edu

Antonio M Soares — soares@ll.mit.edu

Dr. Peter W. Wyatt — wyatt@ll.mit.edu

MIT Lincoln Laboratory

244 Wood Street

Lexington, MA 02420-9108

Phone: (781) 981-7880

FAX: (781) 981-7889




Table of Contents

<u>Section</u>	<u>Page</u>
Extended Abstract	iv
Fact Sheet	v
1.0 Introduction	1
2.0 Technology Description	1
2.1 What Is It?.....	1
2.2 Key Validation Objectives	3
2.3 Expected Performance Envelope	3
2.4 Detailed Test Description.....	3
2.5 Technology Interdependencies.....	4
2.6 Testing	4
2.7 Test Result Comparison.....	6
3.0 Technology Validation Summary	16
4.0 Technology Application for Future Missions	16
5.0 Acknowledgments	16
6.0 List of References	17

Figures

<u>Figure</u>	<u>Page</u>
Figure 1. Schematic Cross Section of <0.25- μ m FDSOI n-channel and p-channel Transistors.....	1
Figure 2. Ring-oscillator Stage Delay vs. Voltage with Fanout = 1 for the 0.25- μ m FDSOI Technology.....	2
Figure 3. Total-Dose Test Results for an 8.0- μ m/ 0.25- μ m n-Channel Device.....	2
Figure 4. Photograph of the Low-power Experiment 6u VME-style Test Board	3
Figure 5. Transistor Test Schematic	3
Figure 6. Transistor/Ring Oscillator Test Overview.....	4
Figure 7. LPE Board Overview	5
Figure 8. N-Channel Transistor (V_{th}); ~2-mV A/D Converter Resolution (left-hand Y-axis).....	7
Figure 9. N-channel Transistor (Leakage1)	7
Figure 10. N-channel Transistor (Leakage2)	8
Figure 11. N-channel Transistor (Leakage3)	8
Figure 12. N-channel Transistor (Drive Current1).....	9
Figure 13. N-channel Transistor (Drive Current2).....	9
Figure 14. N-channel Transistor (gmSat)	10
Figure 15. N-channel Transistor (Gds).....	10
Figure 16. P-channel Transistor (V_{th}); ~2-mV A/D Converter Resolution (left-hand Y-axis)	11
Figure 17. P-channel Transistor (Leakage1)	11
Figure 18. P-channel Transistor (Leakage2)	12
Figure 19. P-channel Transistor (Leakage3)	12
Figure 20. P-channel Transistor (Drive Current1)	13
Figure 21. P-channel Transistor (Drive Current2)	13
Figure 22. P-channel Transistor (gmSat).....	14
Figure 23. P-channel Transistor (Gds).....	14
Figure 24. L = 0.25- μ m 97-Stage Ring Oscillator 1	14
Figure 25. L = 0.25- μ m 97-Stage Ring Oscillator 2	15
Figure 26. L = 0.25- μ m 97-Stage Ring Oscillator 3	15
Figure 27. L = 0.25- μ m 97-Stage Ring Oscillator 4	16

Appendices

<u>Appendix</u>	<u>Page</u>
Appendix A: Telemetry Channels	18
Appendix B: Date of Turn-on/Frequency of Data Capture	19

EXTENDED ABSTRACT

The future of deep-space exploration is dependent on the research and development of new technologies that will allow designers to build low-power, lightweight space systems and peripherals. The focus of this technology validation experiment is to characterize the effects of the space environment on a DARPA-sponsored, sub-0.25- μm , fully depleted silicon-on-insulator (FDSOI) complementary metal-oxide semiconductor (CMOS) technology developed at MIT Lincoln Laboratory. FDSOI technology offers the advantage of providing high performance (>1 GHz operation) from a sub-2.0-V power supply. The resulting reduction in power consumption (~ 5 times less power than the corresponding 0.25- μm bulk CMOS technology), coupled with the SOI technology's inherent resistance to latchup, make this an attractive choice for the design of integrated circuits used in hardware systems for deep-space exploration. In addition, the increased transistor-packing densities realized with SOI technology allow for the fabrication of smaller, lighter, higher-performance devices.

A first step towards validating the sub-0.25- μm FDSOI process as a key technology for deep-space application lies in the examination and analysis of FDSOI behavior at the transistor level. The collection of key parametric data from the measurement of 8.0- $\mu\text{m}/0.25\text{-}\mu\text{m}$ n-channel and p-channel transistors will serve as a sound predictor for how well circuits developed with this technology will perform in space.

One of the major risks associated with electronic technologies in the space environment is operational failure due to total dose radiation. Our approach with the Low Power Flight Experiment (LPE) is to observe the properties of test devices where no attempt has been made in either processing or packaging to optimize performance for the radiation environment. We are instead interested in characterizing the sub-0.25- μm FDSOI baseline process developed at MIT Lincoln Laboratory and verifying that the inherent radiation-hardened qualities of the technology that have been examined through ground testing hold true in the space environment.

FDSOI parametric testing on Deep Space 1 (DS1) is performed by a board designed to emulate the tasks of a semiconductor parameter analyzer. The sub-0.25- μm FDSOI test chip is mounted on this board. All board components, with the exception of the test chip, are radiation-hardened so that all changes in behavior can be isolated to the test chip.

After nearly one year in space, the technology continues to function, yielding parametric data that is very similar to data taken before the launch. The total ionizing dose (TID) exposure of the test chip has had very little effect on function and performance.

The results of the LPE have shown that transistor characteristics and performance are minimally affected by the space environment. This insight into the fundamental building blocks of circuit design will prove to be invaluable when creating more complex SOI test circuits for further space qualification.

Low Power Flight Experiment

FACT SHEET

What Is It?

The Low Power Flight Experiment was designed to monitor and record key operating parameters of sub-0.25- μm , fully depleted silicon-on-insulator (FDSOI) CMOS (complementary metal oxide semiconductor) test devices.

Why Is It Exciting Technology?

The 0.25- μm FDSOI process developed at MIT Lincoln Laboratory has some key advantages over circuits developed in bulk process.


- Low-power operation, 1.0-V supply, \sim 0.3-V threshold.
- Fully depleted device design for reduced parasitic capacitance and near ideal subthreshold swing.
- This “no well,” mesa-isolated island technology allows for increased packing densities with no bulk CMOS latchup.
- Further device scaling realized through the use of the world’s most advanced optical lithographic technology and techniques.
- High performance; 3.9-GHz operation demonstrated.
- Resistant to single-event upset (SEU); more tolerant to total ionizing dose (TID).

Where Are Some Important Applications?

- Space electronics.
- Wireless communication.
- Mobile computing.

When Will It Be Demonstrated?

The technology has just completed the first phase of space qualification through testing onboard Deep Space 1, the first launch of the New Millennium Program.



Sub-0.25- μm , Fully Depleted Silicon-on-Insulator Technology

Points Of Contact

Dr. Craig L. Keast, *Group Leader*
Email: keast@ll.mit.edu
Phone: 981-781-7884

Dr. Peter W. Wyatt, *Associate Group Leader*
Email: wyatt@ll.mit.edu
Phone: 981-781-7882

MIT Lincoln Laboratory
Advanced Silicon Technology Group
244 Wood Street
Lexington, MA 02420-9108
Fax: 981-781-7889

Low Power Electronics DS1 Technology Validation Report Flight Qualifying Sub-0.25- μm , Fully Depleted SOI CMOS Technology

Richard P. D'Onofrio, Craig L. Keast, Antonio M. Soares, Peter W. Wyatt
Lincoln Laboratory, Massachusetts Institute of Technology
244 Wood Street, Lexington, MA 02173-9108
(781) 981-2573

1.0 INTRODUCTION

The new millennium brings with it exciting technical challenges in the area of circuit design. Deep-space travel, wireless communication, and mobile computing are just a few examples of important applications that demand a core of low-power, high-performance electronic components. The design rule constraints of bulk complementary metal-oxide semiconductor (CMOS) technology have limited just how far researchers can go in reducing power consumption while maintaining and improving performance.

Fully depleted silicon-on-insulator (FDSOI) technology promises to be an important area of research for the continued advance of low-power, high-performance electronics. The combination of transistor mesa-island isolation along with the feature-size scaling available through advances in lithographic equipment and techniques, allow for the design of smaller, faster, lower-power circuits. In addition, the technology's inherent resistance to latchup makes it particularly attractive for space application. The New Millennium Program has provided the opportunity to begin the process of flight-qualifying this technology for deep-space application.

2.0 TECHNOLOGY DESCRIPTION

2.1 What Is It?

2.1.1 Fully Depleted 0.25- μm SOI Technology—Figure 1 shows a schematic cross section of an n-channel and p-channel transistor fabricated in the 0.25- μm FDSOI CMOS technology. The starting silicon active layer thickness is thinned to 61 nm by thermal oxidation. After processing, the final active-area thickness is approximately 50 nm. Device isolation is by mesa-etching followed by sidewall-oxidation. The 10 nm of SiO_2 topped by 100 nm of Si_3N_4 are patterned and plasma-etched along with the silicon layer. A 45° sidewall “channel stop” implant is followed by a 40-nm sidewall oxidation. After island doping by implantation through a 7-nm sacrificial oxide, a 7-nm gate oxide is grown at 850° C followed by a 225-nm undoped amorphous Si gate deposition. The gate electrode is then patterned, plasma etched and reoxidized at 800° C. A medium-doped drain implant is followed by a 120-nm spacer oxide deposition and etch followed by a source/drain implant. The source/drain implant, which also dopes the polysilicon gate electrodes, is activated with a 950° C, 30-s annealing.

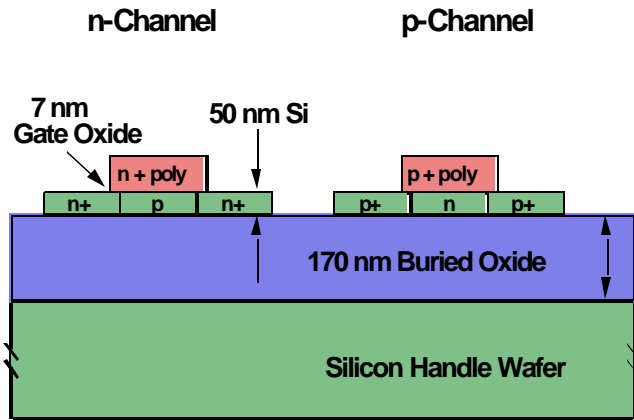


Figure 1. Schematic Cross Section of $<0.25\text{-}\mu\text{m}$ FDSOI n-Channel and p-Channel Transistors

A titanium-capped cobalt salicide process contacts the 50-nm-thick silicon regions and straps the p+, n+, and undoped polysilicon gates [1]. The back end consists of a fully planar, three-level, metal interconnect process that incorporates damascene hot aluminum plugs at contacts, via 1 and via 2, and chemical mechanical polished (CMP) plasma-enhanced chemical vapor deposition (PECVD) tetraethylorthosilicate (TEOS) oxide intermetal dielectric.

Figure 2 shows the inverter stage delay vs. power supply voltage for a 97-stage ring oscillator fabricated in the FDSOI technology. This process results in a 25-ps stage delay at 2 V, and when clocked at the same level of performance as a 2.5-V, 0.25- μm bulk CMOS technology, the FDSOI CMOS offers a five-times-less reduction in power.

To date 85 different digital, analog, and mixed-mode circuits have been fabricated in this technology as part of the DARPA-funded low-power, high-performance multiproject-run research fabrication service at MIT Lincoln Laboratory. Typical digital operation is from 600 mV to 2 V for this 400-mV threshold technology, with clock speeds generally over 100 MHz at 1.0 V and in excess of 1 GHz with a 2-V power supply [2]. The 0.25- μm FDSOI CMOS technology has been used to fabricate a data generation/acquisition process-benchmarking test circuit. The 0.25- μm FDSOI technology had similar performance to the Vitesse

H-GaAs-3 process technology (950 MHz vs. 1 GHz operation). However, the FDSOI circuit consumed 45 times less power than the GaAs circuit (43 mW vs. 2 W).




Figure 2. Ring-oscillator Stage Delay vs. Voltage with Fanout = 1 for the 0.25-µm FDSOI Technology. Also shown is the data point for a commercially available 0.25-µm, 2.5-V bulk CMOS technology.

2.1.2 Preliminary Radiation Performance—The 0.25-µm FDSOI CMOS process was designed for low-power, high-performance operation. Radiation characteristics of the process were not critical design parameters during the process development cycle; i.e., nothing was done to optimize radiation performance. However, given that the process uses thin gate oxides (7 nm), has fully oxide-isolated transistors, and is not susceptible to parasitic bipolar latch-up (no wells), there is the potential for good total dose radiation resistance.

In order to get a baseline on the total-dose-radiation performance, testing was performed on an ARACOR Model 4100 Semiconductor Irradiation System with an 10-keV X-ray source. The dose rate was 10 krad (Si) per minute for 0 to 200 krad and 130 krad (Si) per minute for 200 to 1000 krad. The devices were measured immediately after irradiation. Figure 3a shows the I_d vs. V_{gs} curves for an 8-µm/0.25-µm n-channel device biased with 1.0 V on the gate and 0.0 V on the source, drain, and substrate. The threshold shift was ~130 mV after 1 Mrad (Si). Figure 3b shows the same device; however, in these curves the channel of the transistor from the radiation-induced measurement was performed with the addition of a -30-V substrate-wafer bias. This -30-V wafer bias accumulates the back channel

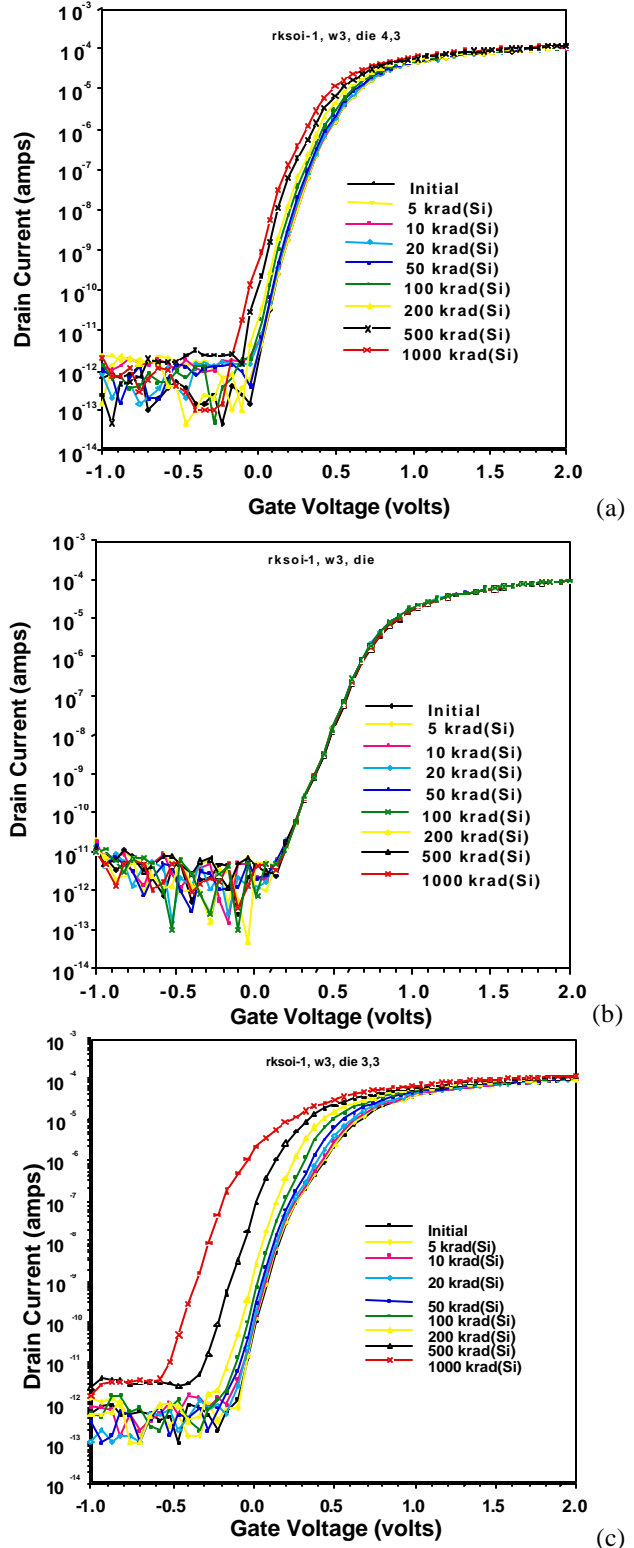


Figure 3. Total-Dose Test Results for an 8.0-µm/0.25-µm n-Channel Device: (a) biased in the “on” state during irradiation, (b) “on” state bias measured with -30 V on the substrate (c) “pass-gate” bias

of the transistor, effectively shielding the front charging effects of the buried oxide. Under these bias conditions the IV curves do not shift after radiation treatment, indicating that the radiation-induced shift in Figure 3a is the result of buried oxide charge. Figure 3c shows a “pass gate-biased” device with 1.0 V on the drain, 0.0 V on the source, gate, and substrate-measured with 0 V on the substrate. For this bias condition the threshold shift is ~500 mV at 1 Mrad (Si).

2.2 Key Validation Objectives

The key objective of the Low Power Flight Experiment (LPE) is to monitor changes in FDSOI device characteristics over the course of the Deep Space 1 (DS1) mission, and to correlate those changes with total-dose-radiation measurements sampled at the time the experiments were performed.

2.2.1 N-channel Transistor Characteristic Measurements—

1. Threshold Voltage
2. Drain-Source Leakage ($V_{ds} = 1.0\text{ V}$, $V_{gs} = 0.0\text{ V}$)
3. Drain-Source Leakage ($V_{ds} = 2.0\text{ V}$, $V_{gs} = 0.0\text{ V}$)
4. Drain-Source Leakage ($V_{ds} = 2.0\text{ V}$, $V_{gs} = -0.5\text{ V}$)
5. Drive Current ($V_{ds} = 1.0\text{ V}$, $V_{gs} = 1.0\text{ V}$)
6. Drive Current ($V_{ds} = 2.0\text{ V}$, $V_{gs} = 2.0\text{ V}$)
7. Saturation Transconductance
8. Drain-Source Output Conductance.

2.2.2 P-channel Transistor Characteristic Measurements—

1. Threshold Voltage
2. Drain-Source Leakage ($V_{ds} = -1.0\text{ V}$, $V_{gs} = 0.0\text{ V}$)
3. Drain-Source Leakage ($V_{ds} = -2.0\text{ V}$, $V_{gs} = 0.0\text{ V}$)
4. Drain-Source Leakage ($V_{ds} = -2.0\text{ V}$, $V_{gs} = 0.5\text{ V}$)
5. Drive Current ($V_{ds} = -1.0\text{ V}$, $V_{gs} = -1.0\text{ V}$)
6. Drive Current ($V_{ds} = -2.0\text{ V}$, $V_{gs} = -2.0\text{ V}$)
7. Saturation Transconductance
8. Drain-Source Output Conductance.

2.2.3 Performance—In addition to monitoring key transistor properties, the LPE also addresses the issue of performance monitoring. By sampling the output frequency of four 97-stage ring oscillators, we can evaluate how stage delay is affected by the space environment.

2.3 Expected Performance Envelope

It is expected that the sub-0.25- μm FDSOI transistor properties, as well as ring-oscillator performance, will be minimally affected by exposure to the total dose radiation seen by the spacecraft.

2.4 Detailed Test Description

2.4.1 Overview—In order to begin the process of space qualification, MIT Lincoln Laboratory fabricated a test integrated circuit consisting of n-channel and pchannel transistors as well as a group of 97-stage ring oscillators.

This low-power test chip was integrated into a test system that was designed to periodically monitor and record any changes in the basic characteristics of the transistors as well as evaluating changes in switching speed by sampling ring oscillator output frequencies as they are exposed to the space environment. The output of dosimeter and temperature-sensing circuits are sampled and recorded at each step of the test sequence to correlate the effects of thermal variation and total dose radiation.

The test system is fabricated on a 6u VME-style board using radiation-hardened components (Figure 4). The board is a category 3 experiment attached to a non-essential bus of the DS1 via a dual-redundant 1553B interface, with a Boeing SMARTIO integrated circuit being used as the protocol controller.




Figure 4. Photograph of the Low-power Experiment 6u VME-style Test Board

2.4.2 Device Testing—A series of MOS transistor measurements are made through the independent control of the gate, drain, and source nodes of each transistor, with connectivity achieved through a low noise, low leakage, programmable switching matrix (Figure 5). Programmable

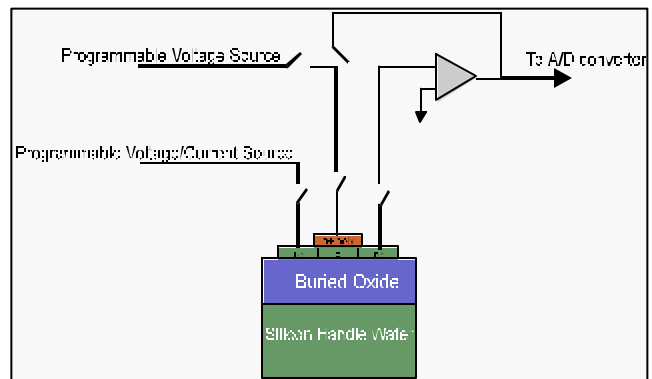


Figure 5. Transistor Test Schematic

voltage and current sources supply the “transistor under test” with appropriate bias conditions for the measurement being performed. Transistor characteristics such as threshold, conductance, and leakage are sampled and processed by an A/D converter. The SMARTIO ASIC is equipped with frequency-to-digital conversion capability, providing an accurate means of sampling ring oscillator output frequency. (See Figure 6 and Figure 7).

The spacecraft begins an LPE test by configuring the SMARTIO ports and sending a “begin” instruction. At this point, experiment control is transferred to the LPE onboard sequencer that cycles through the test instructions stored in ROM. The instructions for any given test set up the appropriate switch and voltage/current configurations. Results from all experiments are stored in onboard memory along with dosimeter and temperature information, where they are then transferred to the spacecraft’s solid-state recorder for transmission down to Earth.

2.5 Technology Interdependencies

The LPE is “piggybacked” with the power activation and switching module (PASM), to which the LPE supplies power. In addition, the LPE’s Boeing SMARTIO protocol

controller provides the communication conduit between the PASM and the spacecraft.

2.6 Testing

The LPE monitors the following eight key transistor parameters that will provide insight into the health of the devices.

1. Threshold Voltage—Transistor “turn-on” voltage defined as $V_{gs} @ I_{ds} = W/L * 0.1 \mu A$.
2. Drain-Source Leakage1—Transistor subthreshold leakage with 0.0 V applied to gate, 1.0 V (– polarity for p-channel devices) applied to drain, source grounded.
3. Drain-Source Leakage2—Transistor subthreshold leakage with 0.0 V applied to gate, 2.0 V (– polarity for p-channel devices) applied to drain, source grounded.
4. Drain-Source Leakage3—Transistor drain diode leakage with –0.5 V (+ polarity for p-channel devices) applied to gate, 2.0 V (– polarity for p-channel devices) applied to drain, source grounded.
5. Drive Current1—Transistor current drive capability with 1.0 V (– polarity for p-channel devices) applied to gate, 1.0 V (– polarity for p-channel devices) applied to drain, source grounded.

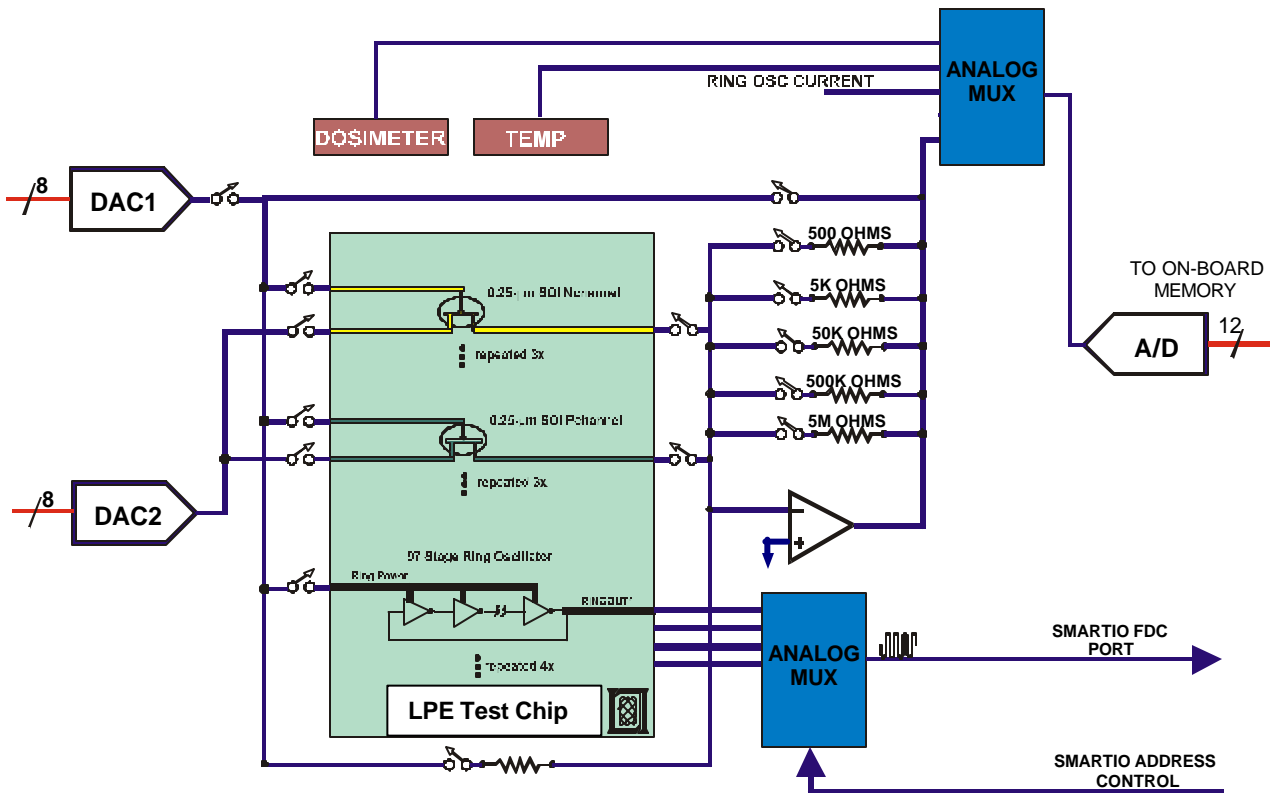


Figure 6. Transistor/Ring Oscillator Test Overview

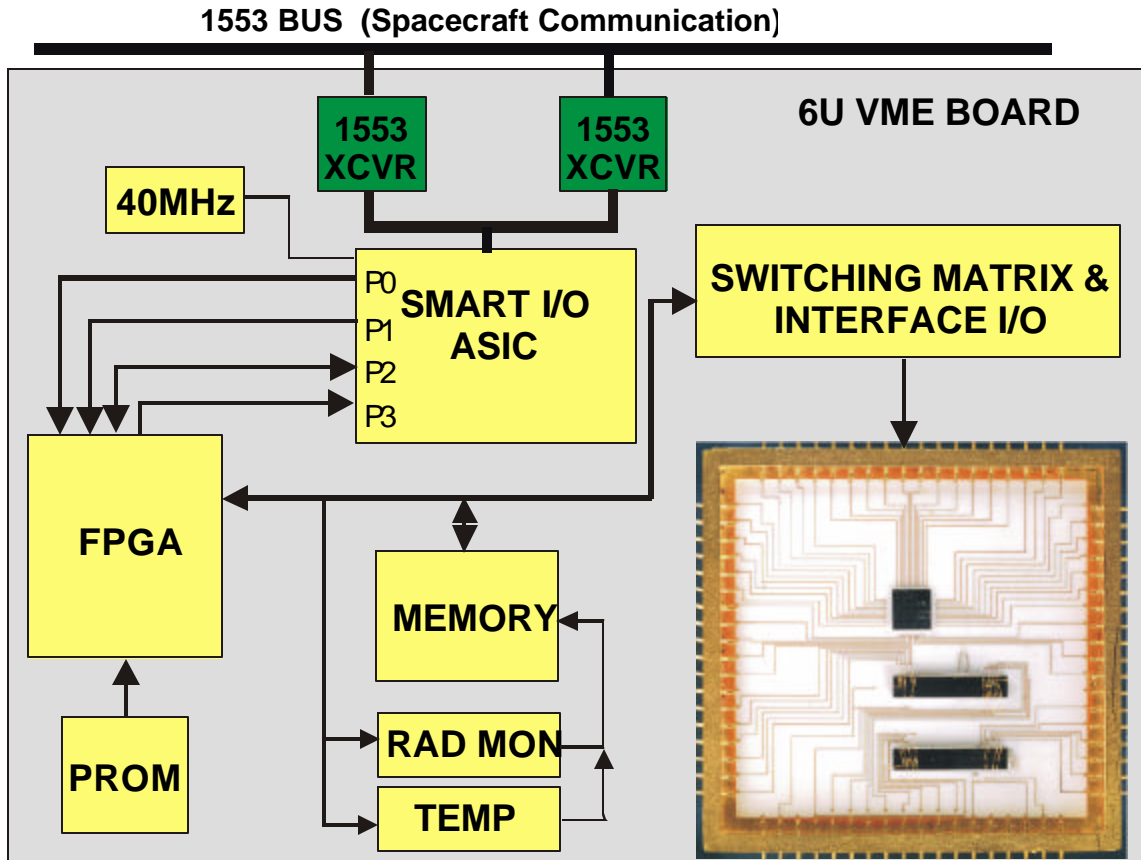


Figure 7. LPE Board Overview

6. Drive Current₂—Transistor current drive capability with 2.0 V (– polarity for p-channel devices) applied to gate, 2.0 V (– polarity for p-channel devices) applied to drain, source grounded.
7. Saturation Transconductance—The effect of gate voltage on output current in the saturation region at $V_{ds} = V_{gs} = 2.0$ V (– polarity for p-channel devices).
8. Drain-Source Output Conductance—Transistor channel conductance in saturation at $V_{ds} = V_{gs} = 2.0$ V (– polarity for p-channel devices).
4. (Lkg3) Drain-Source Leakage ($V_{ds} = 2.0$ V, $V_{gs} = -0.5$ V) = 13.4 nA
5. (Drv1) Drive Current ($V_{ds} = 1.0$ V, $V_{gs} = 1.0$ V) = 1.0 mA
6. (Drv2) Drive Current ($V_{ds} = 2.0$ V, $V_{gs} = 2.0$ V) = 2.9 mA
7. (Gsat) Saturation Transconductance = 1555 μ S
8. (Gds) Drain-Source Output Conductance = 123 μ S.

The LPE also monitors the output frequency of four 97-stage ring oscillators at a 2.0-V power supply. Frequency-to-digital conversion is performed by the Boeing SMARTIO integrated circuit

2.6.1 Ground Testing—The results of 8.0- μ m/0.25- μ m n-channel transistor ground measurements are as follows:

1. (V_{th}) Threshold Voltage = 220 mV
2. (Lkg1) Drain-Source Leakage ($V_{ds} = 1.0$ V, $V_{gs} = 0.0$ V) = 415 nA
3. (Lkg2) Drain-Source Leakage ($V_{ds} = 2.0$ V, $V_{gs} = 0.0$ V) = 7.1 μ A
4. (Lkg3) Drain-Source Leakage ($V_{ds} = 2.0$ V, $V_{gs} = -0.5$ V) = 6.1 nA
5. (Drv1) Drive Current ($V_{ds} = 1.0$ V, $V_{gs} = 1.0$ V) = 339 μ A
6. (Drv2) Drive Current ($V_{ds} = 2.0$ V, $V_{gs} = 2.0$ V) = 1.2 mA
7. (Gsat) Saturation Transconductance = 790 μ S
8. (Gds) Drain-Source Output Conductance = 99 μ S.

The results of 8.0- μ m/0.25- μ m p-channel transistor ground measurements are as follows:

The results of L = 0.25 μm ring-oscillator performance ground measurements @ 2.0 V are as follows:

1. Oscillator #1 Stage Delay = 40.7 ps
2. Oscillator #2 Stage Delay = 41.2 ps
3. Oscillator #3 Stage Delay = 41.7 ps
4. Oscillator #4 Stage Delay = 43.1 ps.

2.6.3 Flight Testing—The results of 8.0-μm/0.25-μm n-channel transistor flight measurements are as follows:

Test	25-May-99	30-May-99	5-Jul-99	11-Jul-99	8-Aug-99	15-Aug-99	29-Aug-99	5-Sep-99	12-Sep-99
Vth (mV)	213	213	218	213	213	220	216	225	228
Lkg1 (A)	4.64×10 ⁻⁷	4.15×10 ⁻⁷	4.15×10 ⁻⁷	4.15×10 ⁻⁷	4.15×10 ⁻⁷	4.15×10 ⁻⁷	3.66×10 ⁻⁷	4.15×10 ⁻⁷	4.15×10 ⁻⁷
Lkg2 (A)	7.35×10 ⁻⁶	6.91×10 ⁻⁶	6.91×10 ⁻⁶	6.91×10 ⁻⁶	6.96×10 ⁻⁶	6.91×10 ⁻⁶	6.86×10 ⁻⁶	6.91×10 ⁻⁶	6.91×10 ⁻⁶
Lkg3 (A)	7.1×10 ⁻⁹	8.1×10 ⁻⁹	5.1×10 ⁻⁹	2.2×10 ⁻⁹	7.3×10 ⁻¹⁰	5.6×10 ⁻⁹	3.7×10 ⁻⁹	6.1×10 ⁻⁹	3.2×10 ⁻⁹
Drv1 (A)	1.07×10 ⁻³	1.08×10 ⁻³	1.08×10 ⁻³	1.08×10 ⁻³	1.08×10 ⁻³	1.08×10 ⁻³	1.08×10 ⁻³	1.08×10 ⁻³	1.08×10 ⁻³
Drv2 (A)	2.96×10 ⁻³	2.98×10 ⁻³	2.98×10 ⁻³	2.98×10 ⁻³	2.98×10 ⁻³	2.98×10 ⁻³	2.98×10 ⁻³	2.98×10 ⁻³	2.98×10 ⁻³
Gsat (μS)	1531	1556	1555	1555	1531	1531	1556	1531	1555
Gds (μS)	123	123	123	123	123	148	123	123	123

The results of 8.0-μm/0.25-μm p-channel transistor flight measurements are as follows:

Test	25-May-99	30-May-99	5-Jul-99	11-Jul-99	8-Aug-99	15-Aug-99	29-Aug-99	5-Sep-99	12-Sep-99
Vth (mV)	-304	-304	-304	-314	-306	-316	-302	-314	-304
Lkg1 (A)	2.45×10 ⁻⁹	2.45×10 ⁻⁹	2.45×10 ⁻⁹	2.45×10 ⁻⁹	2.45×10 ⁻⁹	2.43×10 ⁻⁹	2.43×10 ⁻⁹	2.45×10 ⁻⁹	2.45×10 ⁻⁹
Lkg2 (A)	2.43×10 ⁻⁸	2.43×10 ⁻⁸	2.43×10 ⁻⁸	2.43×10 ⁻⁸	2.43×10 ⁻⁸	2.43×10 ⁻⁸	2.43×10 ⁻⁸	2.43×10 ⁻⁸	2.43×10 ⁻⁸
Lkg3 (A)	6.1×10 ⁻⁹	5.6×10 ⁻⁹	5.6×10 ⁻⁹	6.1×10 ⁻⁹	5.6×10 ⁻⁹	5.6×10 ⁻⁹	6.1×10 ⁻⁹	5.6×10 ⁻⁹	5.6×10 ⁻⁹
Drv1 (A)	3.29×10 ⁻⁴	3.34×10 ⁻⁴	3.34×10 ⁻⁴	3.34×10 ⁻⁴	3.34×10 ⁻⁴	3.34×10 ⁻⁴	3.34×10 ⁻⁴	3.39×10 ⁻⁴	3.34×10 ⁻⁴
Drv2 (A)	1.18×10 ⁻³	1.19×10 ⁻³	1.18×10 ⁻³	1.19×10 ⁻³	1.19×10 ⁻³	1.19×10 ⁻³	1.19×10 ⁻³	1.19×10 ⁻³	1.19×10 ⁻³
Gsat (μS)	790	790	790	790	790	814	790	790	765
Gds (μS)	99	99	99	99	99	99	99	99	99

The results of L = 0.25 μm ring-oscillator flight measurements in ps are as follows:

Test	25-May-99	30-May-99	5-Jul-99	11-Jul-99	8-Aug-99	15-Aug-99	29-Aug-99	5-Sep-99	12-Sep-99
Stage Delay Osc1	41.6	40.7	40.9	40.9	40.8	40.6	40.7	No Data	41.6
Stage Delay Osc2	42.3	41.5	41.6	41.6	41.5	41.3	41.4	40.9	42.3
Stage Delay Osc3	42.8	42.1	42.2	42.2	42.1	42.0	42.0	42.2	42.8
Stage Delay Osc4	44.4	43.6	43.8	43.8	43.7	43.5	43.6	43.8	44.4

2.7 Test Result Comparison

DS1 launched in October 1998 with the first LPE data downlink in May 1999. Linear interpolation has determined that the LPE received ~8 krad passing through the Van Allen Belt and has since received an additional ~16 rad/day.

Figure 8 through Figure 27 show comparison plots between measurements made before and after launch.




Figure 8. N-channel Transistor (Vth); ~2-mV A/D Converter Resolution (Left-hand Y-axis)




Figure 9. N-channel Transistor (Leakage1)




Figure 10. N-channel Transistor (Leakage2)




Figure 11. N-channel Transistor (Leakage3)




Figure 12. N-channel Transistor (Drive Current1)




Figure 13. N-channel Transistor (Drive Current2)




Figure 14. N-channel Transistor (gmSat)




Figure 15. N-channel Transistor (Gds)




Figure 16. P-channel Transistor (Vth); ~2-mV A/D Converter Resolution (Left-hand Y-axis)




Figure 17. P-channel Transistor (Leakage1)

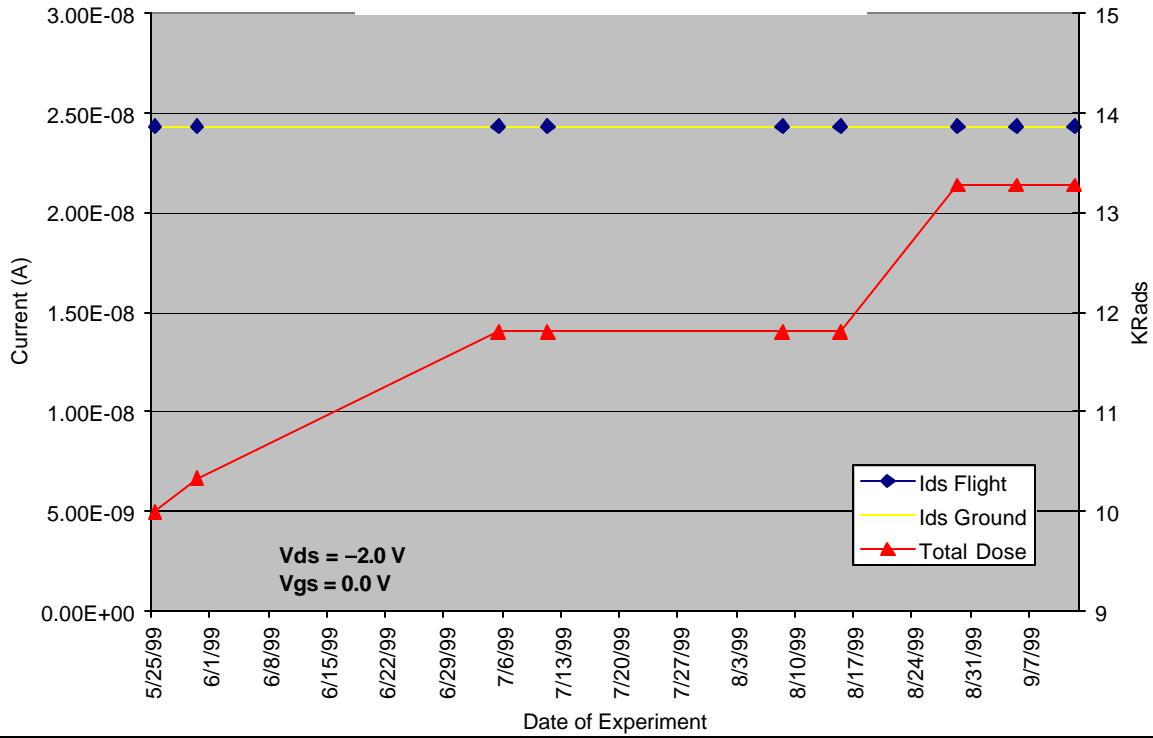


Figure 18. P-channel Transistor (Leakage2)




Figure 19. P-channel Transistor (Leakage3)




Figure 20. P-channel Transistor (Drive Current1)




Figure 21. P-channel Transistor (Drive Current2)




Figure 22. P-channel Transistor (gmSat)




Figure 23. P-channel Transistor (Gds)

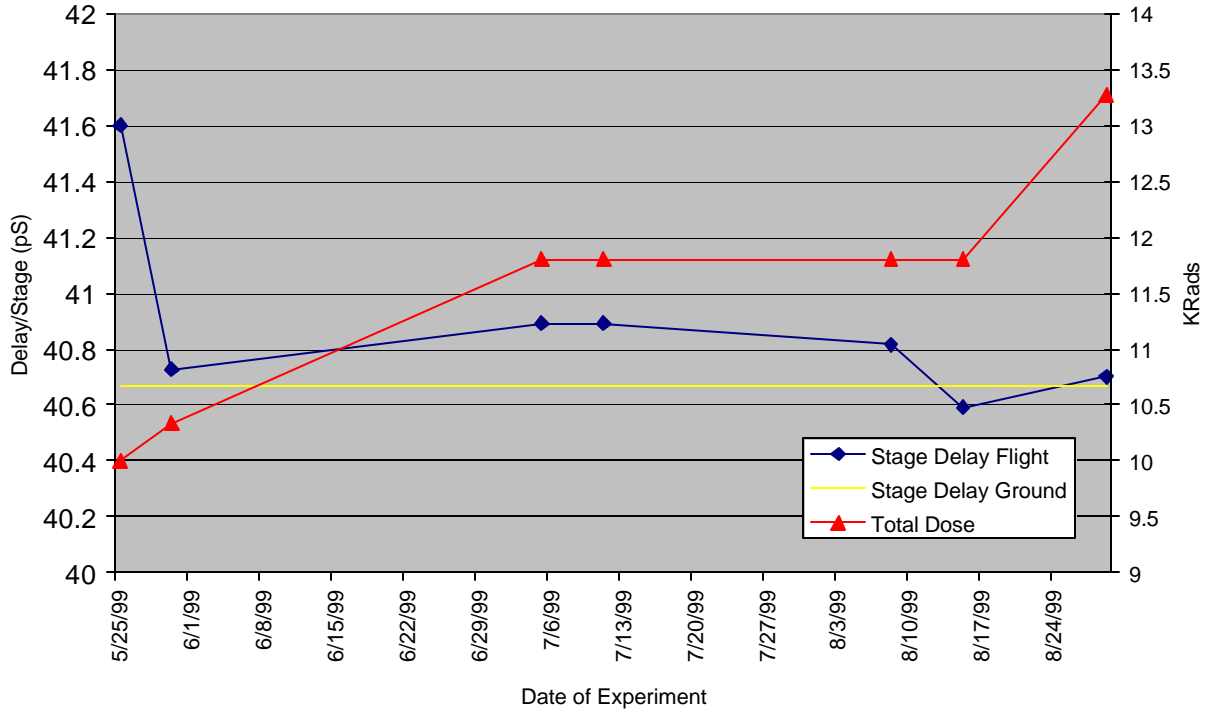


Figure 24. $L = 0.25\text{-}\mu\text{m}$ 97-Stage Ring Oscillator 1




Figure 25. $L = 0.25\text{-}\mu\text{m}$ 97-Stage Ring Oscillator 2




Figure 26. $L = 0.25\text{-}\mu\text{m}$ 97-Stage Ring Oscillator 3




Figure 27. $L = 0.25\text{-}\mu\text{m}$ 97-Stage Ring Oscillator 4

3.0 TECHNOLOGY VALIDATION SUMMARY

The total dose exposure to the LPE board and the 0.25- μm FDSOI test chip was relatively small and had little effect (if any) on the operational characteristics of the devices. Minor fluctuation in measured values represent a small number of counts of the sampling A/D converter and may be attributed to system noise.

These first steps towards qualification have helped to validate 0.25- μm FDSOI as an important technology in the design and advancement of low-power, high-performance electronics for space application.

4.0 TECHNOLOGY APPLICATION FOR FUTURE MISSIONS

Low-power, high-performance electronics is a prerequisite for almost every space-based and terrestrial electronic application. The FDSOI CMOS technology described in this report and validated onboard the DS1 spacecraft provides a glimpse of the future of electronic computation. As the commercial electronics industry continues to follow Moore's Law enabling smaller, faster, and cheaper electronics, SOI technology is starting to pop up on the road maps of many integrated circuit manufacturers. MIT Lincoln Laboratory continues to push this technology with current circuit work targeting 0.175- μm feature sizes, 1.5-V operation, and 15- to 17-ps ring-oscillator stage delays. Cutoff frequencies for n-channel devices in the 0.175- μm process are measuring ~ 85 GHz. Advanced development has already begun on the sub-0.1- μm version of the FDSOI technology and measurement results are expected soon. The unique attributes of fully depleted SOI (reduced device

parasitic capacitances, enhanced subthreshold swing enabling low-threshold and, hence, low-power-supply operation, full oxide isolation between devices eliminating traditional bulk CMOS latchup, and inherent radiation resistance) make it the technology of choice for silicon-based electronic applications in space.

5.0 ACKNOWLEDGMENTS

The LPE Team at MIT Lincoln Laboratory thanks DARPA/MTO for process development and fabrication sponsorship, the packaging expertise of Peter Daniels and his team, and the staff of MIT Lincoln Laboratory's Microelectronics Laboratory for their hard work and dedication.

A special thanks to the efforts of JPL, especially Eric Holmberg, LPE Cognizant Engineer, and Kirk Fleming, Engineer assigned to LPE post-launch activities.

The work described in this report was carried out at the Lincoln Laboratory, Massachusetts Institute of Technology, with sponsorship from the National Aeronautics and Space Administration, and the Defense Advanced Research Projects Agency.

6.0 LIST OF REFERENCES

- [1] H. I. Liu, et. al., "Thin Silicide Development for Fully-Depleted SOI CMOS Technology," *IEEE Trans. on Electron Dev.*, Vol. 45, No. 5, (1998) pp. 1099–1104.
- [2] R. Berger, et. al., "A 1.3 GHz SOI CMOS Test Chip for Low-Power High Speed Pulse Processing," *IEEE Jour, Solid-State Circuits.*, Vol. 33, No. 8, (1998) pp. 1259–1261.

Appendix A. Telemetry Channels

Channel	Mnemonic
P-0300	LPE_PAM_mgr
P-0301	cmd_quality
P-0302	last_cmd_id
P-0303	LPEdataQual
P-0304	LPEdataWord
P-0305	LPE_complete
P-0306	LPE_t_stamp
P-0307	LPEresetWrd0
P-0308	LPEresetWrd1
P-0309	LPEresetWrd2
P-0310	LPEresetWrd3
P-0311	LPEresetWrd4
P-0312	LPEresetWrd5
P-0313	LPEresetWrd6
P-0314	LPEresetWrd7
D-0096	last_pkt_06
D-0097	buf_typ_06
D-0098	buf_min_06
D-0099	buf_max_06
D-0100	pkt_age_06
D-0101	buf_pkt_06
D-0102	sent_pkt_06
D-0103	spac_used_06
D-0104	bytes_ack_06
D-0105	byte_dump_06

Appendix B. Date of Turn-on/Frequency of Data Capture

Date:

25-May-1999

30-May-1999

5-Jul-1999

11-Jul-1999


8-Aug-1999

15-Aug-1999

29-Aug-1999

5-Sep-1999


12-Sep-1999



Multifunctional Structures Technology Demonstration on New Millennium Program (NMP) Deep Space 1 (DS1) DS1 Technology Validation Report

David M. Barnett
Dr. Suraj P. Rawal

LOCKHEED MARTIN



*Aeronautics Division
Denver, CO*

JPL

Table of Contents

<u>Section</u>	<u>Page</u>
Acronyms	iv
Abstract	v
Fact Sheet	vii
1.0 Introduction	1
2.0 Technology Description	1
2.1 MFS Functionality and DS1 Demonstration	1
2.2 Key Technology-Validation Objectives at Launch	1
2.3 Expected Performance Envelope	2
2.4 Detailed Description	2
2.4.1 Interconnect Systems	3
2.4.2 Multi-chip Module Interface Characteristics	3
2.4.3 MFS Flight Experiment and Data-Collection Descriptions	3
2.4.4 Circuit Description	4
2.4.5 Data Collection	4
2.5 Technology Interdependencies.....	4
2.6 Test Program.....	4
2.6.1 Ground Test.....	4
2.6.2 Flight Test.....	6
2.7 Comparison between Ground Test and Flight Test.....	6
3.0 Technology-Validation Summary	7
4.0 Technology Application for Future Missions	7
5.0 Acknowledgments	7
6.0 List of References	8

Figures

<u>Figure</u>	<u>Page</u>
Figure 1. Exploded Assembly View of a Multi-Functional Structure Panel Incorporating Structure, Thermal and Electrical Elements.....	2
Figure 2. Multi-Functional Structures Experiment Installed on a NMP DS1 Flight Panel.....	3
Figure 3. Block Diagram of the MFS Experiment Electronics on the NMP DS1 Flight	4
Figure 4. Vibration Levels Used in the Testing of the NMP DS1 MFS Engineering Development Unit.....	5

Tables

<u>Table</u>	<u>Page</u>
Table 1. Description of the Continuity-Measurement Collection-System Output	5
Table 2. Flight-Continuity Measurement Data	6
Table 3. Flight-Temperature Measurement Data (all in °C)	7

Appendices

<u>Appendix</u>	<u>Page</u>
Appendix A. DS1 Technology Validation Telemetry Channels	9
Appendix B. DS1 Technology Validation Power On Times	10

Acronyms

AFRL	AirForce Research Laboratory
BMDO	Ballistic Missile Defense Office
C&DH	Command and Data Handling
Cu/PI	Copper/Polyimide
DARPA	Defense Advanced Research Project Agency
DOD	Department of Defense
DS1	Deep Space 1
EDU	Engineering Development Unit
EMI/EMC	Electromagnetic Interference/Electromagnetic Compatibility
GFE	Government Furnished Equipment
HDI	High-Density Interconnect
HiLowPDM	High-side/Low-side Power Distribution Module
IC	Integrated Circuit
I/O	Input/Output
I&T	Integration and Test
LMA	Lockheed Martin Astronautics
LSB	Least Significant Bit
LTCC	Low-Temperature Co-fired Ceramic
MCM	Multi-chip Modules
MFS	Multifunctional Structures
NMP	New Millennium Program
PL	Phillips Laboratory
PWB	Printed Wiring Board
SIES	Spacecraft Integrated Electronics Systems

ABSTRACT

The future microspacecraft vision will only be realized through revolutionary changes in current spacecraft architecture coupled with the development of new technologies. This paradigm shift will bring a dramatic cultural change that can only be implemented using a truly concurrent-engineering approach that incorporates advances in structural, thermal, microelectronics, micro-instruments, sensor, power, and propulsion systems.

Addressing technology needs of future microspacecraft, Lockheed Martin Astronautics (LMA) has developed an innovative multifunctional structures (MFS) design that is a new approach to electronics packaging, interconnection, and data and power distribution. MFS integrates these functions with bearing-mechanical loads and provides thermal control.

In particular, the MFS concept involves embedding passive-electronic components within the actual volume of composite materials, new approaches to attaching active-electronic components directly to mechanical surfaces, and using surface areas for mounting sensors and transducers.

The ultimate goal for MFS technology is to maximize the ratio of the volume of the fundamental electronic parts to the total packaging volume. Multi-functional structure technology is a revolutionary design approach that will provide nearly an order-of-magnitude reduction in future spacecraft mass and volume. Significant cost savings are also expected through eliminating touch labor, reuse of flex-circuitry designs for multiple-spacecraft missions, and launch-cost reductions through reduced payload size. MFS is an enabling technology for future microspacecraft missions envisioned by the Department of Defense (DOD) and NASA.

The MFS design approach uniquely combines the advances in the area of electronics (e.g., 2-D/3-D multi-chip modules [MCM]) and flex-circuit interconnects), advanced composites (for structures), and thermal management. MFS eliminates the bulky components (chassis, cables, and connectors) of current spacecraft and enables the integration of electronic subsystems, such as the data-transmission and power-distribution networks, command and data handling (C&DH) subsystem, thermal management, and load handling.

The baseline MFS design consists of a structural-composite panel that has multi-layer copper/polyimide (Cu/PI) patches bonded to one side, heat-transferring devices embedded, and an outer surface acting as a radiator. Electrical interconnects are designed in the Cu/PI layers, circuitry is implemented in MCMs, and flex

jumpers serve as electrical interconnects for power distribution and data transmission. The thermal management devices embedded in the MFS may include miniature heat pipes and various types of high-conductivity thermal doublers and straps.

In an Air Force Research Laboratory/Philips Laboratory (AFRL/PL), Ballistic Missile Defense Office (BMDO) and Defense Advanced Research Project Agency (DARPA)-sponsored program, LMA has successfully developed and demonstrated the design, integration, assembly, and functional performance of the MFS technology and its elements.

LMA has successfully integrated an MFS experiment on the NASA New Millennium Program (NMP) Deep Space 1 (DS1) spacecraft and validated key technology features of MFS design.

Technology and integration risks associated with the MFS-packaging system include:

- The electrical performance of the flex circuit, including the anisotropic electrical interconnects to the flex-circuit jumpers.
- The use of socketed MCMs in a flight environment.
- Connections between the flex circuitry and heritage connectors.
- Integration and test, rework, and repair issues associated with the direct installation of electronics on spacecraft structure without a chassis.

The validation objectives include the successful demonstration of the MFS technology elements (integrating flex interconnect, circuit patches, flex jumpers, thermal doublers, rad-hard composite spot shield, and structural substrate). In the electrical circuit performance area, conductivity measurements were taken during each experiment cycle to independently verify the nominal-trace conductivity, the performance of the anisotropic bonds in a jumper configured for multiple serpentine connections, and a set of daisy-chained connections to the thermal-simulator MCM through a socketed-lead system. A set of temperature measurements were collected to evaluate the thermal performance of the panel underneath the thermal-simulator MCM by using an array of sensors mounted on a flex-circuit tether. Finally, routine health- and status-data were collected to verify proper controller operation during the data collection.

Given the novel nature of the MFS design, extensive development testing was performed prior to any DS1 design effort. This testing included vibration, thermal, x-ray, and electrical performance of a variety of test panels with different configurations of hardware. The technology was fairly well documented leading into the DS1 experiment design. The DS1 components were tested both individually and as a system. The controller board for the flight experiment was tested for

workmanship and the completed panel was tested with the full-up spacecraft assembly.

During the DS1 mission, the MFS experiment was powered up once every two weeks and two experiment cycles were carried out to ensure that a full set of data was collected. The experiment sequence provided a data set containing health and status information, the electrical-conductivity test data, and last (following a warming time period of the panel) thermal-gradient measurements. The experiment was an unqualified success based on the data returned. All health and status data was correct and within normal limits. The electrical-conductivity data never varied by more than one Least Significant Bit (LSB) from the preflight data set. The thermal-gradient data was appropriate for the position of the sensors versus the heat source in the MCM.

MFS technology is eminently suited to use in many missions for the following reasons:

- Offers significant mass (>50%) and volume (>2×) savings over traditional packaging systems.

- Takes full advantage of MCM devices without adding packaging mass due to printed wiring board (PWB) mounting.
- Frees up spacecraft design from traditional form factors.
- Flex circuitry is an enabling technology for wiring density in microspacecraft.
- The technology will readily support mass production of spacecraft for constellations.
- The techniques easily transfer to inflatable structures.

Overall, the NMP-DS1-MFS experiment has been very successful in demonstrating the majority of key features and showing that there are no major roadblocks. Even a minor rework was performed smoothly with the panel in place on the spacecraft bus. The MFS experiment was integrated quite easily on the spacecraft-bus panel and was the first technology experiment to be delivered to DS1.

Based on the successful technology-validation experiment, the MFS technology should be considered fairly mature.

FACT SHEET MULTIFUNCTIONAL STRUCTURES


Strategic Vision

Establish modular multifunctional structures (MFS) technology, integrating electronic, thermal, and structural functions for Next-Generation, Cable-Free Spacecraft



SATELLITE MANUFACTURING TODAY:


- 1000s of heavy individual cables
- Costly
- Complex
- Bulky electronic enclosures
- Wasted space



SATELLITE MANUFACTURING TOMORROW:

- Eliminate cables and connectors
- >10× reduction in mass and volume
- > 2× reduction in cost
- Enabling technology for satellites, launch vehicles, and missiles
- Revolutionary modular design

MFS Concept



Description

- Flex copper/polyimide circuit patch and interconnects for power and data distribution
- Flex material directly bonded to thermal-structural composite panel
- Applicable to intersubsystem cabling
- Revolutionary modular design

Who Needs It?

- All LEO platforms and GEO satellites
- High-density instruments, sensors, and microassemblies
- Nearly all rigid and inflatable structures

Contact:

David Barnett (303-971-8061) or Dr. Suraj Rawal (303-971-9378) Lockheed Martin Astronautics Denver, CO 80201
Dr. Alok Das (505-846-8250) Air Force Research Lab, Kirtland Air Force Base, NM



Multifunctional Structures (MFS) Technology Demonstration on New Millennium Program (NMP) Deep Space 1 (DS1) DS1 Technology Validation Report

*David M. Barnett and Dr. Suraj P. Rawal
Lockheed Martin, Astronautics Division, Denver, Colorado*

1.0 INTRODUCTION

The MFS technology is a revolutionary development in spacecraft packaging that eliminates chassis and cabling by integrating the electronics, thermal control, and structure into a single element. A new system such as this carries the burden of proving itself before it can be considered as a viable design option for flight usage. During the early development of the MFS concept, extensive environmental and electrical testing was performed to demonstrate the robustness of this system. The next step was a flight demonstration; this was accomplished on the NMP DS1 spacecraft.

The following sections describe the MFS system, the validation objectives, potential risks and risk amelioration, the testing program and results, and the future use of MFS in spacecraft design. As a new system, it may take a little while for the design concept to “sink in”; however, the ramifications of this technology for future designs at all levels, but especially in microspacecraft, will be apparent in the description below.

2.0 TECHNOLOGY DESCRIPTION

2.1 MFS Functionality and DS1 Demonstration

The multifunctional structure technology is a new method for constructing spacecraft. An MFS demonstration was proposed for the DS1 spacecraft to incorporate the key design features. Eventually, it is envisioned that entire spacecraft will be fabricated using the MFS system; the DS1 mission provided a starting point. The experiment was designed to demonstrate several features of the technology, including design methods, integration and test (I&T) impacts, functional routing of signals and power, use of flex circuitry in novel ways, and the elimination of chassis and cabling.

The experiment was designed with a spacecraft-interface card to support data gathering, formatting, and transmittal to the main spacecraft computer. The interface card followed an experiment sequence in collecting health and status data (voltages, check sums, initial temperatures, and a subset of electrical conductivity measurements), a full set of conductivity data on a variety of circuit conductors in

multiple configurations, and a temperature-gradient measurement following a 30-minute panel-heating operation. The data set was collected twice in succession to increase the odds of obtaining full data sets. This was in lieu of having any spacecraft data checking to look for dropouts.

The high-level goals for the experiment included successful installation, proper operation of the circuitry over the life of the mission, good thermal performance of the thermal-simulator multi-chip module (MCM) mounting system, and minimal problems dealing with the unique features of the packaging.

2.2 Key Technology-Validation Objectives at Launch

The primary validation objectives included:

- Demonstrate proper electrical performance for the flex-circuitry conductors.
- Monitor the anisotropic flex-to-flex sample bonds for any sign of degradation.
- Verify the stability of MCM electrical connections made using a separable connector attached to the device leads.
- Collect thermal-gradient data that demonstrates proper heat removal from the thermal-simulator MCM.

Given the novel nature of this technology, a variety of intrinsic objectives were also indicated as follows:

- Demonstrate a concurrent engineering effort on the experiment layout and design.
- Demonstrate successful installation of the hardware on another subcontractor’s flight panel.
- Show that rework/repair operations are straight forward even if performed with the panel on the bus.
- Verify the flightworthiness of a new MCM socket that permits rapid removal and replacement of MCMs.
- Demonstrate an instrumentation tether by collecting data from the opposite side of the panel using a flex-circuit element with a linear array of temperature sensors.
- Demonstrate a cover that provides mechanical protection, EMI/EMC shielding, and radiation shielding.
- Demonstrate the use of filled-composite materials for localized radiation shielding of the printed-wiring board (PWB).

2.3 Expected Performance Envelope

The conductivity measurements of the flex circuitry and MCM socket system follow standard analytical techniques for resistance in copper conductors. The criteria for the returned flight measurements primarily centers on the repeatability from ground-to-flight measurements within a specified tolerance: i.e., the launch and flight environment do not cause any degradation that would either increase the resistance or result in a completely open circuit. The allowable tolerance was 20% from ground to flight. A variety of circuit-trace and socket configurations were created to permit testing of traces, anisotropic bonds, and the MCM socket interconnects. Each configuration permitted independently testing for a single type of interface, thereby avoiding contamination between different interconnect systems.

The thermal-gradient part of the experiment measured the temperature distribution over a small area of the spacecraft panel before and after a 30-minute heating cycle. The predicted-maximum rise for safety purposes was approximately 5° C regardless of total heating time. The heat source was a thick-film resistor screen printed in the thermal-simulator MCM package. The resistor footprint was sized to simulate the dissipation from an integrated circuit. The expected performance of the thermal-bonding system from the MCM to the panel was to produce a maximum rise of about 5° C directly under the MCM “hot spot” with an appropriate falling off over near distance from the hot spot.


Secondary performance expectations included no significant degradation in the performance of the spacecraft-interface electronics, no loss of measurements or measurement data, and successful power-up and communication sequences with the spacecraft interface. Also, the survival of the packaging system through the launch phase was obviously a strict criterion given the focus of multifunctional structures.

2.4 Detailed Description

Driven by a spacecraft requirement on a flight program to incorporate a reworkable MCM stack processor on a composite panel, the Spacecraft Integrated Electronics Systems (SIES) program was funded by Air Force Research Laboratory (AFRL)/Phillips Laboratory (PL) to develop methods for efficiently incorporating MCMs into spacecraft without losing the volumetric and mass advantages. The MFS efforts have produced a system that incorporates structure, thermal control, and electronics into a single packaging system while permitting efficient rework and test. Chassis, PWBs, connectors, and cabling have all been eliminated in this system.

The MFS assembly concept is shown in Figure 1. A composite panel with embedded or laminated thermal-

control elements forms the basis for the MFS. For typical-spacecraft construction, clips will be used on the edge of the panel for mechanical attachment to adjacent panels. Flex-circuit patches are then installed on the panel using adhesive. These circuit patches provide local interconnects and can accommodate surface-mount devices. MCM sockets developed in this effort are also installed at this time. Flex-circuit jumpers are added for patch-to-patch and panel-to-panel interconnection and are connected using an in-place bonding system. These jumpers provide signal paths and shielding as required. In the MFS system, it should be noted that traditional chassis, mother boards, cabling, and connectors have all been replaced with the flex-interconnect system. The PWB electronics are reduced to MCMs.



Spacecraft Structural Panel with Integral Thermal Control
Figure 1. Exploded Assembly View of a Multifunctional Structure Panel Incorporating Structure, Thermal and Electrical Elements

The next step is the installation of leaded MCMs into the sockets with a clamping assembly to secure the part and the leads. The clamp also ensures that adequate thermal circuit would normally have functional testing. Test jumpers can be added for testing and then removed or stowed. Finally, a cover is installed that can provide the following protection: physical protection during assembly, electromagnetic interference/electromagnetic-compatibility (EMI / EMC) shielding, and radiation shielding. The entire MFS system is designed to readily support repair/rework with a fundamental requirement that in all cases the MCMs shall be easily removed for reuse with no risk of damage.

The SIES contract provided for several demonstrations, including a flight demonstration on NASA’s NMP DSI mission. This provided an ideal opportunity to validate the technology in terms of produceability and long-term flight worthiness during a multi-year mission. An experiment was designed to apply and test the following features of the MFS system: Circuit patches, in-place jumpers, socketed MCM, soldered MCM, flex tether with embedded sensors, and flex interfaces to traditional connectors. The MFS DSI experiment is shown in Figure 2.




Figure 2. Multifunctional Structures Experiment Installed on a NMP DS1 Flight Panel. The PWB carries the spacecraft interface circuitry and contains the controller for managing the experiment.

The MCMs used in the experiment require close attention to detail in terms of electrical interconnections, thermal and structural interfaces, and compatibility with the remainder of the MFS concepts. The remainder of the paper will discuss these aspects of MCM usage in the system in greater detail and outline the development, limitations, and hardware testing.

2.4.1 Interconnect Systems—In their most basic form, interconnects provide an insulated electrical path from a source circuit to a destination. Enhancements include shielding and connectors or, in the case of local circuitry, circuit traces, and electrical joints, such as solder.

Flexible circuitry was selected for the MFS concept for a variety of reasons:

- Replaces both PWBs and cabling.
- Local electrical-bonding systems can eliminate all connectors.
- Lightweight/low volume.
- Standard product.
- Conductors can be sized/added to meet voltage-drop, shielding, and isolation requirements.

It is very desirable to eliminate connectors for several reasons. They are bulky compared to the conductors and add significant weight. They are usually labor intensive between assembly, calibration of assembly tools, inspections, and test. They can be the source of many additional failure modes.

In a pure implementation of the MFS design, all connectors are eliminated. In situ flex-to-flex bonding is performed to link circuit patches and flex cabling. The only routine

connector left is a prototype MCM socket that has been qualified in extreme vibration environments and is being demonstrated on NMP DS1. The MCM leads are clamped into the connector; the MCM package is then clamped to the panel. This approach supports the easy removal and reinstallation of MCMs during re-work and minimizes the loss of expensive parts.

2.4.2 Multi-chip Module Interface Characteristics—Primary interface considerations for the use of MCMs with MFS are the electrical interconnects and the thermal interface. These are linked in various MCM packages because some lead styles are in the normal-thermal path through the base of the package. Electrical interconnects fall into four categories: leaded packages, pin grids, ball grids (and column grids), and flex-circuit extensions. The best performance is obtained from the leaded and flex-extension packages since they can have the base of the package in good thermal contact with the panel when mounted. The package can either be directly in contact with the panel facesheet or additional heat spreaders/doublers can be used for higher dissipation levels. Ball-grid-array attachment presents greater challenges in thermal-dissipation management.

In the DS1 experiment, two MCMs are used. One MCM is a high-density interconnect (HDI) type of device; the other unit is fabricated in low-temperature, co-fired ceramic (LTCC). The HDI device is fabricated using integrated circuit (IC) dice mounted in a ceramic carrier, with the local interconnects made using multiple layers of flex circuit that is repeatedly laser drilled to the IC die-bonding pads, metallized, and then etched to leave traces. External connections can be made with a flex interposer or conventional lead frames (used in the DS1 part). The HDI device functions as a high-side/low-side switching power-distribution module (HiLoPDM).

The LTCC device has a number of thick-film resistors of different geometries simulating the dissipation from different IC dice. The interconnects are formed with a conventional lead frame whose pitch matches the MCM socket strips. The HiLoPDM device is soldered in using conventional methods. Both of these devices are shown in Figure 2 (although the LTCC device is concealed under a clamping plate).

2.4.3 MFS Flight Experiment and Data-Collection Descriptions—The MFS experiment on NMP DS1 is the first flight demonstration of the MFS technology. The experiment met several guidelines, including: not flight critical, minimal data set collected once per two-week cycle, basic RS232 interface, basic command protocol, low power, and no failure modes that could either cause excessive thermal dissipation or cause electrical-bus faults. The electronics block diagram is shown in Figure 3.

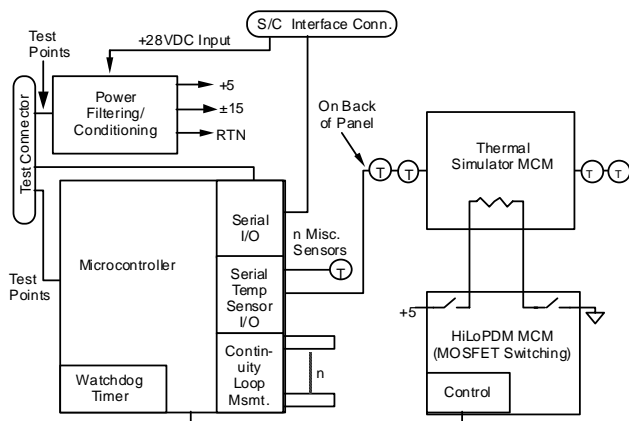


Figure 3. Block Diagram of the MFS Experiment Electronics on the NMP DS1 Flight. The basic experiment verifies the electrical performance of the various interconnect systems used in MFS and demonstrates distributed sensors measuring an induced-thermal gradient in the panel.

2.4.4 Circuit Description—A microcontroller PWB was provided for the spacecraft and test interfaces. The interface board provides the Input/Output (I/O) for the experiment operations, including conductivity measurements of flex-circuit traces and MCM–socket contacts, control of the HiLoPDM power switch for the thermal-gradient experiment, and the temperature sensors that measure the gradient.

The conductivity measurements cover the following: copper “control” traces for nominal conductor performance, traces through the flex jumpers, which includes the anisotropic bonds at each end, and the socket contacts, which are daisy chained in and out of the thermal-simulator MCM contacts. The original design only included an open/connected determination; however, this was modified to a regular measurement with a high and low range to determine if the conductors are degrading. The three types of connections have been designed to keep the types of connections independent: i.e., there are no anisotropic bonds in the MCM socket-pin path, etc. This avoids confusion in data interpretation if there is a systematic failure in one type of interconnect.

In the diagram, a series of temperature sensors are shown passing beneath the thermal-simulator MCM. There are 12 sensors in a 4 × 3 array on the back of the panel under the footprint of the MCM. These devices use a three-wire interface with serial communications and unique addressing. They are mounted on a serpentine-flex circuit that passes from the front of the panel to the rear and is then attached with film adhesive. Several sensors are also used for other measurements on the front of the panel, including the PWB.

2.4.5 Data Collection—Data collection is initiated by the spacecraft through the following sequence: Power on, command #1 to MFS, response with health-and-status information including software version and checksums, command #2 to MFS, response with conductance measurements and 30-minute heating cycle when the thermal simulator MCM is started, pause approximately 30 minutes, command #3 to MFS, response with thermal-gradient data, power is turned off. This cycle is repeated twice in succession to ensure a complete data set. The sequence will be repeated every two weeks during the mission until the link efficiency starts to fall off; thereafter, the sequence will be repeated at larger time intervals.

2.5 Technology Interdependencies

There are no interdependencies between MFS and other spacecraft subsystems.

2.6 Test Program

2.6.1 Ground Test—

2.6.1.1 Development and Protoflight Testing—Thermal and vibration testing was performed during the MFS development efforts to verify performance and electrical integrity. The MCM socket was felt to be especially critical to the value of the MFS system and was therefore subjected to extreme levels of vibration while being monitored by “chatter” detectors for intermittents on the connections. Basically, the MCMs and socket assembly satisfied the typical vibration environment. Subsequently, the vibration levels were increased to test to failure. While the imposed vibration levels generated localized delamination in the panel, the MCM’s socketing approach was robust, with no indication of degradation.

The NMP DS1 engineering development unit (EDU) panel was subjected to much lower levels since there was a PWB mounted on standoffs with potential resonances. During the design phase, the board was analyzed against the flight requirements and the design was adjusted to ensure survival. Figure 4 is a graph of the flight-environment envelope for the protoflight vibration testing. The panel did not have any failures.

Thermal testing was performed during the development phases of the MFS designs. The primary goal was to determine if there were anticipated MCM-dissipation levels that exceeded the capability of a well-designed composite panel with associated thermal controls. The conclusion was that there is probably not a heat load for which design techniques cannot keep the baseplate temperature within the limits necessary to meet junction-temperature requirements. The ultimate limitation is the capability of the spacecraft to dissipate the total heat load.




Figure 4. Vibration Levels Used in the Testing of the NMP DS1 MFS Engineering Development Unit.

2.6.1.2 Continuity and Thermal Gradient Measurements—

The continuity measurements were the primary measurements for meeting the MFS demonstration goals. During ground tests, measurements were taken and were found to comply with the analytical predictions listed in

Table 1. It should be noted that digitized measurements have a built-in error of $\pm 1/2$ LSB; therefore, there are cases in the flight data where a change in one LSB will be seen between different samples. In addition, every measurement was taken on both a unity-gain scale and a $\times 10$ scale. Given the low values of the resistances measured, the $\times 10$ scale is more representative of the measurements. The continuity measurement is non-linear due to the electronics; therefore, the corresponding resistance ranges are also listed for reference. There are three types of continuity measurement: flex-circuit trace only, anisotropic bond, and MCM socket. These are also distinguished below. Layout variations on the flex-circuit patch are responsible for the variations in similar measurements.

The temperature sensors were commercial devices that are designed for a resolution of 0.01°C . The devices were attached to the flex-circuit tether and placed in a temperature chamber to obtain linearity curves. The part-to-part variation was on the order of $\pm 0.2^\circ\text{C}$. The original intent was to fabricate the panel in house. However, the panel ended up being Government Furnished Equipment

Table 1. Description of the Continuity-Measurement Collection-System Output

BYTE	NAME	DESCRIPTION	VALUE (DEC)	NOTES
1	LOOPCAL1	CALIBRATION MSMT	17+/-3	Rcal = 49.9 OHM
2	LOOP1	GAIN = 1 RESISTANCE MSMT	0+3-0	FLEX TEST
3	LOOP2	GAIN = 1 RESISTANCE MSMT	0+3-0	JUMPER/BONDING
4	LOOP3	GAIN = 1 RESISTANCE MSMT	0+3-0	JUMPER/BONDING
5	LOOP4	GAIN = 1 RESISTANCE MSMT	0+3-0	JUMPER/BONDING
6	LOOP5	GAIN = 1 RESISTANCE MSMT	0+3-0	JUMPER/BONDING
7	LOOP6	GAIN = 1 RESISTANCE MSMT	0+3-0	JUMPER/BONDING
8	LOOP7	GAIN = 1 RESISTANCE MSMT	0+3-0	MCM SOCKET
9	LOOP8	GAIN = 1 RESISTANCE MSMT	0+3-0	MCM SOCKET
10	LOOP9	GAIN = 1 RESISTANCE MSMT	1+3-1	MCM SOCKET
11	LOOP10	GAIN = 1 RESISTANCE MSMT	1+3-1	MCM SOCKET
12	LOOP11	GAIN = 1 RESISTANCE MSMT	1+3-1	MCM SOCKET
13	LOOP12	GAIN = 1 RESISTANCE MSMT	1+3-1	MCM SOCKET
14	LOOP13	GAIN = 1 RESISTANCE MSMT	0+3-0	FLEX TEST
15	LOOP14	GAIN = 1 RESISTANCE MSMT	0+3-0	FLEX TEST
16	LOOPCAL2	NOT USED	0	
17	LOOPCAL1	NOT USED	0	
18	VER1	GAIN = 10 RESISTANCE MSMT	2+/-2	FLEX TEST
19	VER2	GAIN = 10 RESISTANCE MSMT	3+/-2	JUMPER/BONDING
20	VER3	GAIN = 10 RESISTANCE MSMT	3+/-2	JUMPER/BONDING
21	VER4	GAIN = 10 RESISTANCE MSMT	3+/-2	JUMPER/BONDING
22	VER5	GAIN = 10 RESISTANCE MSMT	4+/-2	JUMPER/BONDING
23	VER6	GAIN = 10 RESISTANCE MSMT	4+/-2	JUMPER/BONDING
24	VER7	GAIN = 10 RESISTANCE MSMT	4+/-2	MCM SOCKET
25	VER8	GAIN = 10 RESISTANCE MSMT	4+/-2	MCM SOCKET
26	VER9	GAIN = 10 RESISTANCE MSMT	6+/-2	MCM SOCKET
27	VER10	GAIN = 10 RESISTANCE MSMT	6+/-2	MCM SOCKET
28	VER11	GAIN = 10 RESISTANCE MSMT	7+/-2	MCM SOCKET
29	VER12	GAIN = 10 RESISTANCE MSMT	9+/-2	MCM SOCKET
30	VER13	GAIN = 10 RESISTANCE MSMT	3+/-2	FLEX TEST
31	VER14	GAIN = 10 RESISTANCE MSMT	3+/-2	FLEX TEST
32	LOOPCAL2	GAIN = 10 RESISTANCE MSMT	169+/-5	Rcal = 49.9 OHM

(GFE); this impacted the analytical accuracy of the thermal-gradient predictions. The normal experiment sequence for the thermal gradient was as follows: collect a subset of five temperatures during initialization, energize the resistance heater in the thermal simulator MCM for 30 minutes, collect all temperature measurements, power down the experiment for several minutes, power up the experiment, and repeat the data-collection sequence. The flight data clearly reflects the soak temperature, the first rise, the cool down, and the second heat rise. The data also supported the analytical prediction that the maximum heat rise under any condition including faults was 10° C.

2.6.2 *Flight Test*—Flight data collected during two experiment sequences on 26 February 1999 are shown in Table 2. Data was similarly collected approximately every two weeks from February through September and the results never varied by more than one LSB.

Thermal-gradient temperature measurements are shown from the same date in Table 3. The first column is the pre-

heating set of measurements as described earlier. The larger temperature variations in the post-heating data sets reflect the effect of having a concentrated heat source placed in the middle of a field of temperature sensors with varying horizontal distances from the heating source on the reverse side of the panel.

2.7 *Comparison Between Ground Test and Flight Test*

First, it should be noted that the health and status data collected in each measurement cycle was within normal limits, with power-supply outputs always coming in within tolerance and the check sum for the first data set being correct. The conductivity flight data in Table 1 is representative of all further data sets collected. The data did not vary by more than one least significant bit (LSB); this would be a good indication of the stability of the interconnect system used in MFS for this flight.

The thermal data was well within normal limits and varied appropriately in such a fashion to show all temperature sensors working correctly. Varying ambient conditions due

Table 2. Flight-Continuity Measurement Data

BYTE	NAME	DESCRIPTION	VALUE (DEC)	DATA SET 1	DATA SET 2
1	LOOPCAL1	CALIBRATION MSMT	17+/-3	17	17
2	LOOP1	GAIN = 1 RESISTANCE MSMT	0+3-0	0	0
3	LOOP2	GAIN = 1 RESISTANCE MSMT	0+3-0	1	0
4	LOOP3	GAIN = 1 RESISTANCE MSMT	0+3-0	0	0
5	LOOP4	GAIN = 1 RESISTANCE MSMT	0+3-0	0	0
6	LOOP5	GAIN = 1 RESISTANCE MSMT	0+3-0	0	0
7	LOOP6	GAIN = 1 RESISTANCE MSMT	0+3-0	0	0
8	LOOP7	GAIN = 1 RESISTANCE MSMT	0+3-0	0	0
9	LOOP8	GAIN = 1 RESISTANCE MSMT	0+3-0	1	0
10	LOOP9	GAIN = 1 RESISTANCE MSMT	1+3-1	1	1
11	LOOP10	GAIN = 1 RESISTANCE MSMT	1+3-1	1	0
12	LOOP11	GAIN = 1 RESISTANCE MSMT	1+3-1	1	1
13	LOOP12	GAIN = 1 RESISTANCE MSMT	1+3-1	1	1
14	LOOP13	GAIN = 1 RESISTANCE MSMT	0+3-0	0	0
15	LOOP14	GAIN = 1 RESISTANCE MSMT	0+3-0	0	0
16	LOOPCAL2	NOT USED	0	0	0
17	LOOPCAL1	NOT USED	0	0	0
18	VER1	GAIN = 10 RESISTANCE MSMT	2+/-2	1	1
19	VER2	GAIN = 10 RESISTANCE MSMT	3+/-2	3	3
20	VER3	GAIN = 10 RESISTANCE MSMT	3+/-2	3	3
21	VER4	GAIN = 10 RESISTANCE MSMT	3+/-2	3	3
22	VER5	GAIN = 10 RESISTANCE MSMT	4+/-2	3	3
23	VER6	GAIN = 10 RESISTANCE MSMT	4+/-2	3	3
24	VER7	GAIN = 10 RESISTANCE MSMT	4+/-2	3	3
25	VER8	GAIN = 10 RESISTANCE MSMT	4+/-2	3	3
26	VER9	GAIN = 10 RESISTANCE MSMT	6+/-2	5	5
27	VER10	GAIN = 10 RESISTANCE MSMT	6+/-2	5	5
28	VER11	GAIN = 10 RESISTANCE MSMT	7+/-2	5	5
29	VER12	GAIN = 10 RESISTANCE MSMT	9+/-2	7	7
30	VER13	GAIN = 10 RESISTANCE MSMT	3+/-2	2	2
31	VER14	GAIN = 10 RESISTANCE MSMT	3+/-2	2	2
32	LOOPCAL2	GAIN = 10 RESISTANCE MSMT	169+/-5	176	174

Table 3. Flight-Temperature Measurement Data (all in °C)

TEMP SENSOR NO.	MSMT 1	MSMT 2 POST-HEATING	MSMT 3	MSMT 4 POST-HEATING
TEMP1	13.21	16.51	15.79	17.18
TEMP2		16.54		17.66
TEMP3	12.82	16.50	15.58	17.66
TEMP4		17.11		18.25
TEMP5	13.18	17.02	15.95	18.09
TEMP6		17.16		18.23
TEMP7	13.11	17.52	16.06	18.66
TEMP8		17.66		18.83
TEMP9	12.96	17.57	16.06	18.75
TEMP10		17.44		18.53
TEMP11		17.13		18.19
TEMP12		17.38		18.53
TEMP13		17.46		18.65
TEMP14		16.39		17.72
TEMP15		17.17		18.53

to different spacecraft attitudes and flight away from the Sun were reflected in the data with a general trend towards a colder ambient condition. There was no indication of any failed or degraded sensors.

- The design, fabrication, rework/repair, and test of spacecraft panels built without chassis and cabling.
- Hybrid approaches that mix traditional spacecraft chassis/cabling with the MFS design approach.

3.0 TECHNOLOGY–VALIDATION SUMMARY

The following risks were retired with the DS1 MFS demo experiment:

- Use of flex-circuit patches and interconnecting jumpers applied directly to spacecraft panels as an electrical-interconnect system.
- Use of sockets for flight MCMs without risk of opens, shorts, or degradation with time.
- Use of distributed sensors interconnected with flex circuitry to collect data from remote parts of the spacecraft.
- Use of a protective cover that provides an optimum mix of EMI/EMC shielding, radiation shielding and physical protection.

All identified risks that were addressed in the DS1 demonstration were retired. At the time of the experiment conception, the MFS approach was a distinct-paradigm shift from traditional packaging methods to a new system that eliminated the majority of secondary and tertiary packaging to take advantage of the advances in MCM usage. The flight data returned from the experiment did not identify any anomalies and readily met all analytical predictions.

This technology continues to evolve and several organizations are pursuing and/or supporting further improvements and enhancements. Lockheed Martin Corporation (LMA) would be pleased to enter into further efforts to use the MFS packaging system with any interested parties. This would involve the following concerns:

4.0 TECHNOLOGY APPLICATION FOR FUTURE MISSIONS

In the few short years since the MFS experiment was conceived, a number of applications and further demonstrations of the MFS technology have been produced. Hardware using the MFS concepts and “lessons learned” has been supplied to NMP Deep Space Two (flex-circuit interconnects and the tether system), MightySat II Sindri (solar-array interconnect), Space Test Research Vehicle (STRV) Iic, d (experiment and radiation sensor interconnects and entire top panel), Advanced Technology Demonstration Satellite (ATDS) (AFRL/PL demonstration spacecraft), NMP ST5 (in planning), and a variety of further applications in large inflatable structures and nanosats.

LMA is pursuing several enhancements in the MFS technology. These include: demonstration of radio-frequency (RF) pathways in the flex circuitry using alternative dielectrics, optical pathways, production-optimized flex-bonding systems, and integration with inflatable elements.

5.0 ACKNOWLEDGMENTS

The financial support for the MFS development effort was provided by Lockheed Martin Astronautics Independent Research and Development project Lightweight Spacecraft Technology (D-90D), and Air Force-Phillips Laboratory sponsored contract (#F29601-94-C-0167) “Spacecraft Integrated Electronics Structure.” The authors sincerely

thank Dr. Alok Das (AF/PL) for his technical guidance and encouragement. The authors would like to thank their colleagues including Long Nguyen, Keith Sarris, Dan Morgenthaler, and Nathan Harris for their technical contributions.

6.0 LIST OF REFERENCES

- [1] Michael Obal and Janet Sater, "Multi-functional Structures: The Future of Spacecraft Design?" 5th International Conference on Adaptive Structures, Sendai, Japan, December 5 to 7, 1994.
- [2] Joel Sercel et al., "Modular and Multi-functional System in the New Millennium Program," AIAA 34th Aerospace Sciences Meeting and Exhibit, Reno, NV, January 15 to 18, 1996, paper #AIAA96-0702.
- [3] D. M. Barnett and S. P. Rawal, "Multi-functional Structures Technology Experiment on Deep Space 1 Mission," IEEE Aerospace and Electronic Systems, January 1999.
- [4] S. P. Rawal, D. M. Barnett, and F. M. Kustas, "Multi-functional Composite Structures Designs," submitted for publication, 31st International SAMPE Technical Conference, October 1999.
- [5] S. P. Rawal, D. M. Barnett, and D. E. Martin, "Thermal Management for Multi-functional Structures," IEEE Transactions on Advanced Packaging, Vol 22, No. 3, August 1999.
- [6] James Lyke, AFRL/PRS/VSSE, Personal Communications 1995 to 1999.

Appendix A. DS1 Technology Validation Telemetry Channels

MULTI-FUNCTIONAL STRUCTURE

Channel	Mnemonic
O-0051	MFS_mgr_stat
O-0052	MFS_last_cmd
O-0053	MFS_wrds_snt
O-0054	strt_cmd_cnt
D-0192	last_pkt_12
D-0193	buf_typ_12
D-0194	buf_min_12
D-0195	buf_max_12
D-0196	pkt_age_12
D-0197	buf_pkt_12
D-0198	sent_pkt_12
D-0199	spac_used_12
D-0200	bytes_ack_12
D-0201	byte_dump_12

Appendix B. DS1 Technology Validation Power On Times

MULTI-FUNCTIONAL STRUCTURE

MFS initial turn-on was 02/25/99.

Experiment was then conducted bi-weekly from power-off.

Experiment was also conducted weekly with LPE/PASM starting 05/26/99.

Plasma Experiment for Planetary Exploration (PEPE)

DS1 Technology Validation Report



Prepared by:

David T. Young

University of Michigan

College of Engineering

Department of Atmospheric Oceanic and Space Sciences

Tel.: 734-936-7263

e-mail: dtyoung@umich.edu



Jane E. Nordholt

Los Alamos National Laboratory

Space and Atmospheric Science Group

NIS-1

Tel.: 505-667-3897

e-mail: jnordholt@lanl.gov



John J. Hanley

Southwest Research Institute

Space Research and Instrumentation Division

Tel.: 210-522-2884

e-mail: jhanley@swri.edu




Table of Contents

<u>Section</u>	<u>Page</u>
Extended Abstract	v
PEPE Fact Sheet	viii
1.0 Introduction	1
2.0 Technology Description	1
2.1 Overview.....	1
2.2 Key Technology-Validation Objectives at Launch.....	3
2.2.1 Miniaturized 3-dimensional Linear Electric Field (LEF3D) Time-of-flight (TOF) Optics	3
2.2.2 High-Speed TOF Electronics	3
2.2.3 Integrated Ion/Electron Optics	3
2.2.4 Data-Acquisition System.....	3
2.2.5 High-Voltage System.....	3
2.2.6 High-Density Packaging Architecture	3
2.3 Expected Performance Envelope	3
2.4 Detailed Description	5
2.4.1 Miniaturized 3-dimensional Linear Electric Field (LEF3D) Time-of-flight (TOF) Optics	5
2.4.2 High-Speed TOF Electronics	6
2.4.3 Integrated Ion/Electron Optics	6
2.4.4 Data-Acquisition System.....	6
2.4.5 High-Voltage System.....	8
2.4.6 High-Density Packaging Architecture	8
2.5 Technology Interdependencies	9
2.5.1 PEPE Plasma Spectrometer Technology	9
2.5.2 Data-Acquisition and High-Voltage Systems	9
2.6 Test Program.....	9
2.6.1 Ground Test.....	9
2.6.2 Flight Test.....	10
2.7 Comparison Between Ground and Flight Test.....	14
3.0 Technology Validation Summary	14
4.0 Technology Application for Future Missions	14
5.0 List of References	15

Figures

<u>Figure</u>	<u>Page</u>
Figure 1. Photo of PEPE on the Bench Prior to Delivery	1
Figure 2. Cross Section Illustrating the Location of PEPE Subsystems and Layout of PEPE Ion/Electron-Optical System.....	2
Figure 3. Schematic Block Diagram of the PEPE Electronic Subsystems and Their Relationship to Sensor Elements.....	4
Figure 4. Cross-sectional Detail of the PEPE TOF Cylinder.....	5
Figure 5. Data from 26 Jan. 1998 0000-0600 UT Illustrating PEPE's Response to Quiescent Solar-Wind Plasma	7
Figure 6. Data from 1 March 1998 0600-1200 UT Illustrating PEPE's Response to Disturbed Solar-Wind Conditions.....	7
Figure 7. Data from 20 April 1998 0600-1200 UT Showing the Startup (~0905 UT) and Operation of the IPS Xe+ Thruster.....	8
Figure 8. TOF Spectrum Based on Direct Event TOF Data.....	11
Figure 9. Spectra Comparing Solar-Wind Data From PEPE and the WIND/SWE.....	13
Figure 10. Daily Average TOF Counts for a Day Without IPS Firing (Day 009, 1999), a Day With Continuous IPS Firing Early in the Mission (Day 083, 1999), and a Day Much Later in the Mission With Continuous IPS Firing (Day 216, 1999)	14

Appendices

<u>Appendix</u>	<u>Page</u>
Appendix A. PEPE's Telemetry Channels.....	16
Appendix B. PEPE's Date of Turn-on and Provisional Times-of-Data-Capture List	17

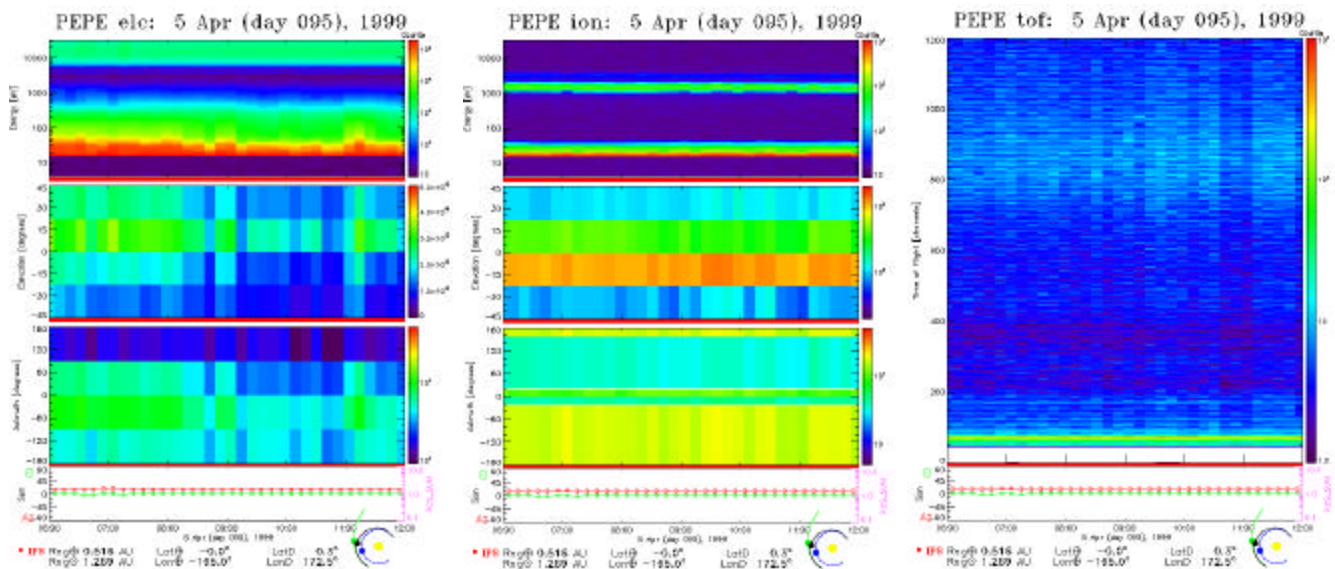
EXTENDED ABSTRACT

The Plasma Experiment for Plasma Exploration (PEPE) is a particle spectrometer capable of resolving the energy, angle, and mass composition of a wide range of plasmas found throughout the solar system. PEPE commenced successful operations on 8 December 1998. As a part of the Deep Space 1 (DS1) mission, the objectives of the PEPE investigation are to demonstrate new instrumentation technologies relevant to low-resource space plasma instrumentation, to show that such instruments can be operated successfully to obtain high-quality scientific data on a spacecraft employing an ion propulsion system (IPS), and to obtain new scientific findings related to the prime scientific targets of the DS1 mission. The three broad categories of new technologies demonstrated in the PEPE instrument include novel electron and ion optical systems, including an electrostatically swept field-of-view and time-of-flight mass analysis, that significantly reduce overall sensor mass and volume relative to performance; a compact, high-reliability, high-voltage system consisting of eight individual supplies ranging from ± 3.6 to ± 15.0 kV; and low-resource, high-performance electronics that perform sub-nanosecond measurements and provide very flexible data acquisition and processing capabilities.


Several categories of risk were associated with the PEPE program from its inception. Technological risks associated with the instrument manufacture included the use of novel materials and processing techniques that had no previous flight history. In particular, previously untried methods had to be developed to metal-plate and chemically etch low-

outgassing exotic plastics. A second new process was required to vapor-deposit a variable-depth, extremely high-ohmic coating on high-purity ceramic cylinders. At the time we were assured by collaborating technologists that no one had ever tried to accomplish these tasks for ground applications, much less for spaceflight. A second risk category was overall system design. Although PEPE was based on previous experience and designs, the entire instrument was built without the benefit of a prototype or an engineering development unit for the optical subsystem, high-voltage subsystem, or for the system as a whole. A few selected subsystems were prototyped, primarily to develop and test digital interfaces. Schedule and budget comprised a third risk category: to our knowledge no instrument of this complexity has been built in a period of 26 months from the contract start date to delivery at the launch site. Schedule risk was managed largely through the contribution of prolonged work hours by a small and dedicated team.

In most cases, validation of the PEPE concept and technologies is being obtained primarily by examining data obtained in space from the instrument and inferring subsystem performance indirectly. A number of subsystems, mostly digital electronics, were tested and validated during ground tests. However, our primary test case consists of the thoroughly studied characteristics of the solar-wind plasma that is simultaneously being observed by the Wind spacecraft and the Advanced Composition Explorer (ACE) spacecraft located near the Earth. Together with solar-wind instruments on other spacecraft located elsewhere in the solar system (Solar and Heliospheric Explorer [SOHO], Ulysses, Cassini), PEPE provides a valuable contribution to



Examples of PEPE Data Returned on 5 April 1999 When the IPS Was Running.
 The first two panels show the energy and angular distributions for electron and ions.
 The third panel shows the time-of-flight spectrum summed over all energies and angles.



PEPE's Mounting Position on DS1. PEPE is color coded to show the electron section in green, the ion and data processing section in blue, and the aperture in red.

the study of large-scale solar-wind structures because of its location in a part of the solar system widely distant from the other spacecraft. Thus, a careful analysis of observations of solar-wind ions and electrons that have been collected since PEPE switch-on in December 1998 serve to validate the overall end-to-end performance of the optical design, high-voltage system, time-of-flight electronics, and other technologies. Particular details of the measured shape and intensity of the solar-wind velocity distributions give specific information about optical alignments, carbon foil and detector efficiencies, high-voltage system performance, and end-to-end system performance.

Although validation of the PEPE instrument and technologies is a primary concern of the program, it was realized early on that because of the compacted development schedule we would have to forgo a considerable amount of ground testing and calibration activities. Individual electronic subsystems received enough ground testing to validate their specified performance; however, combinations of subsystems often received little more than interface checks to validate their compatibility. Because the complete PEPE system came together only very late in the


program, testing at the system level was minimal in the extreme. For example, the thermal vacuum test consisted of a single cycle that was combined with an attenuated calibration period lasting only 2 days! The hot part of the cycle also served as the bakeout period for the instrument prior to calibration. This deficit was to be made up in flight by accumulating a large number of operating hours in different environments and instrument operational modes and by in-flight calibration using targets of opportunity involving similar instruments on other spacecraft.

To a large extent it has been possible to gather the data necessary for validation. The data taken so far with PEPE compare very favorably with solar-wind data obtained by plasma instruments on the Wind, ACE, and Cassini spacecraft after allowances are made for the structure and evolution of the solar wind and the separation distances of the respective spacecraft from DS1. In addition, PEPE data have been used to demonstrate that high quality measurements of plasma at energies above roughly 50 eV can be made with the IPS operating. Below this energy PEPE has obtained measurements of xenon ions as well as secondary electrons related to both the IPS and SCARLETT

solar arrays. These data can be used to map the previously

unobserved local plasma environment of an IPS-driven spacecraft—a topic of considerable interest for future planetary and heliospheric missions.

PEPE technologies can be put to future use in two ways: as an entire sensor technology and as a set of subsystem technologies. As an integrated system, PEPE provides nearly the same capability as the state-of-the-art Cassini Plasma Spectrometer; however, at 5.5 kg and 9.6 W, it requires only 24% of the mass and 46% of the electrical power of the latter. In addition there are qualitatively different approaches used in PEPE that simplify its use in future missions. Typical planetary spacecraft such as Cassini and DS1 are 3-axis stabilized, which presents a problem for plasma instruments that typically need to view as much of the unit sphere as possible. The ideal is 4π steradian coverage, which presents a significant problem to both the spacecraft and instrument designers. On Cassini, this problem was solved by scanning the plasma instrument mechanically using a 3.6 kg motor with attendant problems related to magnetic cleanliness and mechanical stability needed for fine pointing of optical sensors (both problems were solved on Cassini). With PEPE, the problem was approached for the first time by employing an electrostatically scanned field-of-view using no moving parts or magnetic fields. Thus, the PEPE technology has wide appeal for future missions because it eliminates possible magnetic and mechanical interferences. Other PEPE subsystem technologies having wide future applicability include very low resource high-voltage power supplies and compact time-of-flight mass, spectrometer optics and associated electronics.



Comparison Between the Cassini Plasma Spectrometer (CAPS) and PEPE.

ACKNOWLEDGEMENTS

The authors wish to thank the PEPE teams at Southwest Research Institute (SwRI) and Los Alamos National Laboratory (LANL) for their hard work during the development phase of the instrument. These scientists and engineers are: J. L. Burch¹, D. J. McComas², R. A. Abeyta², J. Alexander¹, J. Baldonado², P. Barker², R. K. Black¹, T. L. Booker¹, R. P. Bowman¹, P. J. Casey¹, L. Cope², J. P. Cravens¹, N. Eaker¹, H. O. Funsten², R. Goldstein¹, D. R. Guerrero¹, S. F. Hahn², B. P. Henneke², E. F. Horton¹, D. J. Lawrence², K. P. McCabe², S. Rogillio¹, R. P. Salazar², E. Santiago², M. Shappirio², S. A. Storms², and C. Urdiales¹. We also thank the two institutions and the Department of Energy for their support and dedication to PEPE. We would particularly like to thank SwRI for investment of internal research support to help develop several of the PEPE technologies. Development and operation of PEPE would not have been possible without the support of members of the JPL DS1 team; in particular, L. Livesay, K. Clevin, D. Lehman, R. Nelson, K. Fleming, P. Chadbourne, M. Rayman, and P. Varghese. We particularly want to thank G. Withbroe of NASA HQ for his generous support of the PEPE instrument. We would also like to thank R. Christy and C. Gerber of TriboCoat Inc. for their willingness to try some new things. The development phase of PEPE was supported at SwRI by JPL contract 961207 and at LANL by NASA contracts WO-9049 and WO-9066; the mission operations and data analysis phase is supported at SwRI by JPL contract 960619 and at LANL by WO-9165 and WO-9138.

¹: SwRI


²: LANL

PEPE Fact Sheet

Parameter	Range/Resolution	Performance	Units
Sensor Type	Toroidal electrostatic angular scanning and energy/charge analyzers coupled to linear-electric-field time-of-flight ion mass/charge analyzer.		
Energy	Range	8.0 to 33,500	eV/eV
	Range scanned	120 steps, log-spaced	
	Resolution (ions)	0.046	$\Delta E/E$
	Resolution (electrons)	~0.085	$\Delta E/E$
	Analyzer constant (ions)	13.07	
Mass	Range	1 to 135	amu/e
	Resolution (medium mode)	~4	M/ ΔM
	Resolution (high mode)	~20	M/ ΔM
Angle	EL angle range (scanned) –45 to +45	(°)	
	EL analyzer deflection	$6.7 \times 10^5/(E/Q)$	(°)
	Range scanned	16 steps, linear-spaced	
	AZ angle range (static)	360	(°)
	Solid angle coverage	8.9	sr
	Resolution (electrons)	256 pix @ 5×22	(°) \times (°)
	Resolution (ions)	128 pix @ 5×5	(°) \times (°)
	Resolution (ions)	32 pixels @ 5×22	(°) \times (°)
Temporal	AZ angle \times TOF	0.008/0.032	s
	AZ angle \times EL angle \times TOF	0.128/0.512	s
	AZ angle \times EL angle \times energy \times TOF	16.38/65.54	s
Sensitivity	Electrons ($5^\circ \times 22^\circ$ pixel, $\epsilon \sim 0.5$)	$\sim 1.5 \times 10^{-4}$	cm ² sr \times cts/el.
	Ions ($5^\circ \times 22^\circ$ pixel, $\epsilon \sim 0.5$)	$\sim 8.0 \times 10^{-5}$	cm ² sr \times cts/ion
	(@8.0 kV TOF, $\epsilon \sim 0.2$)	$\sim 3.0 \times 10^{-5}$	cm ² sr \times cts/ion
Dynamic Range	Electrons	0.1 to 10^6	Hz
	Ions (singles)	0.1 to 10^6	Hz
	Ions (TOF analyzed)	0.01 to 10^5	Hz
Resources	Mass	5.5	kg
	Power	9.6	W
	Volume	~7.25	liters
	Density	0.83	g/cm ³
	Telemetry (commandable)	1024, 512, 250, 100, 50, 25	bits/s
	Location on spacecraft	+Z edge of the +X-Y face of s/c	
	Operating range	–20 to +35	C
	Cover (with GN2 purge)	Remove before flight	
Performance	Operating time (as of 11/1/99)	~5600	hours

Contacts: David T. Young Jane E. Nordholt
 University of Michigan Los Alamos National Laboratory
 Ann Arbor, MI 48109 Los Alamos, NM 87544
 dtyoung@umich.edu jnordholt@lanl.gov

John J. Hanley
 Southwest Research Inst.
 San Antonio, TX 78228
 jhanley@swri.edu



Plasma Experiment for Planetary Exploration (PEPE) DS1 Technology Validation Report

*David T. Young, University of Michigan, College of Engineering
Jane E. Nordholt, Los Alamos National Laboratory, Space and Atmospheric Science Group
John J. Hanley, Southwest Research Institute, Space Research and Instrumentation Division*

1.0 INTRODUCTION

Unlike all of the technologies onboard the Deep Space 1 (DS1) mission (except MICAS), Plasma Experiment for Planetary Exploration (PEPE) (shown prior to delivery in Figure 1) is both a spacecraft technology and a self-contained scientific instrument [1, 2]. PEPE itself incorporates a half-dozen technologies that are to our knowledge novel to space-plasma instrumentation. The technologies were developed in response to the need for greatly reduced resources relative to comparable instrumentation on other missions, such as Cassini. Because the two teams that built PEPE (Southwest Research Institute [SwRI] and Los Alamos National Laboratory [LANL]) had also designed and built much of the Cassini Plasma Spectrometer (CAPS) (see [3, 4]), a decision was made to design an instrument that would maintain the performance envelope of CAPS while at the same time reducing the resource envelope by a significant fraction. The PEPE resources were also dictated by their availability and allocation on the DS1 spacecraft.




Figure 1. Photo of PEPE on the Bench Prior to Delivery

PEPE required a number of breakthroughs: a complete redesign of the CAPS particle-optics system, new ways of miniaturizing and incorporating the large number of high voltages required to drive the optics, and miniaturization of critical circuits needed for high-speed (~1 GHz) time-of-flight (TOF) measurements and for data acquisition and compression. Fortunately, the radiation hardness required of the DS1 program did not place any stringent requirements on the PEPE electronic-parts procurements. We also

attempted several new materials treatments and processes, including depositing-uniform coatings of very high ohmic materials on ceramics, metal plating of relatively inert plastics, and complex, multi-layer electrical boards containing ion-optical components.

All of these technologies and their attendant risks are discussed in detail in Section 2. The reference point for much of the discussion is the Cassini Plasma Spectrometer described in [2, 3]. Validation of the instrument technologies has required careful analysis of their performance in a number of situations using the ambient solar-wind plasma, the spacecraft-photoelectron sheath, and the products of the xenon ion propulsion system as test opportunities. Unfortunately, the time-of-flight system has not been able to operate at its planned high-voltage level. This topic will be discussed along with other validation topics in Sections 2.7 and 3. In Section 4 we will discuss the use of PEPE and related technologies for future missions. Appendices A and B give technical details on the PEPE data channels and data collection periods.

2.0 TECHNOLOGY DESCRIPTION

2.1 Overview

PEPE is a charged-particle spectrometer capable of measuring and resolving the velocity distribution of electrons and ions and the mass composition of ions that make up the wide variety of plasmas found in the solar system. In order to coincide with the scientific objectives of the DS1 mission, the particular design chosen for PEPE focuses on measuring solar-wind plasma and the plasma populations resulting from solar-wind interactions with intrinsic plasmas associated with the outgassing of asteroids and comets. However, the general concepts and technologies used in PEPE can be adapted readily to a wider variety of objectives and missions, in particular missions to study planetary magnetospheres. A major driving factor in the design of PEPE was to reduce its resource requirements relative to those of instruments with comparable capabilities. In this case, we turned to the Cassini Plasma Spectrometer, a very high capability instrument presently operating on the Cassini-Orbiter spacecraft. Because the core design teams of the two instruments are the same, the goals for the PEPE design consisted of trying to duplicate the main performance features of the Cassini instrument in a much lower resource instrument. It was recognized at the outset that PEPE could not exactly duplicate these features

because the DS1 mission did not require it and because performance compromises would have to be made in some areas in order to meet resource targets.

With reference to the cross section of the PEPE instrument assembly shown in Figure 2, PEPE is made up of four functional components that are integrated using a novel architecture: (1) a series of charged-particle optical elements; (2) a system of high-voltage supplies that establish bias voltages needed for particle-optical elements and detectors; (3) high-speed pulse electronics that make

timing measurements used to discriminate ion mass; and (4) digital electronics that provide data-acquisition and instrument-command-and-control functions. These are integrated using a packaging architecture that draws the subsystems together in a single, compact, low-resource instrument. The functional components of PEPE make use of several technology applications that are either newly developed for PEPE or are new applications of existing technologies, such as the system of Field Programmable Gate Arrays (FPGAs) used to make up PEPE's powerful and flexible data-acquisition system.




Figure 2. Cross Section Illustrating the Location of PEPE Subsystems and Layout of the PEPE Ion/Electron-Optical System

2.2 Key Technology-Validation Objectives at Launch

The primary-validation objective for PEPE, can be summarized as end-to-end functionality that meets requirements for scientifically useful data products. The objectives of developing a reduced resource instrument can be validated by simple measurement of volume, power, mass, and data rate if the functionality requirement is met. Each of the PEPE technologies can be analyzed and validated against the descriptions and requirements that will be described in the following sections. However, the paramount issue in validating PEPE is the contribution of each technology to overall performance. With that overarching goal in mind, we can reconsider each of the 6 PEPE technologies.

2.2.1 Miniaturized 3-dimensional Linear Electric Field (LEF3D) Time-of-flight (TOF) Optics—The LEF3D-conceptual design is based on that of the LEF3D used in the CAPS instrument’s Ion-Mass Spectrometer (IMS) [5, 6, 7]. The cylinder technology rests primarily on the use of high-resistance surface coatings in place of a set of discrete-ring electrodes on the CAPS/IMS. A second departure was to redesign interfaces between the high-voltage supplies and the LEF3D’s optical elements. Validation objectives include being able to apply the target ± 15 kV high voltage to the cylinder without high-voltage breakdown and operating the cylinder stably for a period comparable to typical mission lifetimes of ~ 2 years. The LEF3D optics should deliver TOF spectra with mass resolution equivalent to $M/\Delta M \sim 20$ based on ray-tracing and experience with the IMS. Resolution and mass range as well as species rejection of the LEF3D optics must also be validated.

2.2.2 High-Speed TOF Electronics—The TOF electronics consist of a high-speed front-end electronics (FEE) that includes amplifiers, discriminators and logic, and a time-to-digital converter circuit with associated logic. The TOF electronics should deliver performance with ~ 1 GHz bandwidth that is consistent with the TOF resolution required of the cylinder. This performance should be nearly identical with that of the IMS [6, 7] although the resources required should drop by $\sim 50\%$. Built-in test functions and in-flight validation of the FEE and time-to-digital converter (TDC) will be carried out using TOF data from the LEF3D.

2.2.3 Integrated Ion/Electron Optics—Using the solar wind as a well-studied and constantly-monitored plasma source, we shall confirm the energy and angle resolution, the correct angular orientation and location of elements of the field of view (FOV), the energy and field-of-view scanning functionality, and the absolute detector response of the sensor. The efficiency of the anti-reflection surface treatments will be validated by measurement of the extent to which solar UV and particles are scattered into the sensor.

2.2.4 Data-Acquisition System—An onboard pulser will be used to create fixed-pattern artificial TOF spectra that can be acquired and compared with ground-based calibration. A second and more stringent validation will be achieved by processing the high-counting-rate random electron and TOF events caused by the solar wind and other naturally occurring plasmas. Solar-wind data from other space-borne instruments on the WIND and ACE spacecraft located near the Earth will be compared with the processed PEPE data to determine that all components of the PEPE data product are correct and free from artifacts.

2.2.5 High-Voltage System—The high-voltage system will be activated and brought up to full operating levels singly and in combinations required for spectrometer operation. Data from the high-voltage supply monitors and from the background intervals during the high-voltage scans will be used to measure and track detector-noise levels in order to ascertain long- and short-term operation criteria for drift stability and ripple. Automatic high-voltage (HV) turn-on sequences will be prepared and executed without operator intervention. The goal was to be able to turn on the HV system automatically within a period of 4 hours or less without intervention.

2.2.6 High-Density Packaging Architecture—If the optical and high-voltage systems’ (sections 2.2.1, 2.2.2, 2.2.3, and 2.2.5) performances are nominal, the packaging architecture will be considered validated. The PEPE data acquisition technology (section 2.2.4) is largely unaffected by the architecture. The fact that PEPE’s instrument density of 0.83 g/cm^3 is significantly higher than similar plasma spectrometers (values range from ~ 0.25 to $\sim 0.5 \text{ g/cm}^3$) indicates that this design feature has been validated, provided PEPE functions correctly in flight.

2.3 Expected Performance Envelope

The scientific objectives for PEPE or any other plasma spectrometer require that it measure the N-dimensional particle-phase space consisting of 3 velocity coordinates and N-3 mass/charge coordinates in the frame of reference of the spacecraft. The time coordinate and 3 spatial and 3 attitude coordinates are required to put the plasma data in the proper context. These fiducial data are obtained from the spacecraft and, although they could affect PEPE’s ability to deliver valid measurements, by convention the spacecraft team is responsible for that aspect of performance. The specific elements of performance of plasma instruments such as PEPE concern: a) the range of parameter space coverage, b) resolution within that range, and c) sensitivity for detecting charged particles within that range. The parameters that have to be measured are those that determine the details of the particle-distribution functions: namely, energy/ charge, mass/charge, and angle of the direction of particle arrival. Because the plasma environment is highly dynamic, it is

also important that the entire range of measurement be covered in as short a time interval as possible.

The performance envelope is dictated not only by the measurement objectives mentioned above, but by the capabilities of the measurement technology. PEPE consists of an electrostatic ion-optical system (magnetic systems are impractical under the circumstances) and its supporting electronics. The optics in turn depend on the correct shape and location of the optical electrodes and the application of the correct voltages to correctly bias the optical elements at a given instant in time. Electrode shape and position were maintained through the usual design and manufacturing processes to an estimated ± 0.005 inches. The magnitude and number of distinct voltages required is determined by the optical design. In the case of PEPE, eight HV supplies are required with voltages ranging from ± 3.6 kV for the two micro-channel plate (MCP) detectors to ± 15 kV for the main TOF voltages. The latter are the highest voltages in PEPE and, in part, determine the mass-resolution performance.

In order to cover a range of particle angles of arrival and their energies, two sets of coupled power supplies drive the deflection and energy-analysis optics (Figure 3). One set of supplies produces two fast-slewing (3×10^6 V/s) bi-polar supplies that deliver ± 5.0 kV to deflect particles $\pm 45^\circ$ in elevation. The deflection supplies float with respect to two “bulk” supplies (Figure 3) as do the ESA supplies. The nominal dwell time at a single elevation step is 0.032 s, giving a rate of 0.512 s for a single, full-elevation scan of 90° . Similarly, the HV set driving the energy analyzer section is also swept, but at a lower rate of one step for each elevation scan. In this way, an entire scan of the PEPE range of energies and angles (TOF- mass measurements occur at every sample) takes place in 65.54 s. These times are sufficient for monitoring the solar wind (usual temporal resolution is 1 to 5 minutes). Higher time resolution needed for rapid flybys of asteroids and comets can be obtained by using shorter dwell times, restricted scan ranges that can be covered more rapidly. Power-supply scan patterns and dwell times are programmable (as explained in section 2.4.4).

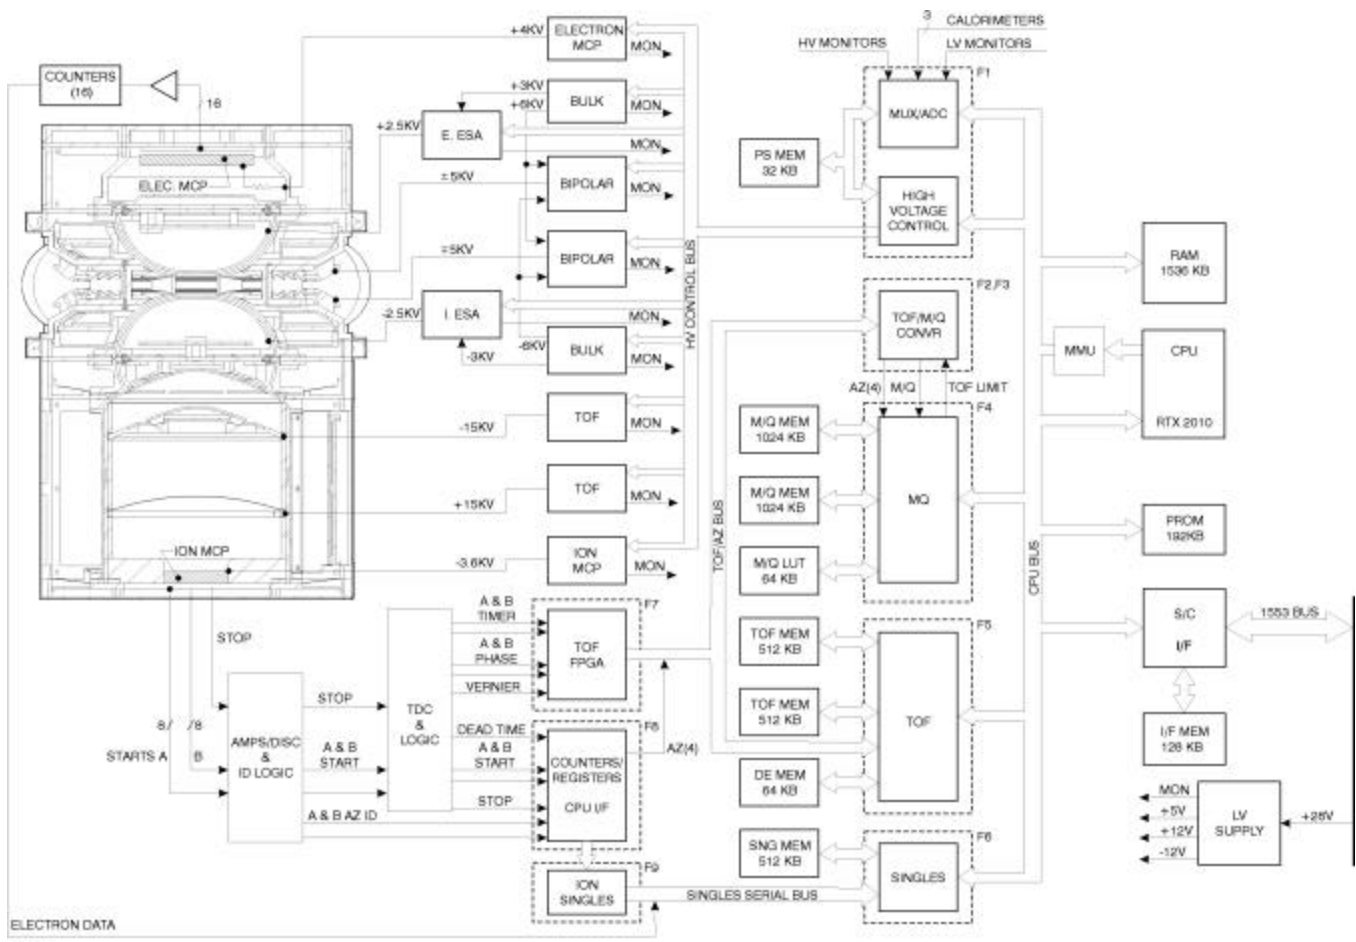


Figure 3. Schematic Block Diagram of the PEPE Electronic Subsystems and Their Relationship to Sensor Elements

Instrument sensitivity is determined by: 1) The instrument aperture (related in turn to the available volume and mass allocated to the instrument) and 2) the energy/charge and angular-resolution requirements. The latter set the size of the gaps between the field elements and determine the high voltages needed to establish the electrical forces across them (since electrical force \sim voltage/optical-element separation). Sensitivity is usually given in units of $[\text{cm}^2 \text{sr}] \times$ [detector counts/incident target particle] (see the PEPE Fact Sheet). Thus, sensitivity and resolution are directly related to the mass and electrical power allocated to the instrument. The design of PEPE was meant to optimize sensitivity for a given set of resources; however, it is difficult to normalize the PEPE performance per unit resources relative to that of other instruments. The ultimate validation is the fact that PEPE obtains excellent solar-wind measurements at relatively high resolution with a fraction of the resources of existing instruments. The comparable plasma analyzers on the WIND spacecraft [8], for example, have about the same range of energy and angular acceptance as PEPE, but do not have either mass/charge analysis capability or a swept FOV (WIND is a spinner).

The WIND instruments are combined with several others in a package so that only very rough estimates can be made of their weight and power; however, it appears that they are comparable to those of PEPE. Since PEPE includes the added features of TOF mass spectrometry and a swept field-of-view (WIND is a spinning spacecraft and does not require a swept FOV), we conclude that PEPE has perhaps a factor of 2 advantage in performance for the same mass. One other figure of merit is the mass density of the packages: PEPE's ratio of mass to volume is 0.75 g/cm^3 , whereas that of the WIND instrument is $\sim 0.25 \text{ g/cm}^3$, similar to that of CAPS. An informal survey of plasma spectrometers shows that instrument density is typically 0.25 to 0.5 g/cm^3 , indicating that PEPE's goal of producing relatively high-packaging density has been achieved.

2.4 Detailed Description

2.4.1 Miniaturized 3-dimensional Linear Electric Field (LEF3D) Time-of-flight (TOF) Optics

Figure 4 shows a cross-section of PEPE's LEF3D cylinder together with characteristic particle trajectories and major features of the instrument. CAPS' Ion Mass Spectrometer (IMS) contains an LEF3D spectrometer similar to PEPE but larger in volume by a factor of ~ 8 . The IMS is made up of 30 discrete field-rings joined by a series of resistors. This arrangement produces the electric field configuration needed to make high-resolution TOF measurements. The approach taken with PEPE was to eliminate the rings altogether, replacing them with a monolithic ceramic cylinder coated with a layer of high-resistance, vapor-deposited chromium oxide. The volume of the cylinder materials was reduced by a factor of ~ 5 . The cylinder itself was made smaller by a factor of ~ 2 and its aspect ratio (height to width) reduced to produce a

design that is volumetrically smaller by a factor of ~ 8 . The key performance factor—namely, the resolution of ion-flight-path timing—was reduced by a factor of 2 owing primarily to shorter flight paths, but was still acceptable for PEPE's science objectives.




Figure 4. Cross-sectional Detail of the PEPE TOF Cylinder

When ions exit the curved analyzer plates (section 2.4.3), a large negative potential accelerates them into thin ($1 \mu\text{g/cm}^2$ or $\sim 50 \text{ \AA}$ thickness) carbon foils. The foils form the entrance to the time-of-flight mass/charge analyzer. Electrons that are released by the passage of the ion through the foil are accelerated onto the outer annulus of the MCP stack at the bottom of the TOF section. They all take a uniform time of about 2ns to get to the MCP. There they start a clock that will be used to determine the time-of flight for the ion that released them. Ions are generally neutralized by their passage through the foils. In this case, they continue down the TOF section until they (with high probability) strike the center section—or “stop” section—of the MCP. This stops the clock and standard time-of-flight mass spectroscopy. The knowledge of the ion's initial energy to determine its velocity is used to determine the ion's mass. Mass resolution $M/\Delta M$ is only about 5 for this process. If the ion remains charged and its energy is not too large to be turned around by the high voltage on the curved grid (heavy-curved line in Figure 4), it will “bounce” in the linearly-increasing electric field just as a mass on a spring would. The time for one-half oscillation of this bounce is independent of energy or angle of flight of the ion; therefore, the TOF is proportional to the square root of only the mass/charge and, thus, the mass resolution for TOF of this type is much higher than in simple field-free TOF systems. Ions that do bounce hit a secondary emitter at the top of the TOF section and the resulting electrons are drawn to the center “stop” portion of the MCP. The mass resolution in the case is calculated at ~ 20 .

All three regions (acceleration, LEF3D, and deceleration) of the TOF section use ceramic cylinders with resistive coatings to produce in them a uniform-electric field; however, only the center section requires the coating to have varying thickness to produce the linearly-increasing electric field.

2.4.2 High-Speed TOF Electronics—The CAPS instrument relied on TOF electronics capable of 750 ps (10^{-12} s) resolution and pulse-pair resolution of 40 ns. The timing electronics required high-speed amplifier discriminator chains and logic (referred to as front-end electronics [FEE]) and a time-to-digital converter (TDC) that required a significant amount of power and component-board space. With reference to the functional block diagram shown in Figure 3, the PEPE design maintains the CAPS functionality and, in addition, doubles the number of angular-position channels encoded in order to capture finer details of the solar-wind ion distribution. Because PEPE's TOF optics operate in substantially the same way as those of the CAPS instrument, they required similar performance but with reduced resources. The FEE and TDC circuits were redesigned using chip-on-board technology. In addition, a direct digital-encoding scheme was incorporated to register the increased number of angular channels. The PEPE timing circuits required about 50% less power than the IMS unit and occupied about 40% less board space.

2.4.3 Integrated Ion/Electron Optics—At the time that CAPS was designed, two entirely separate instruments were needed to measure electrons and ions. This required two sets of housings, separate-entrance collimators and fore-optics, separate high-voltage supplies to drive the two electrostatic analyzers, and separate mounting locations to obtain clear fields-of-view. The duplication of functions required a fairly high investment in resources. In addition, because Cassini, like DS1, is a 3-axis stabilized spacecraft, CAPS required some way to articulate its field-of-view in order to sample the wide range of viewing space occupied by target plasma distributions. On Cassini, the solution was to use a motor/actuator that rotated the entire CAPS instrument (weighing 20 kg) over a range of $\pm 104^\circ$. PEPE was designed so that electron and ion optics share a common entrance aperture that eliminates duplication of this optical element (Figure 2). After crossing the collimator, electrons and ions enter an electric field region created by the inner electrodes of two electrostatic-energy analyzers (ESA) (Figure 1). In a manner similar to the CAPS design, the energy analyzers are cylindrically symmetric, an arrangement that allows the instrument to view over a range of 360° in the plane perpendicular to Figure 2. Unlike CAPS, however, a set of toroidal electrodes was placed just outside the PEPE entrance aperture to deflect incoming ions and electrons in such a way that the instrument FOV can be scanned over a range of $\pm 45^\circ$ in the plane of Figure 2. The result is that PEPE covers a larger range of observation

space with dramatically reduced resources allocated to the optical system. The toroidal-deflection electrodes create an electric field that is terminated at the surface of the PEPE instrument by a toroidal-shaped wire-mesh grid. In order to ensure grid shape and stability, the wire-mesh material was formed in a vacuum-driven jig that shaped the grid and held it in place while the edges of the grid material were epoxied to the grid frame. The resulting toroid mesh was then plated with a thin Ni coating to bond the wires into place and further guarantee the shape. The electron optics made use of a dynode structure that converts the incoming electron flux into low-energy secondary electrons that can be easily concentrated onto a small MCP detector. This device reduces the size and complexity of the MCP that would otherwise be required to cover the exit aperture of the electron-energy analyzer. After the PEPE optical electrodes were machined and metal plated with nickel and copper, they were treated with the Ebanol-C process that develops a thin layer of anti-scattering microscopic crystals that are both rough and black.

2.4.4 Data-Acquisition System—The primary data product of CAPS' IMS TOF system is two 2048-channel TOF spectra generated every 62.5 ms. CAPS' maximum data rate is 16 kbits/s. In contrast, the maximum PEPE downlink rate is only 1024 bits/s for a similar amount of raw data from the TDC. This requires a very high degree of onboard capability for restructuring the angle/energy sweep program and the compression of the resulting data products. In order to provide this kind of a flexible program for optimum data return under all conditions likely to be encountered during the mission, PEPE was equipped with programmable control-of-sample dwell time (factor of 8), a voltage-scan program, and data-acquisition and processing capability. The PEPE system is functionally comparable to that of the CAPS instrument but uses about one third of the resources. However, PEPE relies more heavily on FPGA's (ACTEL 1280) than did CAPS and less on its low-resource processor (RTX2010).

Spectral scanning is carried out by setting the deflection and ESA supplies to their highest commanded levels and then stepping the supplies down. The deflection supply completes a scan and then the ESA is stepped. The nominal method of scanning is the "survey" mode (see Figure 5, Figure 6, and Figure 7) that covers the full 16-step elevation \times 128-step energy scan (8 of the energy scan steps are used for background measurement). If the target-plasma population is restricted in velocity space (e.g., the solar-wind or cometary ionospheres) or the DS1-mission strategy requires that PEPE restrict its data rate, then PEPE can be programmed in several ways that provide more optimal data return. The simplest way to lower the data rate is to integrate the spectra over longer intervals (up to 20 minutes, equivalent to 50 bits/s). More efficient scans can be made by targeting a restricted volume of phase space and creating a

smaller region-of-interest (ROI) angle/energy scan that can be scanned at a higher rate than nominal. A second way of producing a ROI is simply to select a subset of the current spectrum (whether full or ROI) and send back only those products.

Data products (electrons, ion singles and TOF, housekeeping) from a completed energy \times angle \times angle \times TOF spectrum are acquired into separate memory arrays (Figure 3). After acquisition of the current spectrum, the

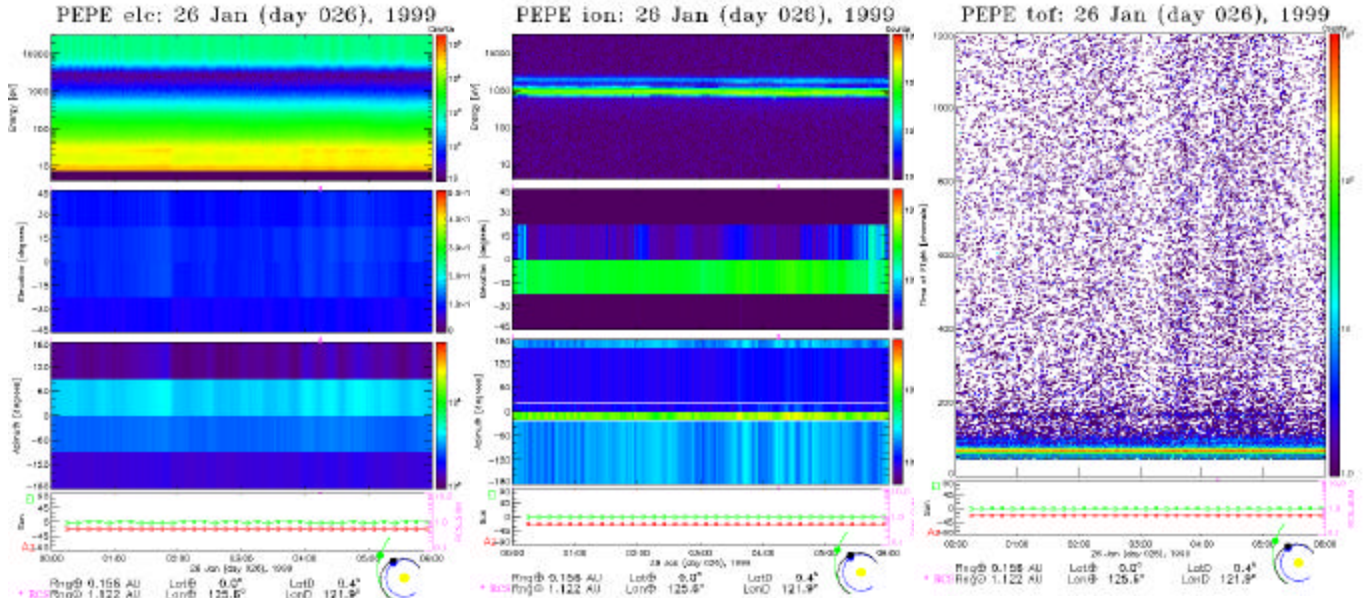


Figure 5. Data from 26 Jan. 1998 0000-0600 UT Illustrating PEPE’s Response to Quiescent Solar-Wind Plasma (Note the lack of interference from solar photons that would appear at all energies near elevation zero in both the ion and electron spectra.)




Figure 6. Data from 1 March 1998 0600-1200 UT Illustrating PEPE’s Response to Disturbed Solar-Wind Conditions (Note the bright vertical bars in the electron data that are caused by attitude-control-thruster firings and the related change in DS1 orientation. The thrusters cause a large cloud of photoelectrons to be created from gas molecules released during the firing.)

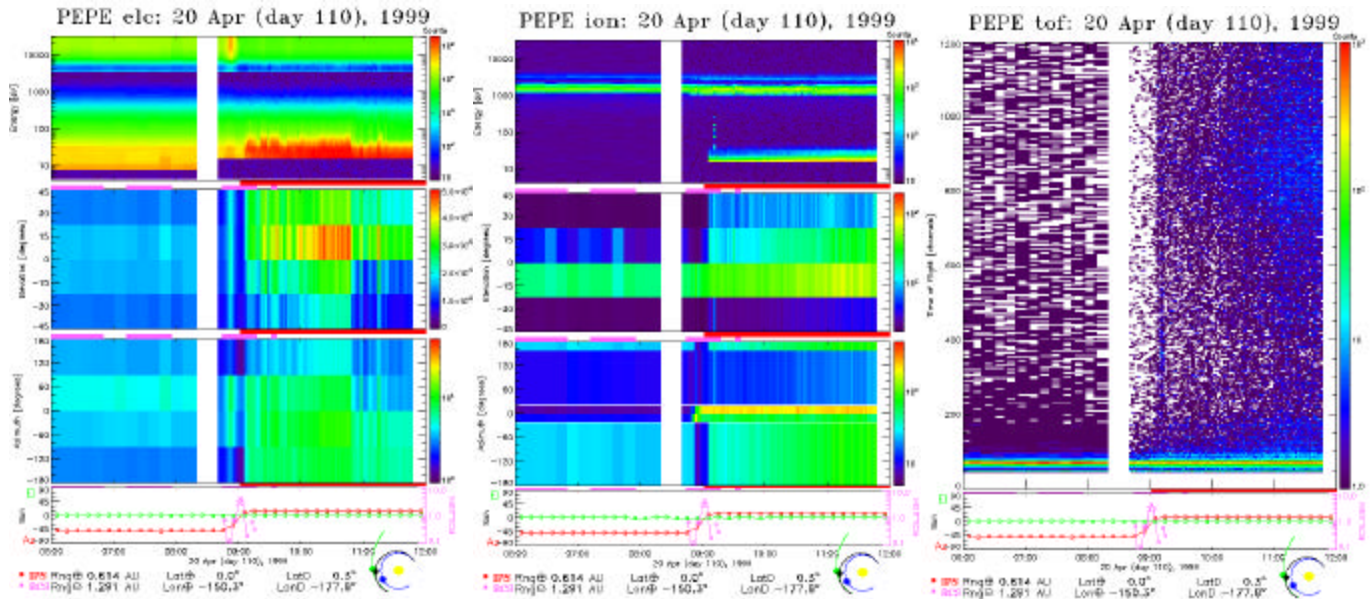


Figure 7. Data from 20 April 1998 0600-1200 UT Showing the Startup (~0905 UT) and Operation of the IPS Xe+ Thruster (The data gap is caused by data-mode changes on DS1. Electron fluxes mask the solar-wind and spacecraft electrons while the Xe+ ions are clearly seen at energies just above the PEPE cutoff of 16 eV.)

data memory is read out while a second spectrum is acquired into a second, identical memory. The two memories (consisting of 1.25 Mbytes each) are operated in ping-pong fashion to ensure the continuous availability of data for compression operations prior to transmission. In addition to forming raw TOF spectra, the system also histograms the TOF data, selects certain pre-programmed, regions of the spectra, and assigns counts within that region to M/Q channels. When operational, the M/Q histogrammer compresses the TOF data efficiently and reduces the amount of data needed to transmit composition information to the ground. Output-data rates can be varied from 1024 bits/s to 25 bits/s. The makeup of this data stream (e.g., the emphasis placed on electron vs. ion data, or TOF vs. singles data) can be changed on command

Over the course of the DS1 mission it has been necessary to reprogram parts of the data-acquisition and processing system to meet the needs, for example, of observations during the flyby of the small asteroid Braille at a velocity of 15 km/s. In that case, a new data sample period four times faster than the designed speed was implemented that could meet the requirements for faster sampling of a reduced region of interest in energy and angular space. Onboard programs used to bring up the PEPE high-voltage system have also been modified several times, as have the data products associated with particular scanning programs.

2.4.5 High-Voltage System—The PEPE electron and ion optics are driven by a system of 8 high-voltage supplies (Figure 3) ranging from programmable but relatively static

voltages of ± 3.6 kV and ± 15 kV to fast-settling 12-bit controlled supplies of ± 2.5 kV and ± 5.0 kV. There are two sets of ± 5.0 kV bi-polar supplies that bias the toroidal deflection electrodes. Performance requirements for the supplies represented compromises in which some characteristics, such as drift and ripple, were allowed to increase slightly in order to simplify circuit design and parts counts. The compromise requirements were based on an assessment of the minimal relaxation that would produce acceptable performance. For example, the deflection and energy-analyzer optics have finite transmission passbands that can be relaxed somewhat in favor of reduced requirements on power-supply stability.

2.4.6 High-Density Packaging Architecture—PEPE contains eight high-voltage power supplies that deliver voltages throughout the instrument (see Figure 2). In order to save the weight and volume associated with bulky high-voltage connectors and cabling, the high-voltage supplies were co-located with the optical elements requiring biasing. Several of the supplies, including the ± 15 kV supplies, were designed as single units that could easily be installed during final instrument buildup. However, other supplies were located deep within the optical system. These were generally hardwired to the optical elements and made use of optical or structural elements for both housing and mounting the supplies. In the CAPS design and the design of many conventional-plasma sensors, the optical elements and electronic components are usually separated into different compartments (if not entirely separate boxes). As is

apparent from Figure 2, the PEPE-optical components are tightly packaged with PEPE-electronic components, particularly in the area of high-voltage supplies and detector electronics. In some cases, the optical paths pass through electronics boards and electronics circuits are placed inside the optical elements, such as the domes of the two electrostatic energy analyzers. This folded-up configuration saved a considerable amount of volume and, therefore, mass, compared to conventional packaging. The particular technologies used to produce compact design include: new methods for fabrication and surface treatment of high-voltage optical electrodes, incorporation of high-voltage signal de-coupling capacitors within detector-anode structures, monolithic microchannel-plate (MCP) holders with integral resistor/capacitor dividers, vapor deposition of thin high-ohmic resistive materials that replace discrete resistor chains, fabrication of suspended sections of multilayer printed-wiring assemblies that allow high-areal throughput of the optical beam through a PWA, high-voltage power supply housings manufactured from metal-coated plastics, and the extensive use of parylene coatings on high-voltage multipliers, which allowed the use of unpotted components. As mentioned above, the packaging density of PEPE is 0.83 g/cm^3 , which is the highest value for plasma spectrometers of which the authors are aware.

2.5 Technology Interdependencies

2.5.1 PEPE Plasma Spectrometer Technology—Because PEPE is a highly capable plasma spectrometer, we have been able to demonstrate the effects that the DS1 ion propulsion system (IPS) and the DS1 spacecraft itself have on local plasma populations and on observations of solar-wind electrons and ions incident on the DS1 spacecraft. PEPE data taken during attitude maneuvers clearly show that there is a strong and irregular photoelectron sheath around the spacecraft. The sheath seems to be effected by the presence of large ($\sim\pm 50 \text{ V}$) differential potentials on the SCARLETT solar arrays and intermittently by the use of the attitude-control thrusters. Data also shows the very noticeable effect that the IPS has on the local-plasma population: low-energy, charge-exchanged xenon ions ejected by the interaction of the primary 1 keV beam with neutral xenon are observed by PEPE at energies up to $\sim 40 \text{ eV}$ (see Figure 7 and Reference [9]). Thermal electrons associated with the Xe^+ beam are accelerated up to $\sim 100 \text{ eV}$ and completely dominate the solar-wind electron flux. Nonetheless, the PEPE observations during IPS thrusting show that observations can still be made of solar-wind ions, though not of electrons.

2.5.2 Data-Acquisition and High-Voltage Systems—The flexibility of the PEPE data-acquisition and operating system technologies has allowed a number of unplanned new modes to be introduced in response to unexpected spacecraft-operational situations or measurement opportunities. A new ROI mode made unexpected use of the

PEPE high-voltage supply technology when it was found that the supplies could be programmed to operate at four times their normal speed. The clock rate of the instrument was increased by a factor of four to allow fast scanning to take place. This new fast mode enabled higher time-resolution-measurements to be made during the asteroid flyby (lack of signal was due to the very weak or non-existent outgassing rate of the object) and will be used again during the planned cometary encounter.

2.6 Test Program

This section summarizes test objectives and success criteria that were used to meet the requirements for instrument validation set forth in Section 2.2. It should be emphasized, however, that the restricted schedule and budget under which PEPE was produced and tested often prevented detailed procedures from being drawn up. Moreover, there was little formal documentation of many test results for the same reason. This section will, therefore, address the test program in a quantitative way wherever possible but will resort to qualitative discussion if necessary.

2.6.1 Ground Test—

2.6.1.1 Miniaturized 3-dimensional Linear Electric Field (LEF3D) Time-of-flight (TOF) Optics—The TOF resistive-cylinder technology was tested by measuring the amount of current drawn with high voltage applied. The resistances of the 3 sections of the cylinder were consistently above 10 Gohm , the value required to meet high-voltage supply-load requirements. The high-voltage stand-off capability of the resistive cylinder was tested repeatedly. Several cylinder combinations were tested at 20% overvoltage ($\pm 18 \text{ kV}$) with varying results. The ultimate performance of the cylinder on the ground was very much effected by the amount of test time in which the system could be pumped to sufficiently high vacuum ($\sim 10^{-8} \text{ Torr}$) for periods of several weeks. In the end we were not able to achieve a stable-applied cylinder voltage above $\sim 8 \text{ kV}$, which was set as the initial on-orbit operating value. (Ironically, just a few weeks after PEPE was delivered, we were able to demonstrate a technology for potting the ceramic cylinders in a way that permitted $\pm 18 \text{ kV}$ to be achieved rather easily.) We planned to operate at the $\pm 8 \text{ kV}$ level initially and then boost the voltage after extensive outgassing was obtained on orbit.

The LEF3D optics and associated high-speed TOF electronics were tested in the LANL ion beam prior to integration with the rest of the PEPE instrument. The tests produced TOF spectra that were difficult to interpret because the beam was not collimated to reduce scattering and neutrals before it entered the TOF section. Ray-tracing simulations of the TOF optics indicated that the goal of $M/\Delta M \sim 20$ could be reached at $\pm 15 \text{ kV}$.

2.6.1.2 High-Speed TOF Electronics—The TOF electronics were integrated with the optics prior to final testing before

delivery. Detailed examination of the TOF spectra indicated that the TOF resolution per channel was 0.75 ns as required. The pulse-pair resolution of 40 ns was also achieved. Logical functions associated with the rejection of non-coincident events and correlations of coincident events were demonstrated. The performance of the built-in-test (BIT) pulser functions was consistent with results taken in the ion beam. Because of the compactness of the final delivery schedule, functional tests associated with measuring circuit dead times under the conditions of randomly-arriving events presented by the ion beam were not carried out.

2.6.1.3 Integrated Ion/Electron Optics—Because of schedule concerns, it was not possible to carry out tests of ion- and electron-optical components at the subsystem level (except for the TOF cylinder discussed above), which is the usual procedure before integration of an instrument. Therefore, the entire optical system was integrated and tested/calibrated at one time in the ion/electron-calibration system at SwRI. The pumpdown period prior to PEPE calibration also served as a single-cycle thermal vacuum test during which the instrument was first cycled hot to +60° C for 48 hours. This satisfied the hot-cycle requirement and also provided a high-vacuum ($\sim 10^{-6}$ Torr) bake-out period as well. A hot start and functional test were performed with the instrument in equilibrium at +45° C. The temperature was then lowered to -35° C and 3 cold starts were performed successfully. Cooling the instrument for these tests had the added advantage of reducing the chamber pressure to 4×10^{-8} Torr, thereby permitting internal-instrument surfaces to outgas more rapidly. The reduced-chamber pressure was an absolute must in order to allow the instrument interior to reach an estimated internal pressure in the 10^{-7} Torr range, where it would be safe to operate high voltage.

Ion beams of several energies were fired at the instrument and successfully recorded as singles events. The ion data indicated that the PEPE energy and angle passbands were in the correct locations and that the PEPE energy-analyzer constant (relating applied voltages to the incident ion energy) was correct. The energy-analyzer constant of 13.07 was close to that determined by ray-tracing (12.8). The angular-deflection constant (see Fact Sheet) could not be verified in the ion beam, although the functionality of the deflector system was verified. Tests of the ability of the multiple-anode system showed that it was operational. An interface problem between the data-acquisition system and the TDC prevented us from obtaining ion TOF-data during an end-to-end test of the optics and electronics. The test was deferred to flight.

A broad energy/angle-electron source was used to stimulate the electron side of the PEPE optics simultaneously with the ion side and this qualitative test was successful. The swept-FOV function was demonstrated for both species, as was the ability to produce scanned-energy spectra. The

ion/electron-beam tests also demonstrated qualitatively that both MCP detectors were operational and operating at roughly nominal efficiencies.

2.6.1.4 Data-Acquisition System—During bench testing, the system successfully acquired all of the data types generated by the internal-pulser system. During vacuum testing, the system successfully acquired and formatted electron- and ion-singles data. The interface problem referred to above was corrected on the bench, but tests of the repaired system were deferred to flight operations.

2.6.1.5 High-Voltage System—The eight power supplies making up the PEPE high-voltage system were first tested on the bench and shown to operate as specified. In particular, the individual voltage levels, voltage waveforms, and transition slew rates were all shown to be within specifications. Once the supplies were integrated with the sensor, it became impossible to test them directly at the output because the electrodes would not tolerate operation at full voltage in air (this is a standard complication). After system assembly, the supplies are monitored on the primary side in PEPE housekeeping data; however, the only way to validate supply operation is through data produced by plasma populations in flight.

Vacuum testing of the high-voltage system was monitored by the ion and electron detectors and by voltage monitors located on the primary side of the supply transformers. Because of problems encountered earlier during testing of the TOF cylinder, tests of the integrated TOF-HV system in vacuum were limited to ± 8 kV. Even at this level, it was noted that the positive TOF-HV monitor tended to sag to slightly lower values. There was no further opportunity to re-test this problem and no fix was attempted on the bench. Because of the relatively high-vacuum pressures during ground testing, the testing of automated-HV turn-on sequences was deferred to flight.

2.6.1.6 High-Density Packaging Architecture—The final ground assembly of the unit proved that the high-density packaging concept worked to the extent that all the parts were inside. Successful high-voltage tests in vacuum (except for the positive TOF-HV sag noted above) proved that the optical and high-voltage systems were packaged correctly. Voltage breakdown during the test would have been indicated by high background rates in the electron and ion MCP detectors; this was not observed. The PEPE instrument density of 0.83 g/cm^3 was determined by dividing the instrument mass by a calculated volume.

2.6.2 Flight Test—Once on-orbit, the PEPE instrument was activated successfully over two DS1 passes on December 8 and 9, 1998. The initial data returned during this period confirmed that end-to-end performance of the PEPE system was nominal (although no detailed-quantitative results could

be obtained immediately). On December 10, the IPS was turned on with PEPE operating. It was immediately clear that the fluxes of Xe⁺ ions and electrons in the lower part of the PEPE energy range were too intense for the PEPE MCP detectors. Subsequent to this, a patch to the PEPE software blocked PEPE energy scans from operating below 16 eV in order to reduce the intensity of the IPS fluxes to tolerable levels on the detectors.

2.6.2.1 Miniaturized 3-dimensional Linear Electric Field (LEF3D) Time-of-flight (TOF) Optics—Once on orbit, the LEF3D optics were used to determine the TOF and M/Q spectra of solar-wind ions. The E/Q ion spectra obtained from the ion singles events also reflect the composition of the solar wind during periods when the solar-wind Mach number is high (>8). These spectra were compared with TOF spectra to demonstrate that the latter were in quantitative agreement. It was also possible to observe xenon and molybdenum ions in the TOF spectra as would be expected during periods of IPS thrusting. Final confirmation of the operation of the TOF system awaits observations of cometary molecules that break up to produce more complex TOF spectra. Resolution of the solar-wind TOF spectra (Figure 8) is consistent with anticipated low-resolution values of $M/\Delta M \sim 5$. High-resolution features have been identified in the spectra

corresponding to $M/\Delta M \sim 17$ to 21. The resolution of the H⁺ peak at a TOF of 70 channels (1 channel = 0.75 ns) is somewhat lower; however, that is expected because of the larger fraction of error that is introduced by the TOF electronics at these short times. Preflight calibration indicated that these were the approximate resolutions; however, because calibration of the TOF section was performed without the collimation provided by the energy-analysis section, it was not possible to determine the exact mass resolution. Figure 8 shows the direct events data (fully resolved and uncompressed) from approximately 3 hours of accumulated time on 2 different days in the solar wind. In Figure 8, the data from both days was summed over energy so that TOF peaks could be more easily picked out. Each peak was analyzed using the energy information also provided in the data. The notation used in the peak labels refers to the ion-charge state before entering PEPE and after the state that exited the foil is shown. The branching ratios for each ion is known and the peaks, as labeled, are consistent with known charge-state branching ratios in thin-carbon foils and the attendant efficiencies. It is somewhat difficult to determine the ultimate-mass resolution for the TOF section because the high-charge states in the solar wind rarely emerge unchanged or more positive from the foils; unless this is the case, ions will not “bounce” and be measured at high resolution.




Figure 8. TOF Spectrum Based on Direct Event TOF Data (This shows solar wind ions at the highest available TOF resolution. Note that the peaks that originated as He⁺ are believed to have derived from the pressurization system of the spacecraft thruster system and not the solar wind.)

2.6.2.2 High-Speed TOF Electronics—The TOF electronics appear to deliver performance consistent with ground measurements. The presence of high-resolution TOF peaks in the spectra confirms the proper functioning of the TOF electronics. The operation of the BIT function has been confirmed in-flight.

2.6.2.3 Integrated Ion/Electron Optics—PEPE data obtained from observations of the solar wind were compared with WIND/SWE instrument data (Figure 9) to obtain calibration factors relating flux intensity and energy/angle response of PEPE to engineering values of density, temperature, and flow velocity. PEPE's values appear consistent within ~10% of measurements made by the WIND instrument. This estimate is based on what appear to be valid correlations between observations made at PEPE and those made at WIND, which is $\sim 10^6$ km distant. The comparison was established by time-shifting the two solar-wind ion spectra until maximum correlation of the density, bulk, and thermal velocities were found. The result is fairly good, as is apparent from Figure 9.

Several other features of the optics have been confirmed as well. PEPE optics and anti-reflective surface treatments were designed to allow the aperture to face directly into the Sun without creating a large flux of internal photoelectrons and resulting background. As seen in Figure 5, a spectrum of the quiet solar wind shows that the PEPE optics are solar blind to the extent that there is no perceptible background above that caused by the spacecraft itself. The cause of the electron background signal at energies >4 keV is not understood. This background is variable, at times disappearing, and does not seem to be related to spacecraft attitude. As is apparent from Figure 5, the background does not interfere with measurements of electrons from either the solar wind or the spacecraft sheath. The data at the bottom of Figure 5 indicate that the spacecraft attitude during which the data were obtained was such that the PEPE aperture was directly viewing the sun at the time. The ability of the deflection optics to keep the solar-wind beam in the instrument FOV despite turns made by the spacecraft is also demonstrated in Figure 6, where the solar wind is quite active but is still tracked by the PEPE deflection system. This shows that the deflection system operates correctly (at least up to solar-wind energies of several keV).

Another important test of the instrument is demonstrated by the data shown in Figure 7, which are taken from a period when the IPS engine was turned on and operated. The ion data clearly show the start-up of the thruster at 0910 UT. The slight disturbance in the electron spectrum at 0850 UT seems to be related to attitude-control thruster firings prior to the main-engine firing at 0910. The electron fluxes intensify just before that time, probably in response to the

plasma neutralizer that emits large numbers of electrons. It is also clear the electron fluxes are highly variable; however, the reason is not understood.

Figure 10 shows the average TOF spectra summed over energy and angle for each of three full days. On the day of year (DOY) 009 of 1999, the IPS was not running. On DOY 083 and DOY 216 of 1999, the IPS ran continuously throughout the day and there was very little, if any, thruster activity. The line spectra clearly show a molybdenum (Mo) peak that appears only on the later day when the IPS had accumulated many hours of firing. The peak at ~ 450 TOF bins is possibly argon; however, possible sources of argon are unknown. This could be a molecular peak with nearly the same total mass. Further investigation is required; however, it is clear that PEPE is capable of providing a very detailed analysis of processes around the spacecraft.

2.6.2.4 Data-Acquisition System—The data system has successfully acquired, compressed, and transmitted more than 75,000 full spectra in the course of operations so far. All of these spectra have been plotted in a summary format shown, for example, in Figures 5, 6, and 7. Examination of the different spectra indicate that the data system acquires and formats the data correctly. High-counting-rate random electron and ion TOF events have been observed throughout this period and have been processed correctly. As noted above, a comparison of PEPE's solar-wind data with that from WIND/SWE indicates that acquisition and processing of PEPE data are being carried out correctly and are free of artifacts that might be introduced by this process. The mass/charge (M/Q) function has not been fully tested because emphasis has been put on the analysis of other data formats.

2.6.2.5 High-Voltage System—The high-voltage system appears to operate correctly. However, a recent anomaly with the operation of the TOF system may be related to sagging of the +HV monitor noted during vacuum testing of the system on the ground. This problem is under investigation. Both ion- and electron-detector backgrounds are very low (<1 count/cm² s [see the spectra in Figures 5 and 6]) and are consistent with the thermionic and cosmic ray backgrounds expected in space. This indicates that very little, if any, noise or ripple is being introduced by the supplies. PEPE's automatic HV turn-on sequence now in use is executed automatically and brings PEPE to full operation within 2 hours (vs. the 4 hours that were originally planned).

2.6.2.6 High-Density Packaging Architecture—The optical and high-voltage systems have performed correctly in flight. This demonstrates that the high-density packaging architecture of PEPE (0.83 g/cm³) is successful.




Figure 9. Spectra Comparing Solar-Wind Data From PEPE and the WIND/SWE Instrument (PEPE data are in red. The two time series have been shifted to obtain the best-time series correlation. Figure courtesy Frank Crary.)

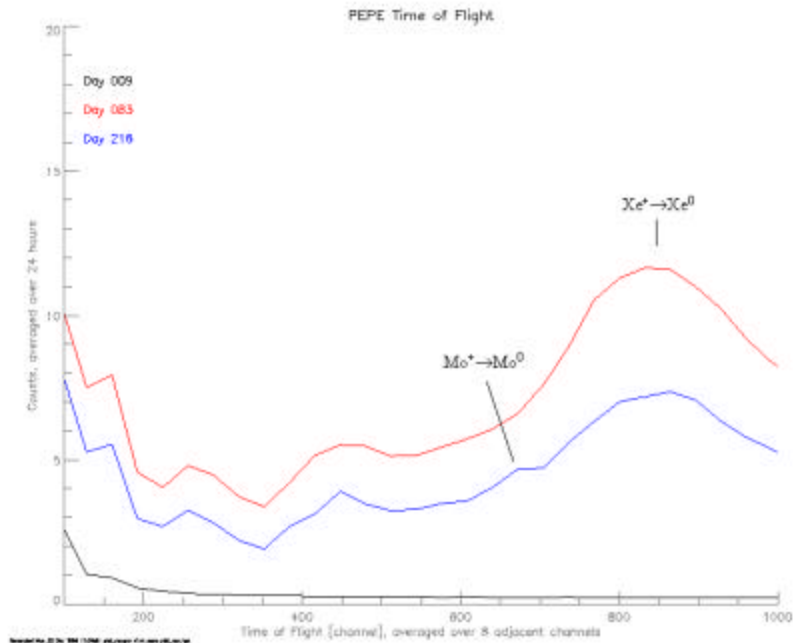


Figure 10. Daily Average TOF Counts for a Day Without IPS Firing (Day 009, 1999), a Day With Continuous IPS Firing Early in the Mission (Day 083, 1999), and a Day Much Later in the Mission With Continuous IPS Firing (Day 216, 1999) (Note the appearance of a molybdenum peak on the last day.)

2.7 Comparison Between Ground and Flight Test

Most of this comparison has been carried out during the discussion of ground and flight tests mentioned above. PEPE has generally proven itself to be a very capable and flexible plasma spectrometer. It is impossible to create any of the space-plasma environmental conditions encountered in space by PEPE on the ground. Ground calibration is restricted to unidirectional, mono-energetic beams of particles that are a poor approximation of the plasmas encountered in space. The single problem that has been encountered in flight (TOF-high-voltage operation) was known from ground testing, but could not be addressed because of the impacted-development schedule. The fact that a successful method was found to address this problem for future applications of the technology by potting the TOF cylinder indicates that the miniaturized-TOF system can be made to work successfully on a future mission.

3.0 TECHNOLOGY VALIDATION SUMMARY

PEPE and its related technologies have been demonstrated to work very well during the flight phase of the mission. All of the six technologies incorporated into PEPE have been validated during the flight phase of the mission with the exception of the operation of the high-voltage system at maximum voltage. However, later tests of an improvement made to the technology on the ground show that the miniaturized TOF system and associated high-voltage subsystem work very well and will be available for future

missions. The PEPE data are of very high quality and are finding their way out to the wider scientific community for further analysis. The ultimate test of PEPE performance must wait for the arrival of the DS1 spacecraft at one or both of the target comets during 2001. That opportunity will allow PEPE to demonstrate the full capability of the six-instrument technologies while contributing to our understanding of cometary physics.

4.0 TECHNOLOGY APPLICATION FOR FUTURE MISSIONS

The PEPE instrument is ideally suited for comprehensive studies of space plasmas on future planetary and magnetospheric missions. The ion/electron optics that perform an analysis of ion and electron directions of arrival and energies have already been incorporated into the Ion Electron Spectrometer instrument scheduled to be flown on the Rosetta Cometary mission in 2007. Individual PEPE technologies, such as the miniaturized high-voltage power supplies, have already served as the basis of improvements in this area. The group at SwRI responsible for the supplies have built a prototype of the MCP supply that weighs 60 grams (vs. the 100 grams for the equivalent PEPE supply). The data acquisition system is being further miniaturized by using more capable gate arrays. In addition, the possibility of custom ASICs designed for this purpose are being investigated for future planetary missions requiring much harder parts.

5.0 LIST OF REFERENCES

- [1] Young, D.T., J. E. Nordholt, J. L. Burch, D. J. McComas, et al., “Plasma Experiment for Planetary Exploration,” to be submitted to *Space Science Reviews*, 2000.
- [2] Nordholt, J. E., D. T. Young, and H. O. Funsten, “Plasma Experiment for Planetary Exploration (PEPE) on DS1,” IEEE.....
- [3] Young, D. T., B. L. Barraclough, J. -J. Berthelier, et al., “Cassini Plasma Spectrometer Investigation,” in *Measurement Techniques in Space Plasmas: Particles*, AGU Monograph Series Vol. 102, R. F. Pfaff, J. E. Borovsky, and D. T. Young, eds., 237–242, 1998.
- [4] Young, D. T., J. -J. Berthelier, M. Blanc, et al., “Casini Plasma Spectrometer Investigation,” accepted for publication, *Space Science Reviews*, 2000.
- [5] McComas, D. J., and J. E. Nordholt, “A new approach to 3-D, high sensitivity, high mass resolution space mass composition measurements,” *Rev. Sci. Instrum.* 61, 3095–3097, 1990.
- [6] McComas, D. J., J. E. Nordholt, J. -J. Berthelier, et al., “The Cassini Ion Mass Spectrometer,” in *Measurement Techniques in Space Plasmas: Particles*, AGU Monograph Series Vol. 102, R. F. Pfaff, J. E. Borovsky, and D. T. Young, eds., 187–193, 1998.
- [7] Nordholt, J. E., J. -J. Berthelier, D. M. Burr, et al., “The Cassini Ion Mass Spectrometer: Performance metrics and techniques,” in *Measurement Techniques in Space Plasmas: Particles*, AGU Monograph Series Vol. 102, R. F. Pfaff, J. E. Borovsky, and D. T. Young, eds., 209–214, 1998.
- [8] Lin, R. P., K. A. Anderson, S. Ashford, et al., “A three-dimensional plasma and energetic particle investigation for the WIND spacecraft,” in *The Global Geospace Mission*, C. T. Russell, ed. 125–153, 1995.
- [9] Wang, J., D. Brinza, R. Goldstein, et al., “Deep Space One investigations of ion propulsion interactions: Overview and initial results,” AIAA paper 99-2971 presented at AIAA/ASME/SAE/ASEE Joint Propulsion Conference, 20–24 June 1999, Los Angeles, CA.


Appendix A: PEPE's Telemetry Channels

W	HKID	New Name	Original Name	Chan ID	Type	Bit pos	Bit len	Calibration	State Title	Units
0	0		CCSDS_HEADER	n/a	n/a	0	96	n/a	n/a	n/a
6	0		mon_mcp_ion	G-0301	unsigned	96	16	u16 * -0.97568-12.3	n/a	volts
7	0	mon_mcp_elc (spelling)	mon_mcp_elc	G-0302	unsigned	112	16	u16 *1.067556-11.9	n/a	volts
8	0		mon_tof_p15k	G-0303	unsigned	128	16	u16 * 3.991378-54.6	n/a	volts
9	0		mon_tof_n15k	G-0304	unsigned	144	16	u16 * -3.98304+37.8	n/a	volts
10	0		mon_sys_p5	G-0305	unsigned	160	16	u16 * 0.001280	n/a	volts
11	0		mon_sys_n5	G-0306	unsigned	176	16	(-1) * u16 * 0.001280	n/a	volts
12	0		mon_sys_p12	G-0307	unsigned	192	16	u16 * 0.003290	n/a	volts
13	0		mon_calor_1	G-0308	unsigned	208	16	u16 * 0.001221	n/a	volts
14	0		mon_calor_2	G-0309	unsigned	224	16	u16 * 0.001221	n/a	volts
15	0		mon_calor_3	G-0310	unsigned	240	16	u16 * 0.001221	n/a	volts
16	0	(Note type change)	eha_generr	G-0311	digital	256	16	n/a	n/a	n/a
17	0	(Note type change)	eha_volterr	G-0312	digital	272	16	n/a	n/a	n/a
18	0	(Note type change)	mon_esa_elc	G-0313	digital	288	16	Note Children sheet	n/a	n/a
19	0	(Note type change)	mon_esa_ion	G-0314	digital	304	16	Note Children sheet	n/a	n/a
20	0		mon_def_pos	G-0315	unsigned	320	16	(u16-2048) * 2.4414	n/a	volts
21	0		mon_def_neg	G-0316	unsigned	336	16	(u16-2048) * 2.4414	n/a	volts
22	0		spare	G-0317	digital	352	1	n/a	n/a	n/a
22	0		ps_fix_stat	G-0318	status	353	1	n/a	OFF ON	n/a
22	0		ps_swp_stat	G-0319	status	354	1	n/a	OFF ON	n/a
22	0		ps_mcp_elc	G-0320	status	355	1	n/a	DISABLED ENABLED	n/a
22	0		ps_mcp_ion	G-0321	status	356	1	n/a	DISABLED ENABLED	n/a
22	0		ps_tof_p15k	G-0322	status	357	1	n/a	DISABLED ENABLED	n/a
22	0		ps_tof_n15k	G-0323	status	358	1	n/a	DISABLED ENABLED	n/a
22	0		ps_bulk_en	G-0324	status	359	1	n/a	DISABLED ENABLED	n/a
22	0		ps_SafeArm_c	G-0325	status	360	4	n/a	SAFED DIV16 ARMED	n/a
22	0		ps_SafeArm_s	G-0326	status	364	4	n/a	SAFED DIV16 ARMED	n/a
23	0	ps1_ps4_regs	ps_status_0	G-0327	digital	368	16	n/a	n/a	n/a
24	0	ps7_abdv_reg	ps_status_1	G-0328	digital	384	16	n/a	n/a	n/a
25	0		if_rti_count	G-0329	unsigned	400	16	n/a	n/a	n/a
26	0		if_sa_error	G-0330	digital	416	16	n/a	n/a	n/a
27	0		if_time_cnt	G-0331	unsigned	432	16	n/a	n/a	n/a
28	0		if_heart_cnt	G-0332	unsigned	448	16	n/a	n/a	n/a
29	0	sc0_ctrl_reg	if_status_0	G-0333	digital	464	16	n/a	n/a	n/a
30	0	sc1_obst_reg	if_status_1	G-0334	digital	480	16	n/a	n/a	n/a
31	0	SPARE	spare	G-0335	digital	496	1	n/a	n/a	n/a
31	0	SPARE	imp_ram_err	G-0336	status	497	1	n/a	OK ERROR	n/a
31	0	SPARE	imp_rom_err	G-0337	status	498	1	n/a	OK ERROR	n/a
31	0	SPARE	imp_shrd_err	G-0338	status	499	1	n/a	OK ERROR	n/a
31	0	SPARE	imp_watchdog	G-0339	status	500	2	n/a	DISABLED ENABLED	n/a
31	0	SPARE	spare	G-0340	digital	502	2	n/a	n/a	n/a
31	0	SPARE	spare	G-0341	digital	504	8	n/a	n/a	n/a
0	1		CCSDS_HEADER	n/a	n/a	0	96	n/a	n/a	n/a
6	1	ps2_swa_reg	imp_status_0	G-0342	digital	96	16	n/a	n/a	n/a
7	1	ps3_swpb_reg	imp_status_1	G-0343	digital	112	16	n/a	n/a	n/a
8	1	SPARE	spare	G-0344	digital	128	1	n/a	n/a	n/a
8	1	SPARE	dpu_ram_err	G-0345	status	129	1	n/a	OK ERROR	n/a


Appendix B: Date of PEPE's Turn-on and Provisional Times-of-Data-Capture List

Note: PEPE operation times are approximate to nearest hour. Intervals lasting less than 24 hours without operations are not recorded.

Start Time (UT)	Date	Stop Time (UT)	Date	Comments
0140	08 Dec 98	0110	10 Dec 98	PEPE Checkout
1900	06 Jan 99	1700	17 Jan 99	
0600	21 Jan 99	0800	01 Feb 99	
2300	02 Feb 99	0300	10 Feb 99	
0000	13 Feb 99	0100	17 Feb 99	
0000	18 Feb 99	1900	23 Feb 99	
0100	25 Feb 99	0100	16 Mar 99	
1200	23 Mar 99	2000	03 May 99	
0300	26 May 99	1900	07 Jun 99	
2300	11 Jun 99	2400	11 Jul 99	
0000	13 Jul 99	1800	05 Aug 99	Turned off due to DS1 anomaly



LOCKHEED MARTIN



Power Actuation and Switching Module DS1 Technology Validation Report

Abbas Salim
*PASM Program Manager
Lockheed Martin Space Systems
Company
L2-30, Bldg. 551
1272 Borregas Avenue
Sunnyvale, Ca. 94089
408-742-9568
abbas.a.salim@lmco.com*




Table of Contents

<u>Section</u>	<u>Page</u>
Extended Abstract	iv
Power Actuation and Switching Module Design and Development Flight Validation Report	1
1.0 Introduction	1
2.0 Technology Description	1
2.1 What It Is; What It Is Supposed To Do	1
2.2 Key Technology Validation Objectives at Launch.....	1
2.3 Expected Performance Envelope.....	2
2.4 Detailed Description.....	2
2.5 Technology Interdependencies.....	6
2.6 Test Program	6
3.0 Technology Validation Summary	6
4.0 Technology Application For Future Missions.....	8
5.0 Acknowledgments	8
6.0 List of References	8
Appendix A. DS1 Technology Validation Telemetry Channels	9
Appendix B. DS1 Technology Validation Power on/off Times	9

Figures

<u>Figure</u>	<u>Page</u>
Figure 1. PASM Module (1.525 × 1.525 × 0.250 in.).....	1
Figure 2. PASM Switch Configuration	2
Figure 3. Switch Current vs. Time	2
Figure 4. Detailed Functional Block Diagram (One of Four PASM Switches)	3
Figure 5. Ceramic Substrate Milling.....	4
Figure 6. Substrate Masking and Metallization.....	4
Figure 7. Populating and Bonding Parts.....	4
Figure 8. HDI Laminating and Etching.....	5
Figure 9. Attaching Surface-Mounted Parts.....	5
Figure 10. Wire-Bonding I/O Leads.....	5
Figure 11. PASM DS1 Flight Configuration.....	6
Figure 12. PASM Flight Performance (Switched Current vs. Time)	7
Figure 13. PASM Flight Performance (Switched Voltage vs. Time).....	7
Figure 14. Future HDI Technology Product Road Map.....	8

Tables

<u>Table</u>	<u>Page</u>
Table 1. PASM Specifications	3
Table 2. Transient Thermal Analysis	6

EXTENDED ABSTRACT

In a unique JPL, Lockheed Martin, and Boeing government/industry partnership a “state of the art” power actuation and switching module (PASM) has been developed, designed, and fabricated for flight qualification on NASA’s Deep Space 1 (DS1) mission in the third quarter of 1998. The features associated with the development of the PASM combine NASA/JPL’s desire to advance the art of power electronics packaging, Lockheed Martin’s proprietary high-density interconnect (HDI) technology, and Boeing Company’s expertise in the application specific integrated circuit (ASIC) design and layout. The PASM development project was organized under JPL’s New Millennium Program (NMP) Microelectronics Integrated Product Development Team (IPDT) and was cost-shared by both the government and industry. The industry assumed the cost of developing the product, and the government paid for its fabrication and test.

The PASM is a quad-switch device. Each of its four stand-alone switches provides the capability to switch power, to isolate faults, and to limit in-rush and fault currents, and supplies voltage and current telemetry. Additionally, it offers the capability for trip time control, di/dt and dv/dt control, and remote on/off control. Each switch can switch anywhere from 3 to 40 V at 3 A maximum and, as a result, can be used in switching the primary as well as the secondary side (conditioned) power. The use of HDI technology for packaging and ASICs for switch control electronics gives PASM a 4 to 1 weight, volume, and footprint advantage over existing hybrid products. It is the advanced packaging technology and utilization of ASICs that makes the PASM unique. It retains, with certain enhancements, all the electrical functions offered in a single-switch hybrid module.

Both HDI and mixed-signal ASIC technologies are rapidly maturing due to their applications in other power and non-power products as well as NASA’s and the Air Force’s commitment to continue to promote and enhance these

technologies. These are strong product risk-mitigation steps that not only have helped to make PASM a successful product, but also have resulted in the successful development of credit card size dc-dc converters at Lockheed Martin.

The key purpose of the validation program was to validate the design and production processes for mixed signal ASICs and the HDI packaging technique and materials by exercising the electrical functions of the switches in the space environment. The validation program included flying two PASM modules as a Category 3 experiment on DS1. The test program included switching 5-V power to a 1-A resistive load through each of the eight switches (four per module). The switches were also operated in parallel (two at a time) to switch 5-V power to the same 1-A resistive load. Certain electrical design flaws in the switch control ASICs prevented them from operating completely. As a result, the in-rush and fault isolation features of the PASM switches were not tested.


The PASM switches were successfully exercised several times during the mission and showed no performance degradation or inability to function.

NASA/JPL has recently awarded a second contract to the Boeing Company to produce a second-generation ASIC to correct the previous design flaws and simplify the design. These second-generation ASICs will be used in PASM modules being procured by JPL for the X2000 program.

The PASM, as well as the technologies used in building it, have succeeded to a large extent in satisfying NASA’s goal to miniaturize power electronics and provide wide-ranging applicability to future NASA science missions as well as other LEO and GEO spacecraft. Additional HDI products in development at Lockheed Martin include a second-generation PASM, dc-dc converters, shunt regulator modules, and lithium-ion battery chargers using PASM technologies. Many of these modules are slated to be delivered to JPL for NASA’s X2000 programs in the near future.




Power Actuation and Switching Module (PASM) Fact Sheet



PASM Module (1.525" x 1.525" x 0.250")

Contact for additional product information:
 Lockheed Martin Commercial Space Systems
 Communications and Power Center
 Newtown, PA 18940
 215-497-1581
<http://www.epc.lmms.lmco.com>



PASM Switch Configuration

What is it?

The power actuation and switching module (PASM) is a quad-switch module. Each of the four switches in the module operates as a circuit breaker by combining both the relay and fusing functions into a single device to safely switch electrical power to the spacecraft loads and to protect and isolate the power source from any load faults.

Why is it exciting technology?


- Lower power hardware manufacturing and test costs
- >4x reduction in weight, volume, and footprint
- Enabling technology for small and large satellites
- Enabling technology for miniaturization of spacecraft power management and distribution modules

When will it be demonstrated?

- The flight demonstration on DS1 was completed in August 1999.
- The technology is being adopted by NASA/JPL X2000 program.
- The same packaging technique as well as the PASM are being used to produce other power modules such as dc-dc converters, shunt regulators, battery chargers, etc.

Who needs it?

- All space missions, as well as commercial consumer electronics and power supplies
- High-density power supplies, instruments, sensors, and micro-power systems or avionics modules.



Switch Current vs. Time

PASM Specifications

Parameter	Specification
Number of switches	four
Switched dc input voltage range (Vin)	3 V to 40 V (28 V nominal)
Housekeeping ± 15 V power (all switches off)	80 mW max.
Housekeeping ± 15 V power (all switches on)	600 mW max.
Rated switch current	3 A max.
Total switch current per module	12 A max. (sum of all four switches)
Switch on resistance (Vin to Vout)	85 mΩ (at 100 °C junction temperature)
Overload trip current	3.5 A ± 7%
Overload trip delay	500 μs min; 500 ms max.
Current limit	4.5 A ± 7%
Turn on time into full rated load	300 μs min; user select max.
Operational temperature range	-40 °C to +100 °C
Storage temperature range	-55 °C to +125 °C

Power Actuation and Switching Module Design and Development Flight Validation Report

Abbas Salim
Lockheed Martin Space Systems Company, Sunnyvale Operations
1272 Borregas Ave. L230, Bldg. 551
Sunnyvale, CA. 94089
(408) 742-9568

1.0 INTRODUCTION

The Power Actuation and Switching Module (PASM) development project came about following the organization of the first two New Millennium Program (NMP) Microelectronics Integrated Product Development Team (IPDT) workshops in early 1996, where a road map for the development of multiple chip modules (MCMs) for power management and distribution (PMAD) electronics was identified. The IPDT industry members proposed various integration and packaging technologies to be developed or advanced jointly with the government, toward the goal of fabricating the MCMs for validation on various deep-space flights. The PMAD MCMs included the dc-dc converters, power regulation and control, and power-switching and distribution electronics. The PASM was conceived to be, and approved for development as the very first product toward fulfilling the NMP roadmap objectives. The Lockheed Martin Corporation's Missile and Space Division and the Boeing Company (the two key members of the microelectronics IPDT) put forth two separate proposals for the development of the PASM. The government (NASA/JPL) opted to combine the best parts of both proposals, thus forming a joint NASA/JPL, Lockheed Martin, and Boeing team for the development of the PASM. Lockheed Martin Corporation was given the overall program responsibility along with the fabrication of the flight-validation modules using their proprietary HDI (high-density interconnect) packaging technology. The Boeing Company was given responsibility for the development and fabrication of the application specific integrated circuits (ASICs) for the PASM. Most design and development effort was funded by the corporations' internal research and development funds, while the government paid for fabricating and testing modules, including the ASICs. The overall design and performance requirements for the PASM were defined by the Lockheed Martin Corporation. The Boeing Company was primarily responsible for the design of the control circuitry for the switch. Work on the production phase started in April 1997 and the flight-validation modules were delivered to JPL in September 1997. A picture of the fully finished de-lidded module is shown in Figure 1.




Figure 1. PASM Module
(1.525 × 1.525 × 0.250 in.)

2.0 TECHNOLOGY DESCRIPTION

2.1 *What It Is; What It Is Supposed To Do*

The heart of each PASM is a switch control ASIC fabricated in Harris Semiconductor's Radiation-Hard Silicon Gate (RSG) process, which is radiation-total-dose tolerant and capable of sustaining high voltages. The PASM is a quad-switch device. Each of the four standalone switches provides the capability to switch power, isolate faults, and limit in-rush and fault currents. Each switch can switch anywhere from 3 to 40 V at 3 A maximum and, as a result, can be used in switching the primary as well as the conditioned (secondary) power. The PASM also includes trip-time control, di/dt control, and provides remote on/off capability and current and voltage telemetry. The use of HDI technology for packaging and of ASICs for switch control electronics gives the PASM a four-to-one weight, volume, and footprint advantage over existing products.

2.2 *Key Technology Validation Objectives at Launch*

The key purpose of the validation program was to validate the design and production processes for mixed signal ASICs and the HDI packaging technique and materials by

exercising the electrical functions of the switches in the space environment. The validation program included flying two PASM modules as a Category 3 experiment on DS1. The test program included switching 5-V power to a 1-A resistive load through each of the eight switches (four per module). The switches were also operated in parallel (two at a time) to switch 5-V power to the same 1-A resistive load. Certain electrical design flaws in the switch control ASICs prevented them from operating completely. As a result, the in-rush and fault isolation features of the PASM switches were not tested.

2.3 Expected Performance Envelope

The PASM switches were successfully exercised several times during the mission and showed no performance degradation or inability to function. Both the load voltage and current telemetry were monitored to assess the performance of each of the eight PASM switches and to ensure that the switch turn on voltage drop is not excessive.

2.4 Detailed Description

Electrical Design—A simplified functional block diagram of the PASM switch configuration is shown in Figure 2.

The module includes four independently configurable switches with independent command, telemetry, and housekeeping power lines. The only common node in the module is the ground. The switches either can be used individually or can be connected in series or in parallel

externally for power switching. Each switch primarily functions as a fault isolation device or a circuit breaker and performs both power switching and fusing functions. It offers current controlled turn on (in-rush current limiting), fault current limiting, trip-time control, and voltage-controlled turn off. These features are graphically depicted in Figure 3.




Figure 2. PASM Switch Configuration

The key performance parameters of the PASM are listed in Table 1. A more detailed functional diagram of a single switch in the PASM is shown in Figure 4, which shows the power-switch FET along with its SCA, various timing capacitors, current sense resistor, output clamp diode and associated input/output functions. Each of the four switches in the module includes a total of nine discrete components interconnected using the HDI technology.

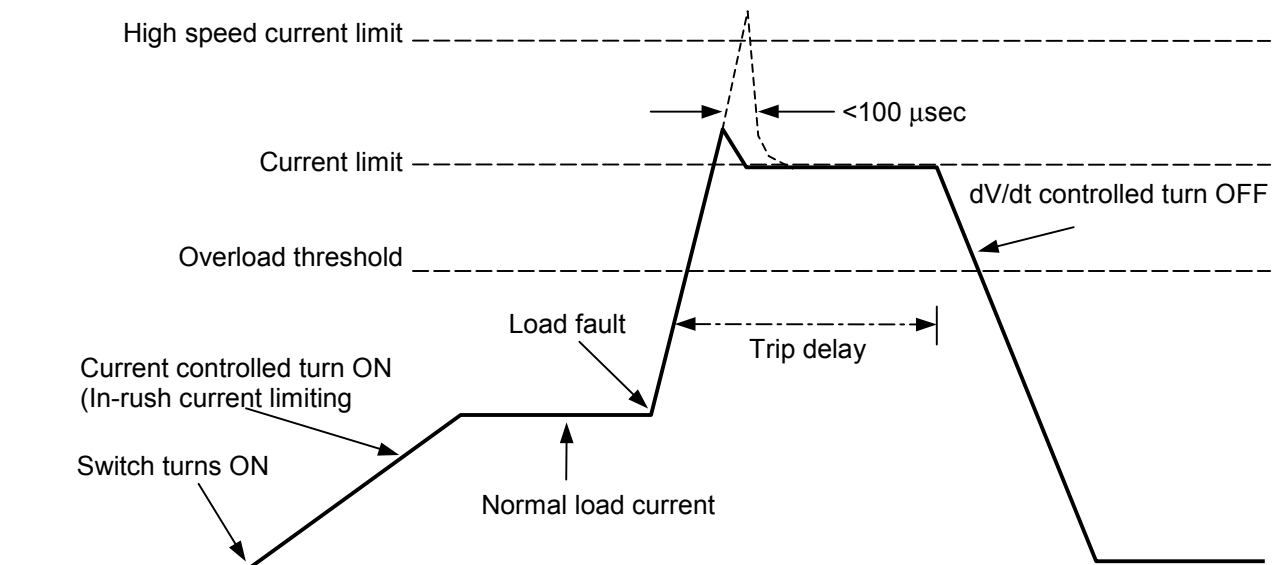


Figure 3. Switch Current vs. Time

Table 1. PASM Specifications

Parameter	Specifications
Number of switches	Four
Switched dc input voltage range (Vin)	3 V to 40 V (28 V nominal)
Housekeeping ±15 V power (all switches off)	80 mW max.
Housekeeping ±15 V power (all switches on)	600 mW max.
Rated switch current	3 A max.
Total switch current per module	12 A max. (sum of all four switches)
Switch on resistance (Vin to Vout)	85 mΩ (at 100 °C junction temperature)
Overload trip current	3.5 A ±7%
Overload trip delay	500 μsec min/500 msec max.
Current limit	4.5 A ±7%
Turn on time into full rated load	300 μsec min; user select max.
Operational temperature range	-40 °C to +100 °C
Storage temperature range	-55 °C to +125 °C

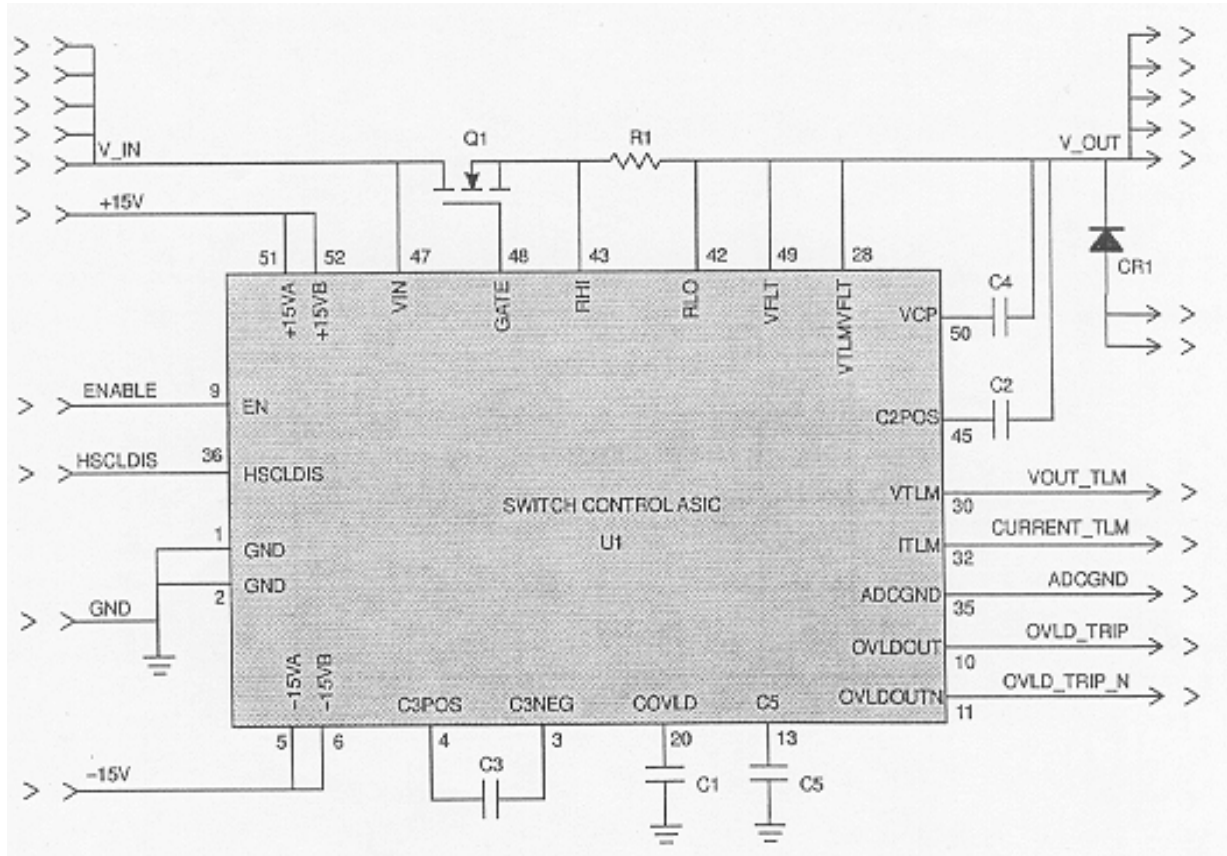


Figure 4. Detailed Functional Block Diagram (One of Four PASM Switches)

ASIC Design—The switch control ASIC (SCA) was custom designed by Boeing and fabricated in Harris Semiconductor’s RSG process. The RSG process is thick-film SOI BICMOS with process enhancements to mitigate the threshold voltage shift post-radiation total dose. The

SCA is 252 by 216 mils with 19 I/O and 6 power/ground leads. Also included were 27 test pads for prototype debug. The power consumption is 150 mW when enabled, and 20 mW when sleeping. Its primary function is to turn on and off a power metallic oxide semiconductor field-effect

transistor (MOSFET) in such a way that load current di/dt is controlled, and the MOSFET is protected from destructive fault conditions. Its secondary functions are to provide load voltage and load-current telemetry, overload status signals, and overload shut down for load fault current. The SCA requires five external capacitors, three of which are selectable for control of turn-off delay, current ramp rate, and overload delay. In Figure 1 there are four SCAs located in the center of the module, underneath the eight capacitors.

PASM Packaging Overview—The PASM uses the Lockheed Martin HDI packaging technology to fabricate a KAPTON™ (polyimide) -based multi-layer interconnect structure. The KAPTON™ structure is laminated one layer at a time to the top surface of the bare die, packaged parts and other active and passive components. Components may be mounted to the topmost layer of the HDI interconnect using standard surface-mount techniques. Components used in HDI are first characterized, which is the physical measurement of components and the mapping of component I/O locations for use during the generation of pads and traces. Pockets to accept the parts are machined into an alumina ceramic substrate (See Figure 5).

Pockets are sized to ensure that the topmost surface of the part is coplanar to the surface of the substrate. The substrate is patterned by sputter deposition, photolithography, and etching to form the required elements prior to component placement (See Figure 6).

The die is attached with thermoplastic resin, thermosetting epoxies (conductive and non-conductive) and various high temperature solders (See Figure 7).




Figure 5. Ceramic Substrate Milling




Figure 6. Substrate Masking and Metallization




Figure 7. Populating and Bonding Parts

The interconnect layer is fabricated upon the populated substrate as follows:

Using a combination of vacuum, heat, and pressure a KAPTON™ film is laminated onto the populated substrate using thermoplastic adhesive. The interconnect bond pads are located using an image processing system. A direct-write laser forms vias through the KAPTON™ to the interconnect bond pads and to I/O pads on the substrate metallization.

The first interconnect layer is formed by sputtering films of titanium, copper, and titanium again. The metals are patterned by exposing a negative photo-resist with a direct-write, computer-controlled laser. The metal is then chemically etched leaving the desired circuit pattern (See Figure 8).

Subsequent layers are formed by laminating additional layers of KAPTON™ onto the substrate using a thermosetting adhesive and repeating the drill, metallization, pattern, and etch process. The module is then populated with surface-mounted components (See Figure 9).




Figure 8. HDI Laminating and Etching




Figure 9. Attaching Surface-Mounted Parts

The completed HDI module is epoxy-bonded into a standard KOVAR™ package and the I/O is wire-bonded (See Figure 10). The package is seam-sealed to complete the module assembly.




Figure 10. Wire-Bonding I/O Leads

Multi-chip module packaging issues—PASM presented three major design issues: current handling capacity of multi-layer thin-film HDI structure, heat dissipation, and CTE (coefficient of thermal expansion) mismatch of large power die-to-substrate. The current handling capacity of HDI in this application was less of a concern than the “ON”

resistance of the switch. The minimization of “ON” resistance was a critical performance characteristic, if the switch was to perform properly as a primary as well as a secondary side or conditioned power switch. The trace size required to carry the specified current with a maximum 10° C rise using 15- μ m copper is approximately 0.3-in. wide. The switch FET die is 0.366 by 0.266 in., which means the trace-width requirement is nearly the length of the FET die. The required trace width indicated that the switch FET must lie parallel to the package interconnect to provide adequate trace access to package input/output pins. The FET die must also be physically next to package output to limit total interconnect length (i.e., lead length, package I/O wire bonds, and HDI interconnect).

Thermal dissipation in the PASM for normal steady state conditions at 3-A maximum-rated current on all four switches was not a significant design driver due to the thermally efficient “chips first” HDI packaging as well as physically large FET die used for switching. The normal operation of the PASM with all four switches in use results in the FET die temperature approximately 4.0° C above ambient.

The real design concern was “failure” operation of the PASM. The PASM is designed to be a smart switch: that is, respond to a current over a specified limit. If a load controlled by the PASM exceeds a preset current for a specified time, the PASM quickly shuts off. The affected load is protected and the failure is prevented from propagating. The PASM is designed to handle large transient currents, shut off, and be available to be commanded back on when required.

The current available from most spacecraft power subsystems in short circuit condition is extremely large. PASM is designed to survive large current transients and turn off without being damaged or damaging the surrounding switches contained in the module. Thermal analysis of the PASM switch components was performed using SINDA 87™ assuming an ambient temperature of 75° C. Electrical power applied to each junction in the switch was modeled as a square wave pulse of 100- μ sec duration. Only the FET had a significant junction temperature increase from the initial 75° C ambient. Results of the transient model are tabulated in Table 2.

The CTE differences between the large FET and the 96% alumina substrate typically used for HDI was evaluated to insure the product would meet mission requirements. The large size of the FET die in the PASM required the use of gold-clad-molybdenum tabs (molytab) or interposers between the base of the FET and the alumina substrate [4]. The FET-die-to-molytab attachment used gold/germanium eutectic; the molytab-to-substrate attachment used Indalloy™ 165 eutectic. The use of the molytab in this

application was dictated by (1) the physical size of the die, (2) the large number of possible thermal and power cycles the module would be exposed to in normal operation, and (3) the inability to use a silver-loaded epoxy.

in-rush and fault isolation features of the PASM switches were not tested.

Table 2. Transient Thermal Analysis*

Component	Power (W)	T _{junction} (C)
Q1 (FET)	3500	96.2
CR1	0	77.0
U1 (ASIC)	0.151	75.0
R1 (Current Sense)	290	75.6
Q31	0.370	75.7**

* One Switch circuit consists of Q,CR,U, and R (Capacitors not shown).
 ** Q31 is an Adjacent FET, Normal Operation (Shown for Comparison).

The design of power electronics is constantly under pressure to reduce size and weight and increase interconnect density to better integrate power products with the end users they serve. The drive to integrate power electronics generally requires technologies that do not easily lend themselves to carrying large currents. The Lockheed Martin HDI technology provides a unique blend of capabilities for power packaging. HDI technology has historically been used for digital or RF applications and has only been applied to power packaging in the last few years. The standard HDI process has been modified to allow use of up to 24 μm (0.001 in.) copper layers for power applications. Processing temperatures have been decreased to allow the use of magnetic as well as packaged parts.

Accommodation of PASM on DS1—A set of PASM s was launched on DS1 in October 1998 for flight-performance verification and validation. The DS1 PASM s are mounted on a printed circuit board (see Figure 11) and housed in a VME cage. The outputs of the modules are connected to dummy loads for on/off characterization of the PASM switches and their performance evaluation during various phases of the mission. The first set of PASM performance test data from DS1 was received in February 1999.

2.5 Technology Interdependencies

PASM s were flown on DS1 as Category 3 electronics experiment and therefore had no direct impact on the performance of other technologies tested on DS1 or on other subsystems of the spacecraft.

2.6 Test Program

The test program included switching 5-V power to a 1-A resistive load through each of the eight switches (four per module). The switches were also operated in parallel (two at a time) to switch 5-V power to the same 1-A resistive load. Certain electrical design flaws in the switch control ASICs prevented them from operating completely. As a result, the




Figure 11. PASM DS1 Flight Configuration

2.7 Comparison Between Ground Test and Flight Test

Figures 12 and 13 show the PASM in-orbit performance over time. Figure 12 shows the load current supplied by each of the eight PASM switches, and Figure 13 shows the voltage at the load supplied by the individual switches. These figures show that over a period of nearly 8 months, all eight switches performed satisfactorily and did not exhibit any degradation in the performance, primarily in terms of the excessive voltage drop in the switch itself. No difference was seen in the PASM ground and in-orbit performance.

3.0 TECHNOLOGY VALIDATION SUMMARY

A state-of-the-art PASM using HDI and mixed signal ASIC technologies has been developed. This program has significantly contributed in validating several production processes which are key to the development and production of future high-density lightweight power electronics. The eventual goal is a self-contained, three-dimensional avionics module for both space and commercial applications.

The PASM performance test data received from the DS1 flight and incorporation of various lessons learned from the design and fabrication phase of this module should help in enhancing the performance of the second generation PASM currently under development.

NASA/JPL has recently awarded a second contract to the Boeing Company to produce a second-generation ASIC to correct the previous design flaws and simplify the design. These second-generation ASICs will be used in PASM s being procured for the X2000 program.




Figure 12. PASM Flight Performance (Switched Current vs. Time)




Figure 13. PASM Flight Performance (Switched Voltage vs. Time)

4.0 TECHNOLOGY APPLICATION FOR FUTURE MISSIONS

The PASM, as well as the technologies used in building the PASM, have succeeded to a large extent in satisfying NASA's goal of miniaturizing power electronics and have added wide-ranging applicability to future NASA science missions as well as other LEO and GEO spacecraft. Lockheed Martin was recently awarded a \$16 million contract to design and build multiple dc-dc converters, shunt regulator modules, and lithium-ion battery chargers using PASM technologies. Supply of a large number of second-generation PASM's is also included in the same contract for NASA's X2000 programs. Figure 14 shows the technology road map for the recently awarded contract and the future product development possibilities.

The PASM design and its technologies are also applicable to consumer electronics—power-switching applications that always emphasize miniaturization and lightweight products.

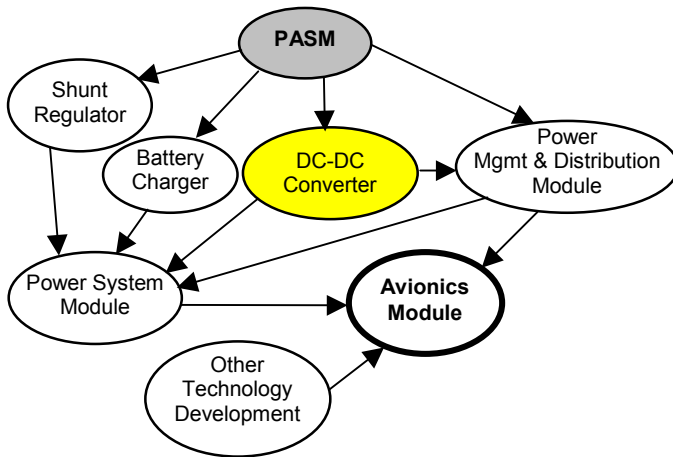


Figure 14. Future HDI Technology Product Road Map

5.0 ACKNOWLEDGMENTS

The work described in this report was carried out at Lockheed Martin Communications and Power Center, Newtown, Pennsylvania; Lockheed Martin Government

Electronics, Morrestown, New Jersey; and at the Boeing Company, Seattle, Washington. The financial support for the design and development of the PASM and its ASICs was provided by Lockheed Martin and the Boeing Company. The program to develop the PASM was organized by JPL for the National Aeronautics and Space Administration, which also paid for its fabrication, test and flight validation on DS1.

Several organizations and individuals have contributed significantly to the successful development of the PASM. They include Dr. Leon Alkalai, Karla Clark, John Treichler and Greg Carr of JPL; Gary Nelson and David Hogue of Boeing Company; and Jim Jud, James Mulvey, Gerhard Franz and Lynn Melino of Lockheed Martin Corporation. Our sincere thanks and recognition are extended to the DS1 Program Office for their willingness to fly the PASM on DS1 as an experiment for flight validation. Chuck Minning of JPL (NMP Microelectronics IPDT Lead) deserves special thanks for his encouragement and support.

6.0 LIST OF REFERENCES

- [1] "Development of a State-of-the-Art Power Actuation and Switching Module," *IEEE International Workshop on Integrated Power Packaging*, Chicago, Ill. Sep. 17–19, 1998.
- [2] "PASM, The Advanced Power Actuation and Switching Module as the Building Block for Space Micropower Systems," *Govt. Microcircuit Applications Conference*, Monterey, CA, Mar. 8–11, 1999.
- [3] "Power Electronics for the Next Century—First Step," *2000 IEEE Aerospace Conference*, Big Sky, Montana, March 18–25, 2000.
- [4] Sergeant, J. and C.A Harper, *Hybrid Microelectronics Handbook*, McGraw-Hill, Inc., 1995.
- [5] Fillion, R., R. Wojnarowski, R. Saia, and D. Kuk, "Demonstration of a Chip Scale Chip-on-Flex Technology," *ICEMCM Proceedings*, 1996.
- [6] *Lockheed Martin Government Electronic Systems Microwave High Density Interconnect Design Guide*, Revision A, October 17, 1995.
- [7] Burdick, W. and R. Fillion, "Extension of the Chip-on-Flex-Technology to Known Good Die," *Microcircuits & Electronic Packaging*, Vol. 19, Number 4, 4th Qtr., 1996.


Appendix A. DS1 Technology Validation Telemetry Channels

Below is a list of all of the telemetry channels that the PASM team collects and uses. (Kirk Fleming, 10/14/99.)

Channel	Mnemonic
P-0315	PASMdataQual
P-0316	PASMdataWord
P-0317	PASM_t_stam;
B-0032	bmPASMgdcddct
B-0033	bmPASMbdcddct
B-0034	bmPASMbdtlct

Appendix B. DS1 Technology Validation Power on/off Times

LPE/PASM initial turn-on was February 25, 1999. The experiment was then conducted weekly from power-off. (Kirk Fleming, 10/29/99)



Remote Agent Experiment DS1 Technology Validation Report

Douglas E. Bernard, Edward B. Gamble, Jr., Nicolas F. Rouquette,
Ben Smith, Yu-Wen Tung
*Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California 91109*

Nicola Muscettola, Gregory A. Dorias, Bob Kanefsky,
James Kurien, William Millar, Pandu Nayak, Kanna Rajan,
Will Taylor
Ames Research Center



<http://rax.arc.nasa.gov>



Table of Contents

<u>Section</u>	<u>Page</u>
Extended Abstract	v
Fact Sheet	vii
1.0 The Remote Agent	1
1.1 Technology Overview	1
1.2 Detailed Validation Objectives	3
1.3 Performance Envelope	4
1.4 Technology Details	4
1.5 Subsystem Interdependencies	10
1.6 Preparing Lisp for Flight	10
2.0 The Remote Agent Experiment	11
2.1 Historical Perspective	11
2.2 Domain Models	12
2.3 Experiment Scenarios	14
2.4 RAX Development	15
2.5 Ground Tests	16
2.6 Ground Tools	18
2.7 Flight Test	20
2.8 Effectiveness of the Development and Test Process	21
2.9 Costing	24
2.10 Lessons Learned	24
2.11 Answers to a Project Manager’s Questions	27
3.0 Future Applications	28
4.0 Acknowledgments	29
5.0 List of References	29

Figures

<u>Figure</u>	<u>Page</u>
Figure 1. Remote Agent Architecture	1
Figure 2. Planner/Scheduler Architecture	5
Figure 3. Temporal Constraints in DDL	5
Figure 4. A Plan Fragment Formed by a DDL Model	6
Figure 5. An Overview of the Remote Agent Executive	7
Figure 6. Multiple Methods in ESL for Achieving Thrust	8
Figure 7. Livingstone Processing Cycle	8
Figure 8. Livingstone Model of the Cassini Main Engine Subsystem	9
Figure 9. Schematic of Livingstone Processing	9
Figure 10. Packetview—Telemetry Packet Display	18
Figure 11. ExecView—Plan Execution Status	18
Figure 12. PS Graph—Planner Progress Display	19
Figure 13. Stanley—Hardware Status Display	19
Figure 14. Timeline Applet	20
Figure 15. Temporal Distribution of Problem Reports	22
Figure 16. Planner PRs by Category	22
Figure 17. Executive PRs by Category	22
Figure 18. MIR PRs by Category	23
Figure 19. RAX Costing	24

Tables

<u>Table</u>	<u>Page</u>
Table 1. Autonomy Levels of RA	3
Table 2. Significant Events for the RAX Project	12
Table 3. Summary of Planner Models for RAX	13
Table 4. DS1 Hardware Modeled as Components in MIR	13
Table 5. DS1 Hardware Modeled as Modules in MIR	14
Table 6. Timelines and Their Respective Tokens by Module (EXEC's perspective)	14
Table 7. Development Testbeds for RAX	16
Table 8. Dates of RAX Readiness on Testbeds	16
Table 9. Number of PRs by Subsystem	22

Appendices

<u>Appendix</u>	<u>Page</u>
Appendix A. Telemetry Channels	31
Appendix B. DS1 Technology Validation Power On Times	32
Appendix C. Acronym Definitions	33

EXTENDED ABSTRACT

Remote Agent (RA) is a model-based, reusable, artificial intelligence (AI) software system that enables goal-based spacecraft commanding and robust fault recovery. RA was flight validated during an experiment onboard Deep Space 1 (DS1) between May 17 and May 21, 1999.

Technology Overview

RA can operate at different levels of autonomy, allowing ground operators to interact with the spacecraft with immediate commands to the flight software, if needed. However, one of the most unique characteristics of RA, and a main difference with traditional spacecraft commanding, is that ground operators can communicate with RA using *goals* (e.g., “During the next week take pictures of the following asteroids and thrust 90% of the time”) rather than with detailed sequences of timed commands. RA determines a plan of action that achieves those goals and carries out that plan by issuing commands to the spacecraft. Actions are represented with tasks that are decomposed on the fly into more detailed tasks and, eventually, into commands to the underlying flight software. When discrepancies are detected between the desired state and the actual state, RA detects, interprets, and responds to the anomaly in real time. More serious anomalies can be addressed with longer response times by generating a new plan of action while the spacecraft is kept idle in a safe configuration. When the new plan is generated, the spacecraft is taken out of the safe configuration and execution resumes normally.

RA differentiates itself from traditional flight software because it is *model-based*. In traditional software programs and expert systems, the programmer decides what the result of a program should be and writes down instructions or rules that attempt to achieve those results. The computer simply executes the instructions or fires the rules with no knowledge of what the intended result was or how it is achieving it. In the RA system, however, each component operates on *models*, general descriptions of the behavior and structure of the spacecraft it is controlling. Each RA component solves problems by accepting goals and using appropriate reasoning algorithms on its models to assemble a solution that achieves the goals. The reasoning algorithms are general-purpose and remain unchanged across different deployments of RA. For different applications, the parts that change are the models and possibly the problem-solving control knowledge needed by some RA modules to tune performance.

Remote Agent Component Technologies

Remote Agent integrates three separate technologies: an onboard Planner/Scheduler (PS), Smart Executive (EXEC), a robust plan-execution system, and the Mode Identification and Recovery (MIR) system for model-based fault diagnosis

and recovery. These component technologies are described briefly below.

PS—PS generates the plans that RA uses to control the spacecraft. Given the initial spacecraft state and a set of goals, PS generates a set of synchronized high-level tasks that, once executed, will achieve the goals. PS consists of a heuristic chronological-backtracking search engine operating over a constraint-based temporal database. PS begins with an incomplete plan and expands it into a complete plan by posting additional constraints in the database. These constraints originate either from the ground, which imposes them directly on the goals, or from constraint templates (e.g., the camera must be pointed at an asteroid to take a picture of it) stored in a model of the spacecraft. PS queries domain-specific planning experts (specialized software modules such as Deep Space 1’s navigation system) to access information that is not in its model.

EXEC—EXEC is a reactive, goal-achieving control system that is responsible for:

- Requesting and executing plans from the planner.
- Requesting/executing failure recoveries from MIR.
- Executing goals and commands from human operators.
- Managing system resources.
- Configuring system devices.
- System-level fault protection.
- Achieving and maintaining safe-modes as necessary.

EXEC is goal-oriented rather than command-oriented. A goal is defined as a system state being controlled that must be maintained for a specified length of time. As a simple example, consider the goal: keep device A on from time X to time Y. If EXEC were to detect that device A is off during that period, it would perform all the commands necessary to turn it back on. EXEC controls multiple processes in order to coordinate the simultaneous execution of multiple goals that are often interdependent. In order to execute each goal, EXEC uses a model-based approach to create a complex command procedure designed to robustly achieve the goal.

MIR—The MIR inference engine provides mode identification (diagnosis) and mode reconfiguration (recovery) functionality. To track the state of each component (called a mode) in the spacecraft, MIR eavesdrops on commands that are sent to the spacecraft hardware by EXEC. As each command is executed, MIR receives observations from spacecraft’s sensors, which are then abstracted by monitors in the spacecraft’s control software. MIR combines these commands and observations with declarative models of the spacecraft components to determine the current state of the system and to report it to EXEC. If failures occur, MIR uses

the same model to find a repair or workaround that allows the plan to continue execution.

The key idea underlying model-based diagnosis is that a combination of component modes is a possible description of the current overall state of the spacecraft only if the set of models associated with these modes is consistent with the observed sensor values. This method does not require that all aspects of the spacecraft state be directly observable, providing an elegant solution to the problem of limited observability.

Risks

RA is flight software and as such poses the same kind of risks posed by conventional flight software. The autonomous behavior implemented by RA is not qualitatively different from that displayed by conventional fault protection or attitude control. In all cases, the spacecraft is commanded on the basis of current state information rather than by direct operator commands. The behavior of RA can be predicted, within an envelope, just as the behavior of fault protection or attitude control can be predicted within certain bounds. Confidence in the RA's responses can be obtained through testing, just as confidence in fault protection or attitude control is obtained now.

A risk addressed by the experiment concerns the integration and testing of the technology. RA in a novel integration of three technologies; the application of these integrated technologies to spacecraft is also new. For this reason, there was no prior experience on development and validation methodologies for such a system. Another risk had to do with the integration of the AI technologies of RA, based on general-purpose search algorithms, together with real-time control software on a flight processor.

Validation Objectives

The first validation objective was to demonstrate RA's ability to autonomously operate a spacecraft with communication from ground limited to few high-level goals. This translated into specific objectives for PS, EXEC, and MIR. The second validation objective was to show that RA could be commanded with different levels of autonomy. This meant supporting all of the possible operation modes: using EXEC to run a traditional sequence of commands,

preparing a plan on the ground and uplinking it to the spacecraft for execution, and providing closed-loop planning and execution onboard the spacecraft. The final validation objective was the first formulation of a development and testing plan for an autonomous flight software system.

Test Program and Results

The Remote Agent Experiment (RAX) consisted of using the RA technology to operate the DS1 spacecraft for several days. A series of operations scenario based on DS1 active cruise mode was developed. In these scenarios, RAX commanded a subset of the spacecraft subsystems: Ion Propulsion System (IPS), Miniature Integrated Camera and Spectrometer (MICAS), Autonomous Navigation (NAV), Attitude Control System (ACS), and a series of power switches.

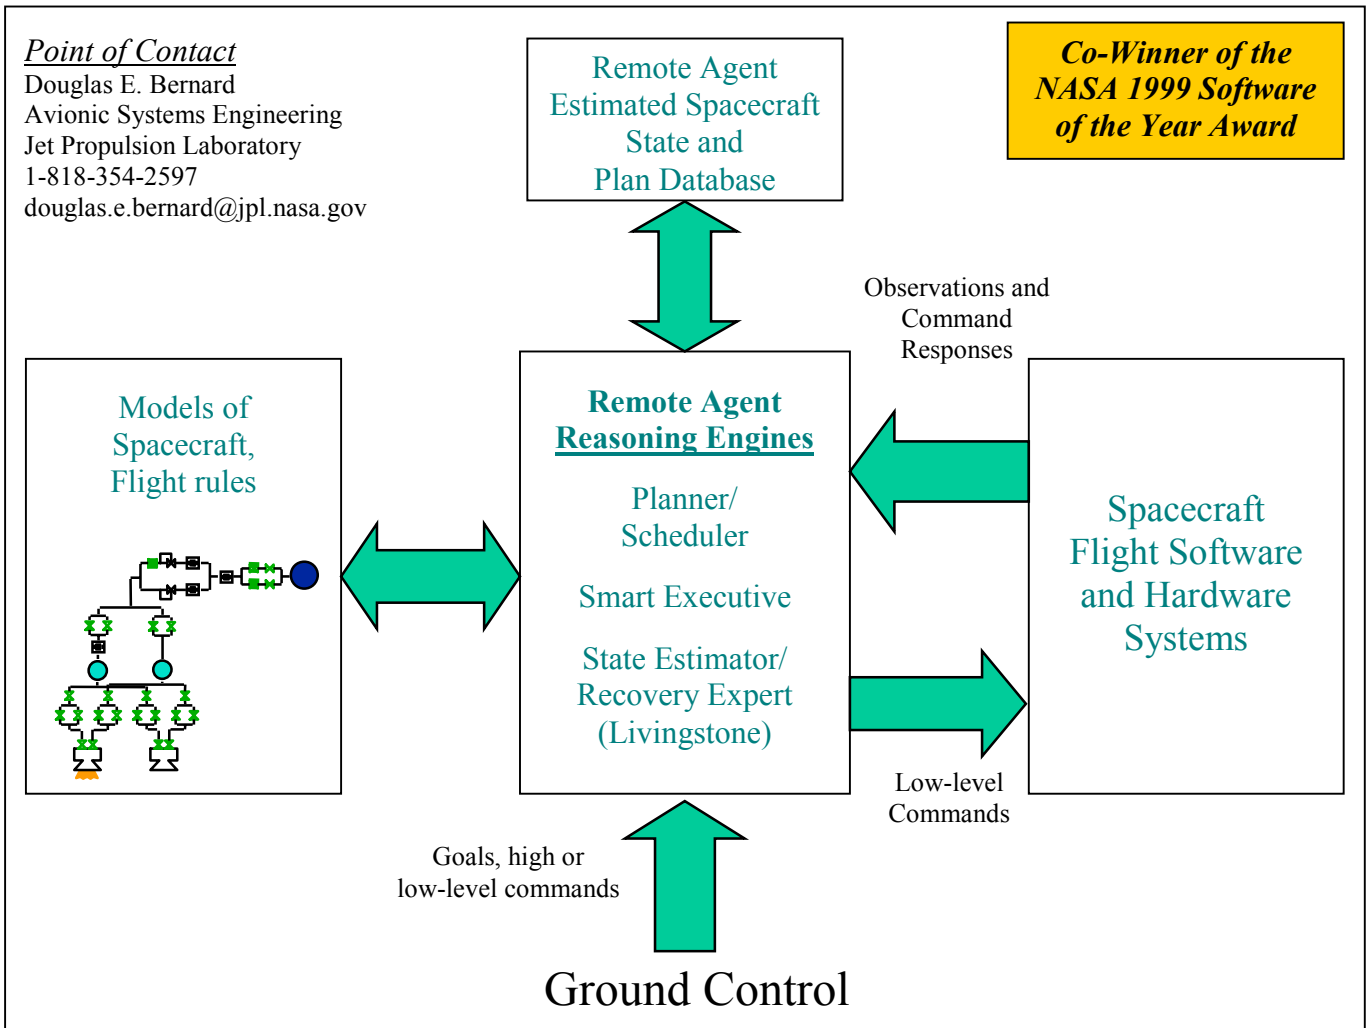
The main scenario goals were to execute an IPS thrust arc, acquire optical navigation images as requested by the autonomous navigator, and respond to several simulated faults. The faults included minor ones that could be responded to without disrupting the current plan and more serious ones that required generating a new plan to achieve the remaining goals. A continuous integration approach was adopted in which new features or bug fixes were integrated in new releases only after the integrated system could successfully run the reference scenarios on all available testbeds. An extensive formal-testing program was conducted, separate from the software development process. Testing was distributed on several different platforms of different speeds, level of fidelity, and availability to the RA team. Test cases were targeted to the most available testbed that could validate them with the reasonable expectation that test results would hold on higher fidelity testbeds.

In spite of a couple of bugs that occurred during the flight experiment, RA successfully demonstrated 100% of its flight validation objectives.

Applicability to future NASA missions

The Remote Agent technology is applicable to any future NASA mission that desires or requires autonomous operations. The RA reasoning engines can be used as-is on future missions. New domain models would be required for each mission.

FACT SHEET



Validation Objectives

- ✓ •Initiate and generate flexible plans on-board
- ✓ •Reject low-priority, unachievable goals
- ✓ •Execute plans generated both onboard and from ground
- ✓ •Confirm execution of commands
- ✓ •Demonstrate model-based failure detection and recovery
- ✓ •Maintain required spacecraft states in the face of failures
- ✓ •Re-plan following a failure
- ✓ •Generate back-to-back plans
- ✓ •Modify mission goals from ground
- ✓ •Execute low-level commands from ground
- ✓ •Update estimated spacecraft-state database from ground

Capabilities

- Robust goal-based commanding
 - Planner expands high-level goals into flexible plans
 - Smart Executive decomposes plans into low-level spacecraft commands and monitors that the states commanded to are achieved and maintained
- Fail-operational model-based fault recovery
 - Livingstone identifies faults and suggests recoveries that the Smart Executive uses to continue plan execution
 - If necessary, Executive requests the Planner/Scheduler to generate a new plan in light of failure

Applicability to Future Missions

Remote Agent technologies are generally applicable to missions that benefit from highly autonomous operation and are currently being applied to prototypes of future NASA missions including a space-based interferometer and an in-situ propellant production plant.

Remote Agent Experiment DS1 Technology Validation Report

*Douglas E. Bernard, Edward B. Gamble, Jr., Nicolas F. Rouquette, Ben Smith, and Yu-Wen Tung
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California*

*Nicola Muscettola, Gregory A. Dorais, Bob Kanefsky, James Kurien, William Millar, Pandu Nayak, Kanna Rajan, and Will Taylor
NASA Ames Research Center*

1.0 THE REMOTE AGENT

Remote Agent (RA) is a model-based, reusable, artificial intelligence (AI) software system that enables goal-based spacecraft commanding, and robust fault recovery. This report describes the RA technology, its development and test history, and the DS1 flight experiment in which RA was validated. Whenever feasible, this report attempts to give guidance on how RA can be fruitfully employed in future science missions. Also highlighted are further technology developments and operational applications the team is currently pursuing.

1.1 Technology Overview

RA integrates three separate Artificial Intelligence technologies: automated planning and scheduling, robust multi-threaded execution, and model-based fault diagnosis and recovery.

1.1.1 Remote Agent Architecture—The RA architecture and its relation to flight software are shown in Figure 1. Viewed as a black-box, RA issues commands to real-time execution flight software (FSW) to modify spacecraft state and receives data from the spacecraft through a set of monitors that filter and discretize sensor values. The RA itself is comprised of three main components: a Planner/Scheduler (PS), a Smart Executive (EXEC), and Livingstone, a Mode Identification and Reconfiguration (MIR) system. An additional component, strictly related with PS, is the Mission Manager (MM). In addition, the RA team provided a clean interface to the rest of the FSW via the Remote Agent Experiment Manager (RAXM), which mediated all communication between RA and FSW and was included from the outset in the FSW design. RAXM provided a messaging conduit between RA and the rest of FSW, including interfaces to the planning experts, as well as to the monitors and the real time sequencer. This mechanism allowed RA to be cleanly bundled on top of the FSW much later in flight and also allowed a clear methodology for testing and validating the RA software on the ground.

The main functionalities provided by RA, how each individual RA component participates in the overall picture, and concrete examples of commanding and operations relative to DS1 are described below.

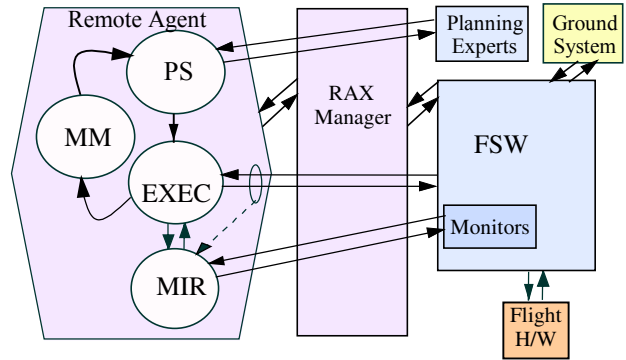


Figure 1. Remote Agent Architecture

RA can operate at different levels of autonomy, allowing ground operators to interact with the spacecraft with immediate commands to the FSW, if needed. However, what makes RA unique is that ground operators can skip formulating detailed timed-command sequences and communicate with RA at the *goal level*. Goals are stored in MM in a *mission profile* covering an extended period.

In principle, a mission profile could contain all goals for a mission, requiring no further uplink from ground. More realistically, mission operations will want to change goals (e.g., scheduled DSN communications can be modified on a week-by-week basis). This is easily done by uplinking commands to edit the mission profile. Goals typically contain few details of how they should be done. For example, the only DS1 Remote Agent Experiment mission profile goals were “Perform AutoNAV orbit determination (OD) activities for 1 hour every day,” and “Thrust the IPS engine for at most 12 hours.”

To translate high-level goals into a stream of commands to flight software, RA follows a two-step process. In the first step, MM selects goals for the next commanding horizon (typically covering several days) and sends them to PS. PS uses its model of the spacecraft to determine which detailed tasks should be selected and scheduled to achieve the goals. For example, in order to perform an OD, PS determines from the model that pictures of beacon asteroids need to be taken. In order to select these asteroids, the model instructs PS to interrogate the AutoNAV software as a *planning expert*. In general, PS will rely on several specialized services provided by software modules external to RA. In DS1, both AutoNAV and ACS provided information to PS

that was incorporated into plans. Going back to our example, observing an asteroid translates, according to the PS model, into taking a series of images of it with the Miniature Integrated Camera and Spectrometer (MICAS).

Therefore, PS schedules a “MICAS take Optical Navigation (OPNAV) module subsystem FSW images” task. Moreover, the model instructs PS that while images of an asteroid are being recorded, the attitude of the spacecraft must be compatible with the MICAS camera pointing at it. If this is not the case, the PS model instructs PS to schedule an appropriate turn, changing the attitude from the previous one to the desired one.

The brief example above points out another fundamental characteristic of all RA components: their fundamental reliance on explicit, declarative *models* of the spacecraft.

Although the level of detail varies between the different components, RA models are fairly abstract and focus on system level interactions—not detailed individual subsystems’ or components’ operation.

This approach has two advantages. First, this provides a method to capture system-level knowledge in a form that can directly command a spacecraft—no costly, error-prone translation into flight software is needed. At best, system requirements are translated into flight rules to check command sequence validity, not generate them.

Secondly, the more abstract models employed are less susceptible to changes when a detailed understanding of the behavior of each subsystem is gained during spacecraft development. Although they need to be adjusted to the new finding, abstract models usually remain structurally unchanged and, therefore, remain the synthesis procedures that RA components use to generate command loads.

Once PS has generated a plan for the next commanding horizon, EXEC receives it and incorporates it into the queues of tasks that it is currently executing. Tasks generated by PS tend to be fairly abstract. EXEC’s responsibility is to synchronize the parallel execution of the plan’s tasks according to the specifications contained in the plan and to further decompose each task, one at a time, into more detailed steps. This task decomposition eventually results in individual commands being sent, one at a time, to FSW. For example, the abstract task “MICAS take OPNAV images” is decomposed into commanding MICAS to take a number of snapshots while checking that MICAS is kept “ON” during the entire process.

Besides its goal-directed commanding and model-centered approaches, RA puts particular emphasis on robustness of execution and flexibility of response to faults. The mode identification (MI) component of MIR observes EXEC

issuing commands, receives sensor observations from monitors, and uses model-based inference to deduce the state of the spacecraft and provide feedback to EXEC. The other component of MIR, mode reconfiguration (MR), serves as a recovery expert, taking as input a set of EXEC constraints to be established or maintained, and recommends a recovery action to EXEC that will achieve those constraints. MIR provides both the MI and MR functions using a single core algorithm and a single declarative model.

Fault protection in RA happens at two different levels:

- *Low-level fault protection loop.* This involves EXEC and MIR in the context of executing a single PS-generated task. Suppose that EXEC is commanding MICAS power on in order to ensure that MICAS is on during the “MICAS take OPNAV images” PS task. It does so by sending an appropriate command to the power driver. MI observes the command and, on the basis of its previous state estimate and its models, predicts the likely next state in which the system will be. This prediction provides a qualitative description of the sensor readings MIR should observe from the spacecraft (e.g., the switch sensor and current sensor should be consistent with MICAS being on). If the expected observations are not received, MI uses its model to hypothesize the most likely cause of the unexpected observations in terms of failures of the spacecraft’s components. The information about the new state of the spacecraft hardware is sent to EXEC, which now asks MIR for an action to correct the problem. MIR now activates MR, which, using the same model, determines the least-cost system state that satisfies EXEC’s request and one that is reachable from the fault mode. MIR then gives EXEC the first action in a possible sequence that will take the system to that state. Such a recovery might involve resetting a device, attempting a command again, or a complex reconfiguration of the spacecraft to enable a functionally redundant system. EXEC executes the recovery action, under the watchful eye of MIR, and receives further actions from MIR if needed by the recovery process. When the recovery is complete, EXEC continues executing the PS task in a nominal fashion. Note that during this entire process the original PS task is still active and in a “nominal” state. This depends on the time allocated to the task including enough slack to tolerate variations during execution that can be handled by low-level fault protection.
- *High-level fault protection loop.* This involves EXEC and PS. Assume that all recovery actions suggested by MR fail and no more recovery actions are available. MIR infers that MICAS is unusable and communicates this to EXEC. This means that there is no way to execute a command necessary for the success of the “MICAS take OPNAV images” task. Moreover, the assumed conditions for other tasks that may be present

in the plan in the future may now be invalidated. Therefore, EXEC terminates task execution with a failure, discards the rest of the plan, and immediately commands the spacecraft to enter an appropriate “RA standby” mode.¹ EXEC then activates PS by communicating to it the current state of the spacecraft and asks for a new plan. After receiving the initial state from EXEC and the goals from MM, PS generates a new plan that achieves the goals as best as possible within the new, degraded spacecraft configuration. When the plan is ready, PS sends it to EXEC. EXEC then exits the “RA standby” state and resumes normal operations by starting the execution of the new plan.

With the above capabilities, RA allows implementation of *fail-operational* behaviors under a much broader range than is possible in traditional spacecraft commanding. Traditionally, only critical sequences (e.g., Saturn orbit insertion for Cassini) are designed to tolerate a large number of faults without requiring “safing” of the spacecraft. This depends on the cost of analysis and implementation of these sequences. Therefore, in less critical mission phases, a fault event usually requires the intervention of the ground operations team to correct it. With RA, the cost of implementing these scenarios is significantly reduced, making possible an increase of mission productivity and a reduction of cost of operations.

1.2 Detailed Validation Objectives

Validation of a technology with the complexity and the pervasive systemic impact of RA required attention to several different aspects and dimensions.

The first and most obvious objective was to validate the fact that RA could autonomously command a system as complex as a spacecraft for an extended period of time. This translated into the following list of objectives for each RA component.

1.2.1 PS/MM Validation Objectives—

- Generate plans onboard the spacecraft.
- Reject low-priority, unachievable goals.
- Replan following a simulated failure.
- Enable modification of mission goals from ground.

1.2.2 EXEC Validation Objectives—

- Provide a low-level commanding interface.
- Initiate onboard planning.
- Execute plans generated onboard and from the ground.
- Recognize and respond to plan failures.
- Maintain required properties in the face of failures.

¹ Note that this is a standby situation only from the perspective of RA. From the point of view of FSW, “RA standby” mode is not a fault mode and does not require FSW fault protection.

1.2.3 MIR Validation Objectives—

- Confirm executive command execution.
- Demonstrate model-based failure detection, isolation, and recovery.
- Demonstrate the ability to update MIR state via ground commands.

1.2.4 Other Objectives—Other validation objectives addressed the impact of the introduction of RA into a “traditional” spacecraft software architecture. From the outset, RA was designed to work in conjunction with existing FSW modules and not to replace them. As a result, fidelity control provided by RA depends on the scope and detail of the spacecraft models. The challenge was to demonstrate that such cooperative arrangement with FSW could indeed be carried out. This consisted of modeling within RA only a specific set of spacecraft subsystems and allowing conventional techniques of FSW control to deal with the remaining control modes of the craft. While there are no software or architectural limitations that would disallow RA to command all subsystems for an extended period of time, the fielding of RA on DS1 was also meant to provide a credible demonstration of the fact that autonomy concepts could be applied within a well-defined scope.

Even within the scope of the autonomy demonstration, it was important to show that adopting RA was not an “all or nothing” proposition and could be commanded with different autonomous-operation levels. Table 1 shows the possible RA autonomy levels, all the way from having EXEC issuing low-level commands from a low-level script analogous to a traditional command (autonomy level 2), to preparing a plan on the ground and uplinking it to the spacecraft for execution (autonomy level 3), to providing closed-loop planning and execution on the spacecraft (autonomy level 6). The DS1 autonomy experiment was designed from the outset to begin at level 3 to build confidence and then migrate to level 6.

Table 1. Autonomy Levels of RA

Level	Ground System	Onboard PS	Onboard EXEC
1	Prepare real-time commands	None	None (executed w/o EXEC involvement)
2	Prepare sequence	None	Execute sequence
3	Prepare plan, upload to EXEC as script	None	Execute plan; “Scripted mode”
4	Prepare plan, upload to planner as goals	Confirm and pass thru the planner	Execute plan; “Planner Mode”
5	Prepare plan, including some unexpanded goals	Complete the plan	Execute plan
6	Define goals	Prepare plan	Execute plan

The final set of validation objectives involved the development process for autonomy software. This covered a number of separate items:

- Integration of RA with the DS1 FSW, a large and complex system written in a language (C) different from RA (Lisp).
- Adaptation of RA models and scenarios to reflect operational constraints imposed by the flight team, even late in the development process.
- Achievement of high-level of confidence by the DS1 spacecraft team by going through a rigorous test regimen dictated by the team on high-fidelity testbeds.

The level of achievement for each validation objective is discussed below.

1.3 Performance Envelope

Note that these performance and resource figures refer to RA as flown on Deep Space 1 in 1999 in Lisp. Each of the RA engines has been or is being re-architected and ported to C or C++. These new systems may exhibit significantly different performance characteristics:

- Memory—32 Mbytes memory peak, 20 average.
- CPU—
 - RAX ran at priority level just below that of DS1 sequencer (very low).
 - 20% of CPU when planner is idle (only EXEC and MIR are running).
 - 45% of CPU while planner is running (PS, EXEC, and MIR all running).
- The time required to generate plans depends on the plan's complexity. RAX plans took 50 to 90 minutes to generate.
- Telemetry—An average of 10 bits per second. This includes notification as each activity in the plan is executed, current diagnosis for each device monitored by MIR, and a summary of the planner's plan-generation progress. Similar telemetry would be needed for future science missions.
- File space—140 KB for support files, plus approximately 100 KB per stored plan, depending on plan complexity (proportional to number of activities in the plan). Compressed binary executable was 4 MB. At most one plan needs to be stored, though all plans were stored during RAX for validation purposes. RAX also generated a 1MB log.

1.4 Technology Details

RA consists of general-purpose reasoning engines and mission-specific domain models. The engines make decisions and command the spacecraft based on the knowledge in the models. This section describes the details of the reasoning engines and how they interact. The DS1 domain models developed for the flight experiment will be discussed in the flight experiment section.

1.4.1 Planner/Scheduler—PS provides the core of the high-level commanding capability of RAX. Given an initial, incomplete plan containing the initial spacecraft state and goals, PS generates a set of synchronized high-level activities that, once executed, will achieve the goals. In the spacecraft domain, planning and scheduling aspects of the problem need to be tightly integrated. The planner needs to recursively select and schedule appropriate activities to achieve mission goals and any other sub-goals generated by these activities. It also needs to synchronize activities and allocate global resources over time (e.g., power and data storage capacity). Subgoals may also be generated due to limited availability of resources over time. For example, it may be preferable to keep scientific instruments on as long as possible (to maximize the amount of science gathered). However, limited power availability may force a temporary instrument shutdown when other more mission-critical subsystems need to be functioning. In this case, the allocation of power to critical subsystems (the main result of a scheduling step) generates the subgoal “instrument must be off” (which requires the application of a planning step).

PS is able to tune the order in which decisions are made to the characteristics of the domain by considering the consequences of action planning and resource scheduling simultaneously. This helps keep the search complexity under control. This is a significant difference with respect to classical approaches both in Artificial Intelligence and Operations Research, where action planning and resource scheduling are addressed in two sequential problem-solving stages, often by distinct software systems (see [14]).

Another important distinction between PS and other classical approaches to planning is that, in addition to activities, the planner also “schedules” the occurrence of states and conditions. Such states and conditions may need to be monitored to ensure that, for example, the spacecraft is vibrationally quiet when high-stability pointing is required.

These states can also consume resources and have finite durations and, therefore, have very similar characteristics to other activities in the plan. PS explicitly acknowledges this similarity by using a unifying conceptual primitive, the *token*, to represent both actions and states that occur over time intervals of finite extension. Examples of token semantics details are given further along in this section.

PS consists of a heuristic search engine that deals with incomplete or partial plans. Since the plans explicitly represent time in a metric fashion, the planner makes use of a *temporal database*. As with most causal planners, PS' beginning plan is incomplete; PS attempts to make the plan more complete by posting more constraints in the database.

These constraints originate from the goals and from constraint templates stored in a domain *model* of the

spacecraft. The temporal database and the facilities for defining and accessing model information during search are provided by the Heuristic Scheduling Testbed System (HSTS). The planning engine searches the possible plans for one that satisfies the constraints and achieves the goals. The action definitions determine the space of plans. The constraints determine which of these plans are legal and heavily prune the search space. The heuristics guide the search in order to increase the number of plans that can be found within the time allocated for planning. Figure 2 describes the PS architecture. Additional details on the planner algorithm and its correctness can be found in [10].

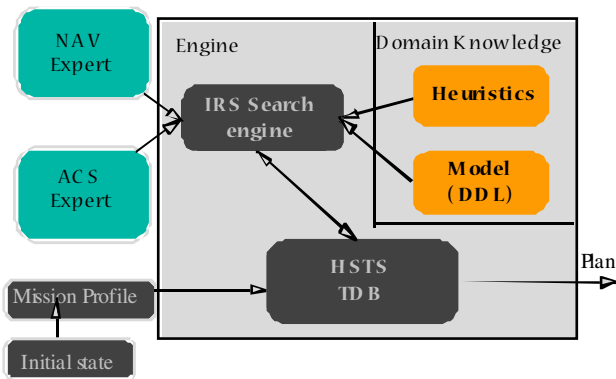


Figure 2. Planner/Scheduler Architecture

The model describes the set of actions, how goals decompose into actions, the constraints among actions, and resource utilization by the actions. For instance, the model will encode constraints such as “do not take MICAS images while thrusting” or “ensure that the spacecraft does not slew when within a DSN communication window.” These constraints are encoded in a stylized and declarative form called the Domain Description Language (DDL).

In conventional modes of writing flight software, the constraints in the domain are mixed with the control information. In the model-based approach of RA, the domain model is a distinct entity that encodes the mission-specific flight rules. This means that (in the case of PS) not only are the core engines (the HSTS Temporal Database [TDB] and the Search Engine) reusable across missions, but that the model can be manipulated independently of any other piece of the flight code. (Note that since the heuristics search control information is model dependant, this module would be impacted also.) In addition, the richness of the representation and the declarative form of DDL ensures that mission/systems engineers can have a substantially easier job of understanding and verifying the implementation of the flight rules in RA than would have been possible in conventional FSW. These are some of the advantages that RA brings to a mission.

Each subsystem in the model is represented in the PS database as a set of dynamic *state variables* whose value is tracked over time. *Timelines* are treated as instantiations of state variables and are used interchangeably with state variables in this report. Each dynamic state variable can assume one or more values. A token is associated with a value of a state variable occurring over a finite time interval. Each value has one or more associated *compatibilities* (i.e., patterns of constraints between tokens). A legal plan will contain a token of a given value only if all temporal constraints in its compatibilities are satisfied by other tokens in the plan. Figure 3 shows an example of a set of compatibilities with temporal constraints.

```
(Define Compatibility
;; compats on SEP_Thrusting
(SEP_Thrusting ?heading ?level ?duration)
:compatibility_spec
(AND
  (equal (DELTA MULTIPLE (Power) (+ 2416
    Used)))
  (contained_by (Constant_Pointing
    ?heading)
  (met_by (SEP_Standby))
  (meets (SEP_Standby)))
)

(Define Compatibility
;; Transitional Pointing
(Transitional_Pointing ?from ?to ?legal)
:parameter functions
(?duration_ <- APE_Slew_Duration (?from
  ?to ?start_time))
(?legal_ <- APE_Slew_Legality (?from
  ?to ?start_time))
:compatibility_spec
(AND
  (met_by (Constant_Pointing ?from))
  (meets (Constant_Pointing ?to)))
)

(Define Compatibility
;; Constant Pointing
(Constant_Pointing ?target)
:compatibility_spec
(AND
  (met_by (Transitional_Pointing *
?target
    LEGAL))
  (meets (Constant_Pointing ?target *
    LEGAL)))
)
)
```

Figure 3. Temporal Constraints in DDL

The first compatibility indicates that the master token (which is at the head of the compatibility) is **SEP_Thrusting** (when the Solar Electric Propulsion [SEP] engine is producing thrust²), which must be immediately preceded and followed by a **SEP_Standby** token (when the

² Solar Electric Propulsion (SEP) is synonymous with IPS.

SEP engine is in a standby mode but has not been completely shut off). The master token must be temporally contained by a constant pointing token; the complete thrusting activity requires 2416 Watts of power. The **Constant Pointing** token implies that the spacecraft is in a steady state aiming its camera towards a fixed target in space. **Transitional Pointing** tokens describe an activity when the spacecraft slews. Figure 4 gives a visual rendering of these compatibilities.

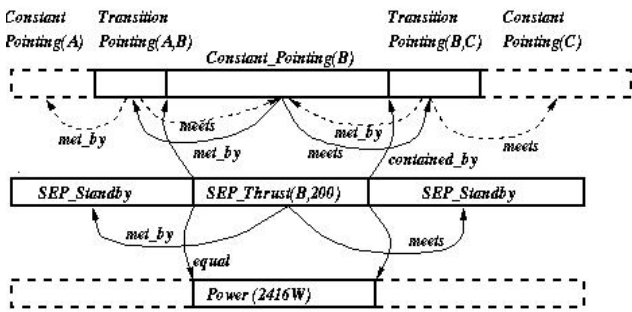


Figure 4. A Plan Fragment Formed by a DDL Model

The timeline approach to modeling is also driven by strong software engineering principles. In a complex domain with different individuals and organizations with varying expertise, timelines provide disparate views of the same domain model across organizational boundaries. For instance, the ground team might want to own and access timelines relating to communication coverage and when DSN access is available, while the attitude control team might want to place high-level goals on the attitude timeline.

Four distinct kinds of state variables are identified. A *goal* timeline will contain the sequence of high-level goals that the spacecraft can satisfy (e.g., the navigate goal described previously). Goal timelines can be filled either by ground operators or by onboard planning experts seen by PS as goal generators. For example, in order to generate the portion of the plan that commands the IPS engine, PS interrogates NAV, which returns two types of goals: the total accumulated time for the scheduling horizon and the thrusting profile to be followed. These two types of information are laid down on separate goal timelines.

Expected device-health information over time is tracked by *health* timelines. The expected profile is communicated by EXEC to PS in the initial spacecraft state. EXEC can communicate that the health of a device has changed even if no fault has occurred. Another kind of state variable is an *internal* timeline. These are only used by the planner to internally organize goal dependencies and subgoaling. Finally, an *executable* state variable corresponds to tasks that will be actually tracked and executed by EXEC.

The RAX PS treats all timelines and tokens within a simple, unified search algorithm. This has advantages. The ground team could force certain behaviors of the spacecraft by including in the mission profile explicit tokens on executable timelines. The additional tokens will be treated by PS as goals, will be checked against the internal PS model, and missing supporting tasks will be automatically expanded to create an overall consistent plan. This will greatly facilitate the work of the ground team. For DS1, such models were understandably more comprehensive and complex, with more timelines, tokens, and compatibilities between differing token types, and required careful consideration during modeling to ensure that interactions between timelines do not result in unanticipated and harmful behaviors generated by the planner.

When a science mission wants to fly the RA planner, primary tasks to be adapted to the mission will be:

- Perform knowledge acquisition to determine all the spacecraft flight rules.
- Encode these flight rules in the DDL model of the spacecraft.
- Design the search control heuristics that will be needed to ensure that the planner is able to produce a valid plan within specified resource (time, CPU) bounds.

Note that this is not to suggest that models can be or ought to be built in an all-or-nothing fashion. On the contrary, the team strongly believes that coming up with a viable plan encapsulating all domain flight rules is an incremental process (You build some and test some).

As mentioned previously, since the underlying search algorithm does not need to be rewritten, the mission will save costs in revalidating the control system and can confine itself to building and validating the model and search control heuristics. Efforts are underway at NASA’s Ames Research Center to implement automated tools that will ensure that full coverage of the behaviors anticipated by the models is simulated during the modeling process. Additional efforts are also underway to automatically generate the heuristics from a given model of the domain. This will further allow mission designers and systems staff to build robust and complex models on their own without relying on the AI technologists themselves.

Additional details about the planner can be found in [5 to 7] and [10 to 12].

1.4.2 Executive—The Smart Executive (EXEC) is a multi-threaded, reactive-commanding system. EXEC is responsible for sending the appropriate commands to the various flight systems it is managing. EXEC can replace the traditional spacecraft sequencer or can be used in conjunction with a traditional sequencer to command a complex subsystem (e.g., interferometer).

EXEC is a multi-threaded process that is capable of asynchronously executing commands in parallel. In addition to a traditional sequencer's capabilities, EXEC can:

- Simultaneously achieve and maintain multiple goals (i.e., system states) by monitoring the success of commands it issues and reactively re-achieving states that are lost.
- Perform conditional sequencing. Commands can be dependent on conditions that occur at execution time.
- Perform event-driven commands, as opposed to traditional sequencers that are time-driven (i.e., taking a sequence of pictures based on the results of monitoring a range sensor).
- Perform high-level commanding and run-time task expansion. EXEC provides a rich procedural language, Execution Support Language (ESL) [1], in which spacecraft software/model developers define how complex activities are broken up into simpler ones. To increase robustness, a procedure can specify multiple-alternate methods to achieve a goal.
- Perform sequence recovery. In the event an executing sequence command fails, EXEC suspends executing the failed sequence and attempts a recovery, either by executing a pre-specified recovery sequence, such as reissuing the command or consulting a recovery expert (e.g., MIR). Once the desired state of the failed command is achieved, the suspended sequence is restarted.
- Execute a temporally-flexible sequence (or plan). In order to decrease the probability of a sequence failing, time ranges can be specified for executing and achieving the desired state for each command.
- Manage resources. As a multi-threaded system, EXEC can work on multiple tasks simultaneously. These tasks may compete for system resources within the constraints not already resolved by ground or the planner. EXEC manages abstract resources by monitoring resource availability and usage, allocating resources to tasks when available, making tasks wait until their resources are available, and suspending or aborting tasks if resources become unavailable due to failures (such as a device breaking). See [1] and [2] for a more detailed discussion.

Figure 5 illustrates key functions of EXEC.

EXEC achieves multiple tasks, sending the appropriate control commands (decomposed from high-level commands) to the flight software. The tasks also lock properties that need to be maintained. For example, if a task commands a switch ON, the switch property will be locked ON. Monitors (and MIR) determine if it is consistent to believe that the switch is ON. Since EXEC stores this state in its state database should the inferred state of the switch change, the database

will be updated and an event created, thereby signaling a change. If the signaled event violates a property lock, an EXEC property thread interrupts those tasks that subscribed to that property lock. It will then attempt to achieve the state of the switch being ON using its own recovery mechanism or by consulting a recovery expert (e.g., MIR). If the switch cannot be turned ON in time, a hard deadline that is being tracked is missed; in response EXEC commands the spacecraft into a safe, wait state while it requests a new plan from the planner that takes into account that the switch cannot be turned ON.

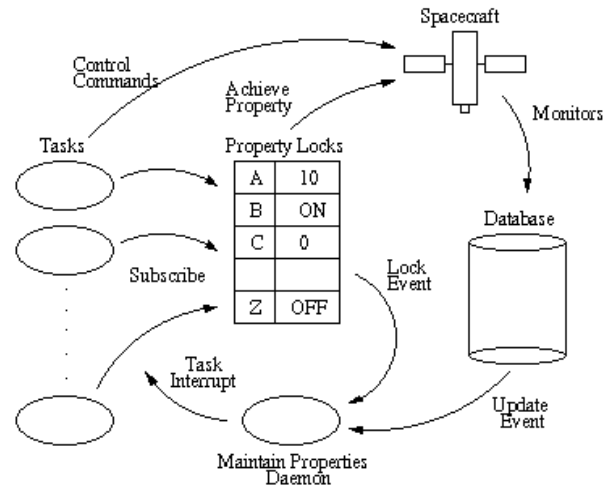


Figure 5. An Overview of the Remote Agent Executive

Recoveries may be as simple as sending another command to turn a switch ON, or may be complex, such as when multiple subsystems are tightly coupled. For example, consider two coupled DS1 subsystems: the engine gimbal and the solar panel gimbal. A gimbal enables the engine nozzle to be rotated to point in various directions without changing the spacecraft orientation. A separate gimbal system enables the solar panels to be independently rotated to track the sun. In DS1, both sets of gimbals communicate with the main computer via a common gimbal drive electronics (GDE) board. If either system experiences a communications failure, one way to reset the system is to power-cycle (turn on and off) the GDE. However, resetting the GDE to fix one system also resets the communication to the other system. In particular, resetting the engine gimbal to fix an engine problem causes temporary loss of control of the solar panels. Thus, fixing one problem can cause new problems. To avoid this, the recovery system needs to take into account global constraints from the nominal schedule execution, rather than just making local fixes in an incremental fashion; the recovery itself may be a sophisticated plan involving operations on many subsystems.

Domain-code developers use ESL to create high-level commands that EXEC decomposes and executes at run-time depending on the spacecraft state. The ESL code in Figure 6 illustrates multiple methods for achieving IPS thrusting at a desired level depending on the current state of execution. If IPS is in standby mode, ACS is commanded to change control modes only after the desired IPS thrust level has been confirmed.

```
(to achieve (IPS THRUSTING ips level)
  ((ips is in standby state p ips)
   (sequence (achieve (power on? 'ega-a))
    (command with confirmation
      (send-ips-set-thrust-level level))
    (command with confirmation
      (send-acs-change-control-mode
        :acs-tvc-mode))))
  ((ips in thrusting state p ips)
   (command with confirmation
     (send-ips-change-thrust-level level)))
  (t (fail :ips achieve thrusting)))
```

Figure 6. Multiple Methods in ESL for Achieving Thrust

EXEC and its commanding language, ESL, are currently implemented using multi-threaded Common Lisp. A new version of EXEC is currently under development in C/C++. The internal EXEC code is designed in a modular, layered fashion so that individual modules can be designed and tested independently. Individual generic device knowledge for RAX is implemented based on EXEC's library of device management routines to support addition of new devices and reuse of the software on future missions.

More details about EXEC can be found in [1 to 3] and [7].

1.4.3 Diagnosis and Repair—The diagnosis and repair engine of RA is the Mode Identification and Reconfiguration (MIR) system. MIR eavesdrops on commands that are sent to the onboard hardware managers by EXEC. As each command is executed, MIR receives observations from spacecraft sensors, abstracted by monitors in lower-level device managers for ACS, Bus Controller, and so on. MIR uses an inference engine called Livingstone to combine these commands and observations with declarative models of the spacecraft's components to determine the current state of the system (mode identification [MI]) and report it to EXEC. EXEC may then request that Livingstone return a set of commands that will recover from a failure or move the system to a desired configuration (mode reconfiguration [MR]). Figure 7 illustrates the data flow among the spacecraft, EXEC, and Livingstone.

MI is responsible for identifying the current operating or failure mode of each component in the spacecraft, allowing EXEC to reason about the state of the spacecraft in terms of component modes, rather than in terms of low-level sensor

values. MR is responsible for suggesting reconfiguration actions that move the spacecraft to a configuration that achieves all current goals as required by PS and EXEC, supporting the run-time generation of novel reconfiguration actions. Though in RA, Livingstone is used only to recover following a component failure. Livingstone's MR capability can be used to derive simple actions to reconfigure the spacecraft at any time. Thus, Livingstone can be viewed as a discrete model-based controller in which MI provides the sensing component and MR provides the actuation component. Livingstone uses a single set of models and core algorithms to provide both the MI and MR functions.

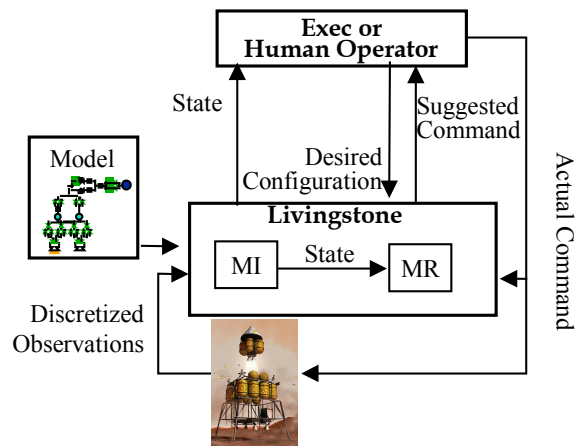


Figure 7. Livingstone Processing Cycle

To use Livingstone, one specifies how the components of interest are connected. For each component type, one then specifies a finite state machine that provides a description of the component's nominal and failure behavior.

Figure 8 graphically depicts a Livingstone model of the Cassini main-engine subsystem. An important feature is that the behavior of each component state or mode is captured using abstract, or qualitative, models [3, 4]. These models describe qualities of the spacecraft's structure or behavior without the detail needed for precise numerical prediction, making abstract models much easier to acquire and verify than quantitative engineering models. Examples of qualities captured are the power, data, and hydraulic connectivity of spacecraft components and the directions in which each thruster provides torque. While such models cannot quantify how the spacecraft would perform with a failed thruster, for example, they can be used to infer which thrusters are failed given only the signs of the errors in spacecraft orientation. Such inferences are robust since small changes in the underlying parameters do not affect the abstract behavior of the spacecraft.

Livingstone's abstract view of the spacecraft is supported by a set of fault-protection monitors that classify spacecraft sensor output into discrete ranges (e.g., high, low, nominal) or symptoms (e.g., positive X-axis attitude error). One

objective of the RA architecture was to make basic monitoring capability inexpensive so that the scope of monitoring could be driven from a system engineering analysis instead of being constrained by software development concerns. To achieve this, monitors are specified as a dataflow scheme of feature extraction and symptom-detection operators for reliably detecting and discriminating between classes of sensor behavior.

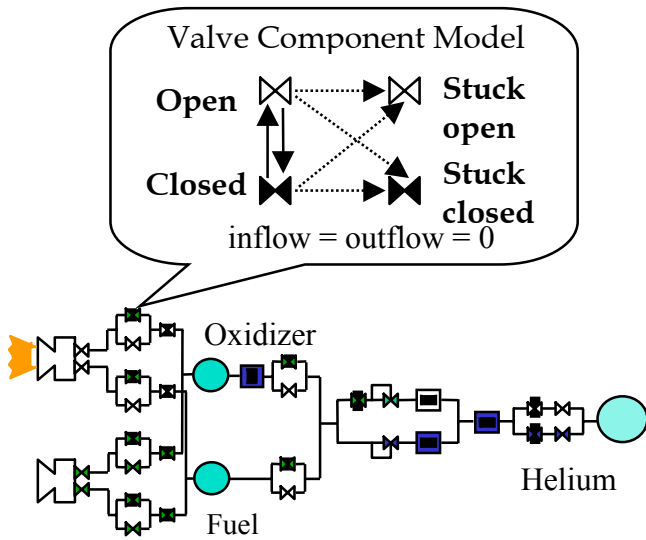


Figure 8. Livingstone Model of the Cassini Main Engine Subsystem

The software architecture for sensor monitoring is described using domain-specific software templates from which code is generated. Finally, all symptom detection algorithms are specified as restricted Harel-state transition diagrams reusable throughout the spacecraft. The goals of this methodology are to reuse symptom-classification algorithms, reduce the occurrence of errors through automation and to streamline monitor design and test.

It is important to note that the Livingstone models are not required to be explicit or complete with respect to the actual physical components. Often, models do not explicitly represent the cause for a given behavior in terms of a component's physical structure. For example, there are numerous causes for a stuck switch: the driver has failed, excessive current has welded it shut, and so on. If the observable behavior and recovery for all causes of a stuck switch are the same, Livingstone need not closely model the physical structure responsible for these fine distinctions.

Models are always incomplete in that they have an explicit unknown failure mode. Any component behavior that is inconsistent with all known nominal and failure modes is consistent with the unknown failure mode. Therefore,

Livingstone can infer that a component has failed, though failure was not foreseen and left unmodeled because it was not possible.

By modeling only to the detail level required to make relevant distinctions in diagnosis (distinctions that prescribe different recoveries or system operations), a system with qualitative "common-sense" models that are compact and quite easily written can be described. Figure 9 provides a schematic overview of Livingstone's processing.

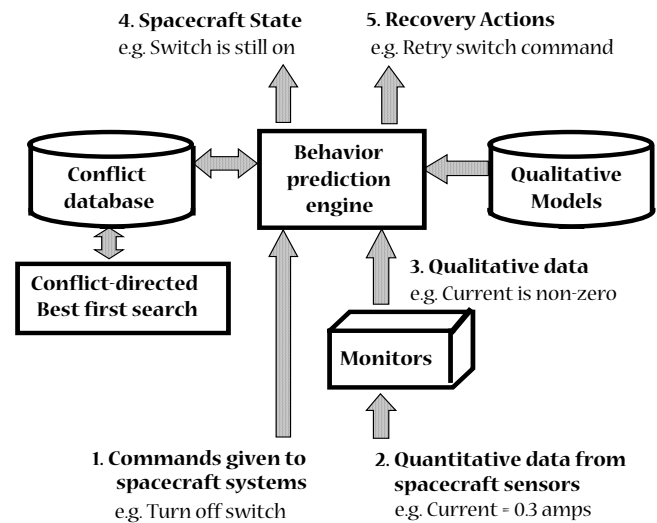


Figure 9. Schematic of Livingstone Processing

Livingstone uses algorithms adapted from model-based diagnosis (see [9]) to provide the above functions. The key idea underlying model-based diagnosis is that a combination of component modes is a possible description of the current state of the spacecraft only if the set of models associated with these modes is consistent with the observed sensor values. Following de Kleer and Williams [8], MI uses a conflict-directed, best-first search to find the most likely combination of component modes consistent with the observations. Analogously, MR uses the same search to find the least-cost combination of commands that achieve the desired goals in the next state. Furthermore, both MI and MR use the same system model to perform their function.

The combination of a single-search algorithm with a single model, and the process of exercising these through multiple uses, contributes significantly to the robustness of the complete system. Note that this methodology is independent of the actual set of available sensors and commands. Furthermore, it does not require that all aspects of the spacecraft state are directly observable, providing an elegant solution to the problem of limited observability.

The use of model-based diagnosis algorithms immediately provides Livingstone with a number of additional features.

First, the search algorithms are sound and complete, providing a guarantee of coverage with respect to the models used. Second, the model-building methodology is modular, which simplifies model construction and maintenance and supports reuse. Third, the algorithms extend smoothly to handling multiple faults and recoveries that involve multiple commands. Fourth, while the algorithms do not require explicit fault models for each component, they can easily exploit available fault models to find likely failures and possible recoveries.

Since the flight experiment, Livingstone has been ported to C++ and significantly improved in the areas of both MI and MR. The improved Livingstone is scheduled to be test flown on both the X-34 and X-37 experimental vehicles. Additional technical details about Livingstone can be found in [4] and at <http://ace.arc.nasa.gov/postdoc/livingstone>

1.5 Subsystem Interdependencies

The Remote Agent Experiment Manager (RAXM) is the flight software interface to the Remote Agent Experiment (RAX) and isolates the RA software from the rest of the FSW via a set of clean application programming interfaces (APIs).

In addition, RAXM provides a terminal in the point-to-point message-passing protocol used by the DS1 flight software (see Figure 1). RAXM in particular is tasked with handling three messages throughout the mission: RAX-START, RAX-STOP, and RAX-ABORT; RA software is operational only during the times between a RAX-START and either RAX-STOP or RAX-ABORT. RAX-START is used by RAXM to decompress the RAX Lisp image and initiate RAX control. The RAX-STOP is implemented to cleanly terminate RAX at the end of the experiment under nominal circumstance, while the RAX-ABORT is intended to kill the RAX process in the event of an abnormality detected by RAXM. At all other times, RAXM discards all incoming messages, allowing all FSW subsystems that interact with RAX to be ignorant of the RAX state.

When RA runs, RAXM handles and dispatches all incoming messages related to RA: some of the messages are handled by RAXM, others are passed through to RAX itself. Similarly, outgoing messages from RAXM can be due either to RAXM or to RAX itself.

Like the code for other flight software subsystems, RAXM is written in the C programming language and is part of the launch load. As a result, the interfaces for RAX needed to be specified early.

The computational resources (CPU fraction, memory space, telemetry buffers and downlink, etc.) required by RAXM when RA was not running were insignificant. This was, by

design, a way to mitigate the impact of the RA technology demonstration on DS1.

1.6 Preparing Lisp for Flight

One important aspect of the RA preparation for flight was the preparation of Lisp for flight. The RA software development and runtime environment was based on Common Lisp, in particular the Harlequin Lispworks product. The use of Lisp was appropriate given the background of the RA developers, the early inheritance of code libraries, and the hardware independence of the high-level software interfaces between RA and the rest of the flight software. However, with the choice of Lisp came some unique challenges. These challenges fell into two rather broad categories: resource constraints and flight-software interfaces.

To fit within the 32 MB memory allocation and the CPU fraction constraints, the RA team thoroughly analyzed their code for memory and performance inefficiencies and applied a “tree-shaking/transduction” process to the Lisp image. The analysis is, of course, common for any high-performance software. However, transduction is Lisp-specific and arises from the tight coupling of the Lisp runtime and development environments. Transduction removes the unneeded parts of the development environment (e.g., the compiler, debugger, windowing system). The result is a significantly smaller image, both in terms of file system and runtime memory. During RA testing, peak memory usage was measured at about 29 MB. Upon completion of the transduction process, the RA Lisp image was compressed by a factor of about 3 to 4.7 MB and uplinked to the spacecraft. Onboard decompression was initiated at the start of each RA run, with the file being inflated directly into the 32-MB RA memory space. Use of this custom compression drastically reduced the file-uplink time and kept RA-file space usage within the agreed limits.

Added to the challenge of working within resource constraints was the challenge of working out the complicated interfaces between RA and the rest of the flight software. The flight software was written in the C programming language and ran on the VxWorks operating system. Lisp and C interacted through Lisp’s foreign function interface. This interface was the source of many early problems, primarily caused by discrepancies between data structure alignments assumed by the Lisp and C compilers. These problems were quickly discovered and resolved with the help of an extensive test suite that analyzed many function-parameter variations.

Another problem arose in preparing the Lisp multi-threading system for flight. Originally, the Lisp thread scheduler relied on a high-frequency, external, periodic wakeup call issued at interrupt level. However, this went against the design principles of the DS1 flight software. Hence, Lisp’s

approach—to thread preemption to use a lower frequency wakeup call implemented with flight software timing services—had to be significantly changed.

Most of the late integration problems with RA Lisp arose because of the VxWorks port. As RA moved from testbed to testbed, ever closer to the final spacecraft configuration, low-level Lisp problems arose. The problems were consistently of two types: a function assumed by Lisp to be present was not present or a function was present but did not perform as expected by Lisp. The first type of problem was resolved by consistent application of a detailed RA and FSW build process. The second type of problem was addressed on a case-by-case basis. Solutions to these problems were made difficult due to the reduced debugging visibility as testbeds assumed the spacecraft configuration. The entire undertaking benefited from the dedicated efforts of both Harlequin and the DS1 FSW team.

2.0 THE REMOTE AGENT EXPERIMENT

During the DS1 mission, the Remote Agent technology was validated with an experiment, the Remote Agent Experiment (RAX). The flight experiment was conducted between May 17, 1999, and May 21, 1999, and achieved all of the technology-validation objectives. However, the story is incomplete without reporting the valuable data gathered during development and testing on the ground. In the case of RA, this is particularly important since the technology is intended as a tool to support system engineering and operations for a mission, rather than simply provide the resulting autonomous capabilities. By quantitatively analyzing the history of RAX's development, we can evaluate how well the current state of the technology supports its ultimate goals. This can also help identify weak points that require further research and development.

RAX and the team attempt to evaluate the development and testing experience with respect to the features of the technology is described here. First, RAX must be put into the larger perspective of RA's technology evolution. Then the subsystems and fault modes modeled, the experiment scenarios, and the expected in-flight behavior are described. Then how RAX was developed and validated and the details of the flight experiment are discussed. Then the effectiveness and cost of development and testing are successively analyzed. The analysis is supported by the actual problem reports filed in the RAX problem-tracking system during development. Lessons learned conclude this section.

2.1 Historical Perspective

Development of the RA technology effectively started in May 1995. At that time, spacecraft engineers from JPL and Artificial Intelligence (AI) technologists from Ames Research Center (ARC) and JPL started working together

on the New Millennium Autonomy Architecture rapid Prototype (NewMAAP), a six-month effort intended to assess the usability of AI technologies for onboard flight operations of a spacecraft [17]. NewMAAP yielded proof of concept of an autonomous agent that formed the fundamental blueprint for Remote Agent. NewMAAP also helped build the team of technologists that continued development of Remote Agent on DS1.

The successful demonstration of NewMAAP in November, 1995 led to the selection of RA as one of the components of the autonomy flight software for DS1. Between December 1995 and April 1997, the RA team was part of the DS1 flight software team. This led to the development of the three engines of the RA component technologies and included a substantial speed up of the MIR inference engine (see [4]), the design and implementation of the ESL language used by EXEC (see [1]), and the design and implementation of the heuristic search engine for PS together with the language to formulate search heuristics.

Regarding the overall Remote Agent architecture, the fault protection protocols were designed and implemented, both at low level (involving EXEC and MIR) and at the high level (involving EXEC and PS). During this period, the team acquired much of the high-level system knowledge needed to model DS1 cruise operations (including image acquisitions of beacon asteroids for AutoNAV, timed IPS thrusting, and file uplink and downlink) and other DS1 capabilities required for asteroid encounter activities.

In March 1997, the DS1 autonomy flight software was substantially overhauled and DS1 adapted the Mars Pathfinder (MPF) flight software as the basis for its flight software. Also, RA was re-directed to become an experiment operating for at most six days during the mission on a cruise scenario, including AutoNAV orbit determination and IPS-timed thrusting. RAX re-used much of the software developed during the previous autonomy flight software phase of DS1. RAX focused on the process of testing each RA component, integrating and testing them into the complete RA, and integrating and testing RA together with the DS1 flight software on the flight processor. Shortcomings found during the development and testing phases required several extensions and re-designs of domain models and the reasoning engines.

This document is focused solely on RAX and makes use of the detailed development and testing records maintained during this phase. However, when the technology readiness conclusion is presented, it will reflect the entire Remote-Agent's development history.

Table 2 shows the highlights of the RAX, starting with the RAX development effort after the redirection of the flight software to MPF. Due to this change, a requirement was

imposed on the RAX team to keep interactions with the flight team to a minimum. From the beginning, RAXM was identified as being the primary interface to RA and part of the launch load of DS1; delivery of RAXM was initiated by December after negotiating all interfaces with FSW. This was the only significant interaction the team had with the DS1 flight team till February 1999, three months prior to activation of RAX. Integration of RAX on the *Radbed* high fidelity testbed was completed during April 1998, which allowed the team to understand the timing characteristics of RA in flight. The RAX Software Delivery Review in September allowed the team for the first time to show the DS1 project the progress the team was making and explain the expected behavior of RAX during flight. November of the same year, barely five months before the experiment, was the first time RAX software ran on a *Papabed* after interfacing with the actual FSW. It took another month to actually produce a plan and execute it on this testbed. The RAX delivery entered the final deliverable phase in February 1999 with code development frozen and bug fixes under a strict change-control regime. RAX was finally initiated on DS1 on May 17, 1999.

Table 2. Significant Events for the RAX Project

Event	Date
Start of RAX development	April 1997
Delivery of RAX Manager to flight software	December 1997
RAX integrated on the flight processor	April 1998
Project Software Delivery Review	September 1998
DS1 launch	October 1998
First run of RAX with FSW on high-fidelity hardware simulation	November 1998
Beginning of M5 DS1 project phase	February 1999
RAX experiment	May 1999

Below is a detailed description of the DS1 subsystems modeled in RAX and the scenarios on which RA was exercised during RAX development and testing.

2.2 Domain Models

The team only developed domain models for the subsystems and fault modes that were necessary for the experiment. Table 3 describes the timelines modeled by the planner. Table 4 and Table 5 list the components and module models developed for MIR while Table 6 shows the modeled EXEC timelines. These models captured the following subsystems and resources:

- Ion Propulsion System
Detect and command standby through thrusting states.
- Attitude Control Subsystem.
PS planned attitude changes requested by NAV (IPS attitudes and beacon asteroids) or specified as goals in

the mission profile. These attitudes were restricted in the model to slews that maintained the solar panels on-Sun. For the experiment, the NAV profiles and goals were specified to further limit the attitudes to either high gain antenna (HGA) at Earth (the default attitude and the IPS thrust attitude) or MICAS bore-sight at a beacon asteroid.

- MICAS.
PS planned data takes and low-voltage power on/off. Switch status and commands were modeled, but the switch commands are not actually issued. (See the scenario description for why this is so.)
- Power.
PS tracked predicted peak-power usage for each activity in the plan (e.g., IPS thrusting, MICAS on) and ensured that the total would never exceed the available power from the solar panels, as predicted by the operations team and supplied in the mission profile. MIR modeled a portion of the power distribution system and its relays in order to confirm operation of the switches commanded by RAX and distinguish between failures in the power system and erroneous sensor readings. MIR modeled switches not commanded by RA so that it could request the experiment be aborted if the power system was in a state out of scope for the experiment.
- Reaction Control System.
MIR modeled the thruster pallets, thrusters, and valves of the RCS system in order to determine the health of the various components from errors in attitude and recommend which control mode to utilize.
- Data System.
MIR modeled the 1553 bus and a subset of the remote terminal devices on it in order to monitor for remote terminal hangs and recommend resets. Resetting was limited to the Power Actuation and Switching Module (PASM) instrument. Other remote terminals were modeled in order to allow MIR to request the experiment be aborted if certain out-of-scope data-system problems occurred.
- Sensors.
MIR modeled a subset of the switch position and current sensors onboard DS1 as fallible components in order to allow sensor failure as an explanation for unexpected observations.
- Remote Agent.
PS models aspects of the operation of RA itself. For example, the Planner timeline allows PS to plan time for its next planning activity. The Special Activities timeline allows PS to schedule execution of scripts that (unbeknownst to RA) will cause simulated failures onboard the spacecraft.

Table 3. Summary of Planner Models for RAX

Subsystem	State Variables	Value Types	Compatibilities	Comments
MICAS	Executable: 2 Health: 1	7	14	Models the health, mode, and activity of the MICAS imaging camera. RAX demonstrates fault injection and recovery for this device as part of the 6-day scenario.
Navigation	Goal: 1 Executable: 1 Internal: 1	5	6	To schedule orbit determination (OD) based on picture taking activity.
Propulsion & Thrust	Goal: 2 Executable: 1 Internal: 1	9	12	Based on thrust schedule generated by the NAV module, the planner generates plans to precisely activate IPS in specific intervals based on constraints in the domain model and is the most complex set of timelines and subsystem controlled by the planner.
Attitude	Executable: 1 Health: 1	4	4	Enables the planner to schedule slews between constant pointing attitudes when the spacecraft maintains its panels towards the Sun. The targets of the constant pointing attitudes are imaging targets, Earth (for communication), and thrust direction (for IPS thrusting).
Power Management	Goal: 1 Internal: 1	2	1	Allows the planner to ensure that adequate power is available when scheduling numerous activities simultaneously.
Executive	Goal: 1 Executable: 1	2	7	Allows modeling of low-level sequences, bypassing planner models, giving Mission Ops the ability to run in sequencing mode with the RA.
Planner	Executable: 1	2	2	To schedule when EXEC requests the plan for the next horizon.
Mission	Goal: 1	2	2	Allows MM and PS to coordinate activities based on a series of scheduling horizons updatable by Mission Ops for the entire mission.

Table 4. DS1 Hardware Modeled as Components in MIR

Component Class	# in Model	Modes
ion propulsion system (IPS)	1	Standby, Startup, Steady State Thrusting, Shutdown, Beam Out, Controller Hung, Unknown
remote terminal	6	Nominal, Resettable Failure, Power-cyclable Failure, Unknown
attitude control	1	TVC, X for Y, Z for Y, X for Y Degraded, Z for Y Degraded, X for Y Failed, Z for Y Failed, TVC Failed, Unknown
switch	12	On, Off, Popped On, Popped Off, Stuck On, Stuck Off, Unknown
switch sensor	12	Nominal, Stuck On, Stuck Off, Unknown
current sensor	3	Nominal (reported value = real value), Unknown (values unconstrained)
thruster valve	8	Nominal, Stuck Closed, Unknown
thruster	8	Nominal, Unknown
propellant tank	1	Non-empty, Unknown (thruster hydrazine out or otherwise unavailable)
bus controller	1	Nominal, Unknown
vehicle dynamics	1	Nominal (this is a qualitative description of force and torque)
power bus	3	Nominal (failure considered too fatal and remote to involve in diagnosis)

Table 5. DS1 Hardware Modeled as Modules in MIR

Module	# in Model	Subcomponents
power relay	12	1 switch, 1 switch sensor
power distribution unit	1	12 relays, 3 power buses, 3 current sensors, 1 remote terminal
generic RT subsystem	3	1 remote terminal (models RT for devices MIR does not otherwise model)
IPS system	1	1 IPS, 1 remote terminal
thruster pallet	4	2 thrusters (X facing and Z facing)
reaction control system	1	4 thruster pallets
PASM subsystem	1	1 remote terminal

Table 6. Timelines and Their Respective Tokens by Module (EXEC's perspective)

Module	Timeline	Token	Description	
ACS	Spacecraft Attitude	constant_pointing_on_sun	Point vector at Target, Solar Panels at Sun	
		transitional_pointing_on_sun	Turn vector to Target, Solar Panels at Sun	
		poke_primary_inertial_vector	Small attitude change	
	RCS_Health	rcs_available	Maintain information on thruster status	
	RCS_OK	maintain_rcs	Set and maintain desired RCS mode	
MICAS (Camera)	MICAS_Actions	micas_take_op_nav_image	Take a set of navigation pictures	
		MICAS_Mode	micas_off	Keep MICAS off
			micas_ready	Keep MICAS on
			micas_turning_on	Turn MICAS off
			micas_turning_off	Turn MICAS on
	MICAS_Health	micas_availability	Ensure MICAS is available for use	
Op-Nav	Obs_Window	obs_window_op_nav	Wait for a specified duration	
	Nav_Processing	nav_plan_prep	Send message to prepare navigation plan	
PASM	PASM Available	pasm_monitor	Monitor the PASM switch	
SEP	SEP	sep_standby	Achieve and maintain IPS standby state	
			sep_starting_up	Achieve and maintain IPS start-up
			sep_thrusting	Maintain a thrust level
			sep_shutting_down	Stop thrusting and go to standby state
		SEP_Time Accum	accumulated_thrust_time	Monitor thrust time accumulated
		SEP_Schedule	thrust_segment	Specifies desired thrust level and vector
		SEP_Thrust Timer	max_thrust_time	Set a timer and stop thrusting if time reached
		thrust_timer_idle	Thrust timer is off	
Planner	Planner_Processing	planner_plan_next_horizon	Request and get next plan from planner.	
		script_next_horizon	Run the next scripted plan	
General	EXEC Activity	exec_activity	Execute a low-level sequence file passed as a parameter	
	EXEC_Eval	exec_eval_watcher	Process a specified script	

2.3 Experiment Scenarios

The RAX experiment proposal contained a 12-hour scenario and a 6-day scenario. The 12-hour scenario was designed as a confidence builder for the DS1 project. The 6-day scenario was to be run following successful completion of the 12-

hour scenario. Together, the 12-hour and 6-day scenarios demonstrate all RAX validation objectives and were used for all RAX integration and testing until the beginning of March 1999. Then the DS1 project levied additional constraints on how the spacecraft could be commanded and

specified that RAX should produce 12 hours of thrust or less. The team responded by developing a 2-day scenario that met the additional commanding constraints and provided 12 hours rather than 4 days of thrusting. The DS1 project viewed very favorably the group's ability to quickly respond with a new scenario for these new constraints. Each scenario is described below.

2.3.1 Twelve-hour Scenario—The twelve-hour scenario involves neither onboard planning nor thrusting with IPS. Rather, the plan is generated on the ground, uplinked to the spacecraft, and executed by EXEC and MIR. The scenario includes imaging asteroids with the MICAS camera to support optical navigation, a simulated sensor failure scenario, and demonstration of low-level commanding from a script through RAX to flip a switch. The planning of optical navigation imaging provides the planner the opportunity to reject low-priority, unachievable goals since the optical navigation windows had time only to image a subset of the asteroid goals.

2.3.2 Six-day Scenario—The 6-day scenario includes both onboard planning and operating IPS and is a full-up test of RA. The scenario is divided into 2 planning horizons. At the start of the scenario, PS generates a plan for the first horizon that included MICAS imaging for optical navigation and IPS thrusting. Execution of the first plan also includes a ground command to modify the goals for the second horizon. At the end of the optical navigation window, PS plans to switch off the MICAS camera. However, a stuck-on-failure injection in the camera switch prevents RA from turning off the camera, leading to a plan failure. Repeated attempts to recover the problem fail. This leads to a replan, which produces a second plan with the camera being left on. The second plan also includes an activity to produce a plan for the second horizon (the third plan in the scenario), which is executed back-to-back with the second plan. While the second plan is being executed, the switch-failure injection is undone and ground informs MIR that the switch is now fixed. The execution of the third plan includes IPS thrusting, optical-navigation imaging, and two simulated failures, a communication failure on the 1553 bus and a thruster-valve-stuck-closed failure.

The MICAS stuck-on failure demonstrates how MIR and EXEC can make repeated attempts to recover a camera switch until it is deemed permanently stuck. The 1553 bus remote-terminal failure illustrates successful recovery of communication with a device by resetting its remote terminal (RT). In the ACS thruster-valve-stuck-closed failure, MIR infers from an attitude error and models of the spacecraft dynamics that one of a particular pair of thruster valves is stuck closed. MIR is then able to recommend that no matter which one of the two valves is stuck, switching ACS control modes will mitigate the problem.

2.3.3 Two-Day Scenario—In March 1999, the DS1 project analyzed the 6-day plan and decided that RA should not switch the MICAS camera off after each use due to concerns about thermal effects. In addition, RA would be required to produce at most 12 hours of IPS thrusting to ensure that DS1 would be on track for its asteroid encounter in July 1999.

The 2-day scenario was created that is similar to a compressed 6-day scenario, except that the simulated MICAS switch failure was active for the whole duration of the scenario. This prevented RA from ever switching off the camera. Furthermore, the mission profile was adjusted so that PS would produce plans with only about 12 hours of IPS thrusting. This scenario is similar to the standard DS1 cruise phase, which consists of IPS thrusting punctuated with periodic optical-navigation activities. This baseline demonstrated RAX's basic commanding capabilities.

This scenario retains the simulated faults that exercise RAX's robust fault-response capabilities. Since the team could not depend on failures occurring during the experiment, failures were simulated by injecting false monitor readings consistent with the failures. While simulations are necessary for demonstration, the RAX is fully responsible for responding to real failures within its limited scope occurring during the experiment. To avoid potential conflicts between RAX and the flight-software fault-protection mechanism (FP) the RAX response threshold is a little lower than that FP to allow RAX to detect and respond to faults before FP does. If RAX fails to resolve a fault quickly enough, the FP response would be triggered (since the fault is still active). The FP response is to terminate RAX and resolve the fault.

2.4 RAX Development

RAX was developed on a number of platforms of decreasing processor speed and increasing level of hardware and software fidelity (see Table 7).

The team adopted a continuous-integration-development process with new software capabilities being first developed on the *UNIX* platform. Before they could be incorporated in a software build and be appropriately tagged, new features or bug fixes had to run to completion a representative set of scenarios. As time progressed, testbeds of higher and higher fidelity became available. As this happened, the requirements for acceptance of software modifications became more and more demanding since the scenarios had to run on all available platforms.

Besides the speed of the processors, another factor effecting productivity was the simulated-clock speed. The *UNIX*, *Babybed* and *Radbed* platform made use of a low-fidelity simulation developed by the RAX team, which essentially only simulated the message traffic and the delays in

receiving responses from flight software. This meant that the simulator was allowed to advance the clock at “warp” speed, simulating in a second, several minutes or hours of actual elapsed time. Time warping allowed the team to run to completion the full 6-day scenario in less than an hour, tremendously increasing the productivity during development and testing on such lower-fidelity testbeds.

Table 7. Development Testbeds for RAX

Platform	Fidelity	CPU/OS	Hardware	Availability	Speed
DS1 Spacecraft	Highest	Rad6000 VxWorks	Flight	1 for DS1	1:1
DS1 Testbed	High	Rad6000 VxWorks	Flight spares + DS1 sims	1 for DS1	1:1
Hotbench	High	Rad6000 VxWorks	Flight spares + DS1 sims	1 for DS1	1:1
Papabed	Medium	Rad6000 VxWorks	DS1 sims only	1 for DS1	1:1
Radbed	Low	Rad6000 VxWorks	RAX sims only	1 for RAX	1:1
Babybed	Very Low	PowerPC VxWorks	RAX sims only	2 for RAX	7:1
UNIX	Lowest	SPARC UNIX	RAX sims only	unlimited	35:1

Since the higher-fidelity testbeds could not be warped in time because of interfaces to the actual FSW code, it induced the team to devise reduced-length scenarios that would exercise in a few hours of actual clock time most or all of the functionalities included in the full, multi-day flight scenarios. These shorter scenarios led to exercising RAX under stress conditions complementary to those addressed by the formal test process. As a consequence, continuous integration over the course of testing and development led to the discovery and correction of a large quantity of RAX software problems. Table 8 shows the highlights of the testing on the various testbeds.

Table 8. Dates of RAX Readiness on Testbeds

Testbed	Date
UNIX	August 1997
Babybed	February 1998
Radbed	April 1998
Papabed	November 1998
Hotbench	March 1999
DS1 testbed	April 1999
DS1 spacecraft	May 1999

2.5 Ground Tests

To qualify RAX to run onboard the DS1 spacecraft, RAX underwent a rigorous program of formal tests. The tests covered nominal and off-nominal situations and exercised at all levels of fidelity available on the ground testbeds each

Remote Agent component individually, the integrated RAX product, and RAX together with the flight software.

Autonomous systems like RA pose testing challenges that go beyond those usually faced by more traditional flight software. In fact, the range of possible behaviors exhibited by an autonomous system is usually very large. This is consistent with the expectation that the system operate robustly over a large range of possible values of system parameters. However, an exhaustive verification of all situations would require an unmanageably large number of test cases.

To make matters worse, the tests should ideally be run on high-fidelity testbeds, which are heavily oversubscribed, difficult to configure correctly, and cannot run faster than real time. For example, in RAX the team could run only 10 tests in four weeks on the DS1 *Hotbench*. To cope with these time and resource limitations, the team employed a “baseline testing” approach to reduce the number of tests. Moreover, the team exploited whenever possible the lower-fidelity testbeds to validate system behaviors for which there was high confidence that the test results would extend to higher-fidelity situations. The high-fidelity testbeds were used mostly in nominal situations and under stress conditions requiring RAX to guarantee spacecraft safety.

The baseline scenario was the scenario that was expected to execute in flight initially the 6-day and 12-hour scenarios and subsequently the 2-day scenario. The team tested a number of nominal and off-nominal variations around these scenarios. These covered variations in spacecraft behavior that might be seen during execution and changes to the scenario that might be made prior to execution. Changes included variations to the goals in the mission profile, variations in when faults might occur, and variations in the FSW responses.

The architecture of RA allowed the team to run certain tests on lower-fidelity testbeds and be confident results would hold on higher-fidelity testbeds. Specifically, RA commands and monitors the spacecraft through well-defined interfaces with FSW. Those interfaces were the same on all platforms, as were the range of possible responses. Only the fidelity of the responses improved with platform fidelity. This let the team exercise a wide range of nominal and off-nominal behaviors on the *Babybeds* and the *Radbed*, test the most likely off-nominal scenarios on the *Papabed*, and test only the nominal scenarios and certain performance- and timing-related tests on the *Hotbench* and on the *DS1 Testbed*. Functional testing RA’s PS component was a special case because it required extensive use of the *UNIX* testbeds.

The rest of this section describes the tests on each testbed.

2.5.1 UNIX—The PS team made extensive use of the *UNIX* testbed for PS unit testing throughout the formal testing process. Use of the *UNIX* testbed was critical since PS is a computationally intensive task and could not take advantage of time warping. Both in nominal- and fault-response situations, PS essentially operated as a batch process with practically no reliance on the underlying real-time system (e.g., timing services). This let the team repeatedly run a batch of 269 tests with several variations of initial states, goals of the planner, and model parameters (e.g., possible turn durations). Tests were repeated for each release of the RA software, providing a certain measure of regression testing for the PS software.

2.5.2 Babybed and Radbed Testing—Each of the RA modules devised a test suite of nominal and off-nominal scenarios that isolated and exercised key behaviors in each module. For PS, this involved a batch of 54 tests comprising some of the tests in the *UNIX* batch plus tests devised to test system-level responses of PS (e.g., response to invalid initial states or to an asynchronous kill message sent by EXEC). The repetition of the tests from *UNIX* both validated the complete functional equivalence of PS between *UNIX* and PPC and verified the acceptability of PS performance on the real-time architecture. MIR was exercised on a batch of 110 tests covering the likeliest failure contexts. The PS and MIR tests were used for testing EXEC. A suite of twenty additional scenarios exercised the system-level interaction of all modules. These tests were run rapidly on the *Babybeds* and *Radbed* with time warping. Running a scenario was a time-consuming and error-prone process. To alleviate this, an automated testing tool was designed that accepted an encoded scenario description as input, controlled the simulator and ground tools to execute the scenario, stopped the test when appropriate by monitoring the telemetry stream, and stored all logs and downlinked files for later examination. This rapid data collection led to a total running time of about one week for all tests, since tests could be scheduled overnight and required no monitoring. Analyzing the results of the tests, however, was still a time-consuming process. These tests were run after each major RAX-software release.

2.5.3 Papabed Testing—*Papabed* was extensively used during development in order to integrate RAX with the DS1 flight software. In the context of the formal testing process, *Papabed* was used only to run six off-nominal system test scenarios on the “frozen” version of the RAX delivered to flight software for the flight experiment. These off-nominal scenarios corresponded to the situations that were most likely or had the potential for highest impact on the outcome of the experiment. No bugs were detected in these scenarios, probably because RA responses to off-nominal situations were well tested on the *Babybed*.

2.5.4 Hotbench and DS1 Testbed Testing—The *Hotbench* and *DS1 Testbed* were reserved for testing the nominal scenarios and for testing a handful of requirements for spacecraft health and safety. RAX was designed with a “safety net” that allowed it to be completely disabled with a single command sent either by the ground or by onboard FSW fault protection. Hence, the only ways in which RAX could affect spacecraft health and safety was by consuming excessive resources (memory, downlink bandwidth, and CPU) or by issuing improper commands. The resource-consumption cases were tested by causing RAX to execute a Lisp script that consumed those resources. The team guarded against improper commands by having subsystem engineers review the execution traces of the nominal scenarios and doing automated flight-rule checking. The nominal scenarios were run in conditions that were as close to flight-like as possible.

2.5.5 Software Change Control—As the date of the flight experiment drew closer, our perspective on testing changed. Throughout 1998, the main goal of testing was to discover bugs in order to fix them in the code. Starting in January 1999, the discovery of a bug did not automatically imply a code change to fix it. Instead, every new problem was reported to a Change Control Board (CCB) composed by senior RAX-project members. Every bug and proposed fix was presented in detail, including the specific lines of code that needed to change. After carefully weighing the pros and cons of making the change, the board voted on whether or not to allow the fix. Closer to flight, DS1 instituted its own CCB to review RAX changes.

As time progressed, the CCB became increasingly conservative and the bias against code modifications significantly increased. This is demonstrated by the following figures. In total, 66 change requests were submitted to the RAX CCB. Of these, 18 were rejected amounting to a 27%-rejection rate. The rejection rate steadily increased as time passed: 8 of the last 20 and 6 of the last 10 submitted changes were rejected.

The reason for this increase in conservatism is easily explained. Every bug fix modifies a system that has already gone through several rounds of testing. To ensure that the bug fix has no unexpected repercussions, the modified system would need to undergo thorough testing. This is time consuming, especially on the higher-fidelity testbeds; therefore, full re-validation became increasingly infeasible as flight approached. Therefore, the CCB faced a clear choice between flying a modified RAX with little empirical evidence of its overall soundness or flying the unmodified code and trying to prevent the bug from being exercised in flight by appropriately restricting the scenario and other input parameters. Often, the answer was to forego the change.

2.5.6 Summary of Testing Resources—About 269 functional tests for PS were conducted on UNIX (repeated for 6 software releases), more than 300-Babybed tests (repeated for 6 software releases), 10-Papabed tests (run once), 10-Hotbench tests (repeated for two releases), and 2 DSI-Testbed tests (on the final release) over a period of 6 months with four half-time engineers. This figure includes design, execution, analysis of the test cases, and development of testing tools.

2.6 Ground Tools

To provide adequate coverage and visibility into the RA’s onboard workings, a ground tools suite was designed to interface with the real-time RA-generated telemetry.

The two major goals of the RA ground tools were:

- To present a summary of the spacecraft status understood easily by the mission operations team.
- To present enough information about the inner workings of the RA software for the experiment team to quickly recognize and debug problems.

To support these goals, telemetry specific to Remote Agent was downlinked during the test. The RA-specific telemetry included contained:

- Planning events (e.g., planning started, finished, and progress messages).
- Sequence execution events (e.g., plan p is starting execution or plan step x started executing at time t1).
- Mode-interpretation events (e.g., valve v changed state from nominal to possibly-stuck-closed).
- Messages between RA components and between RA and other DS1 flight software.

This telemetry was event-oriented and largely incremental (i.e., interpretation of one message depended upon receipt of earlier messages). This design created some problems, as will be detailed later. In addition, there were heartbeat (“I’m still alive”) messages from RA components visible in the normal DS1 telemetry. The next sections describe specific ground tools constructed to examine this telemetry.

2.6.1 PacketView—Remote Agent telemetry is event-based. PacketView displays these events in a simple one line per message format that is easily understandable to all members of the team. As such, PacketView was the most fundamental of the RA ground tools. The PacketView interface (see Figure 10) displayed four types of telemetry packet for each of the three RA modules and IPCO_RAXO onboard messages sent between the three RA components. These were displayed as color-coded scrolling lines of text. Missing packets, detected by gaps in the sequential packet numbering, were displayed as a single line in a unique color. Search capabilities were available based on message type and content.

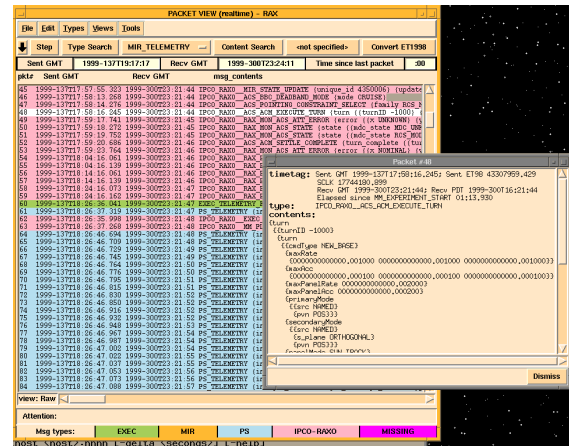


Figure 10. Packetview—Telemetry Packet Display

User-selectable dialogs presented “pretty printed” versions of the single-line packet entries. The “time bar” displayed the most recent “spacecraft sent” Greenwich Mean Time (GMT), the most recent “ground received” GMT time, and a running time since the last packet was received.

2.6.2 ExecView—ExecView visualized the execution status of plans onboard the spacecraft (see Figure 11). Different kinds of activity associated with different parts of the spacecraft appeared on separate timelines. For each kind of activity, specific planned events were represented as tokens appearing along the timeline. Tokens were color-coded to represent their execution status (i.e., in-the-future, currently-executing, completed, and completion-overdue).

As the plan was being executed by EXEC onboard the spacecraft, the start and finish times of the activities would be expected to change. Through the constraints, these changes would impact later activities. ExecView would propagate these changes downstream in the schedule, using the same propagation techniques used by the Planner.

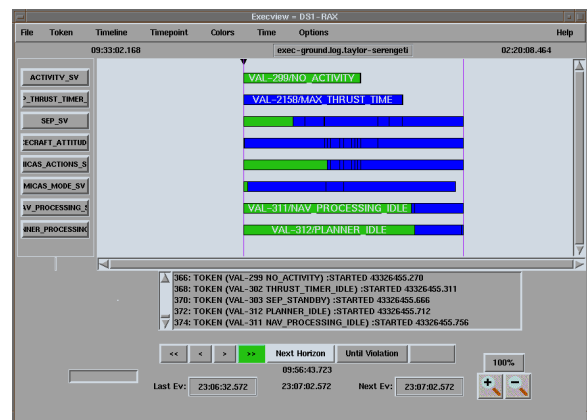


Figure 11. ExecView—Plan Execution Status

ExecView was designed initially as a debugging tool for validating EXEC development. As a result, it did not have support for handling missing telemetry packets during flight. As a result, it produced some erroneous conclusions during RAX concerning the state of plan execution. To make ExecView more useful, it will have to handle such missing data.

2.6.2 The Ground Planner—Of the three technology modules flown as part of RA, the spacecraft team was least comfortable with PS. To allow the DS1 team to gain confidence in the onboard planner, the RAX team used a ground twin of the planner. The ground planner was identical to the one onboard and was able to duplicate the onboard twin by tapping into the real-time telemetry available. It had access to other flight software resources via connection to the *Papabed*. This testbed accurately replicated the software onboard DS1 (although it did not replicate the hardware). Of particular importance to the planner were navigation-module and beacon-asteroid files describing targets for optical navigation and the portion of ACS that predicted the time required to change spacecraft orientation.

The ground planner was a useful tool in predicting the performance of the planner onboard and was especially useful as a confidence builder for mission staff unfamiliar with the working of an autonomous-planning agent.

2.6.3 PS Graph—PS Graph (see Figure 12) displayed the problem-solving trajectory taken by PS for each of the plans generated by the onboard planner. This took the form of an X-Y graph representing the search depth vs. number of search nodes visited for each successive step of the planner’s search. The purpose of these plots was to provide a quick summary of the PS problem-solving process. For example, a trajectory that visits the same depth level several times while the search-node number increases indicates that the planner is backtracking. The persistence of this situation for a large number of steps is an indication that PS may be thrashing and that it will be unlikely to return a solution within the allotted amount of time. Another use of the PS-Graph plots is to compare telemetry-data trajectories generated during simulation runs of the ground planner twin.

Although very simple, the power of this tool’s summarization and the insight level that it can provide during both RA development and operations in a stressful situation was surprising. As will be discussed in the flight experiment section below, PS Graph allowed the team to monitor an unexpected situation with PS and quickly identify the likely problem’s cause. In the future, it is advisable to design several simple visualizations like PS Graph for the reduced ground team to support an autonomy mission.




Figure 12. PS Graph—Planner Progress Display

2.6.4 Stanley and MIR—A version of MIR was also run on the ground. The purpose of this was to infer MIR’s full internal representation of the spacecraft state from the telemetry that contained a much smaller subset. Specifically, it contained the set of independent variables in MIR’s spacecraft model. The Stanley ground tool displayed a hierarchical schematic of the spacecraft’s onboard components whose status was driven by the ground MIR (see Figure 13).

Components could be opened, to show more detail, or closed. The states displayed were blue (“ok - powered off”), green (“ok - powered on”), yellow (“recoverable failure”), purple (“degraded failure”), and red (“permanent failure”). Since Stanley assigned colors to all states, nominal and off-nominal, the user can tell at a glance the conditions of the devices. Stanley did not address the issue of displaying continuous values, such as a battery state-of-charge.




Figure 13. Stanley—Hardware Status Display

In addition to the color changes, detected component faults were reported by popping up an alert box. The alert box allowed the user to click on an entry, resulting in the schematic being opened down to the appropriate hierarchical level to show the local context of the fault. Histories of all state changes, important or not, were available at any time by clicking on components.

2.6.5 Predicted Events—In flying an autonomous agent, like RA, ground operators observing the spacecraft state via its telemetry may not be in a position to know precisely when certain events are to take place. It was nevertheless important to have a prediction of when RA planned to take various actions so that the appropriate subsystem stations at the mission operations center could be staffed for observability. Therefore, the team generated a Predicted Events File (PEF), which reported both the low-level commands RA would issue together with the high-level actions RA was asked to take.

2.6.6 Public Outreach via the Web—E-mailed summaries of events onboard presented in simple English and a Java applet timeline display on the Web, patterned after ExecView, were two additional tools to present RA’s progress to the public. These tools are interesting because they required an even higher target for simplicity and understandability than did the flight controllers’ tools.

Several recent missions have used pagers and e-mail to deliver notifications to the mission-operations team. The DS1 ground system, for instance, alerted operators by pager when a given measurement strayed outside a preset range or when fault-protection telemetry went into an unusual state. This was taken a step further in RAX by producing descriptions of important events in common English. The summarized descriptions were automatically posted to the RAX Web site (<http://rax.arc.nasa.gov>) and e-mailed in batches to a public mailing list. Two thousand subscribers received this e-mail during RAX. Terse descriptions were also sent to team members’ alphanumeric pagers via e-mail.

An alternative Remote Agent-activity description (Figure 14) was also provided using horizontal timelines patterned after ExecView. This was implemented as a Java applet. The timelines in the top window represented major kinds of activity (e.g., attitude or camera-related activity). Along the timelines were tokens indicating particular activities (e.g., a turn), in effect displaying the plans generated onboard on a user’s Web browser. Also included were controls to step through the timelines and an event-based summary similar to that provided in e-mail. The most interesting feature of this applet was its ability to show what RA planned to do at any time. The user could click on any event, and the applet would show what the RA planned to do at that time. This is interesting because the plan changed several times due to simulated faults. Thus the applet provided an historical

overview of RA’s re-planning activity and recreated conditions aboard the spacecraft for the general public.

Due to time pressure, the outreach tools were designed to handle the nominal scenario only (including the simulated faults). They did not accurately reflect the RAX software problems that occurred. They did, however, summarize activity during the new scenario without modification. These summaries are still available at the RAX web site: <http://rax.arc.nasa.gov>.

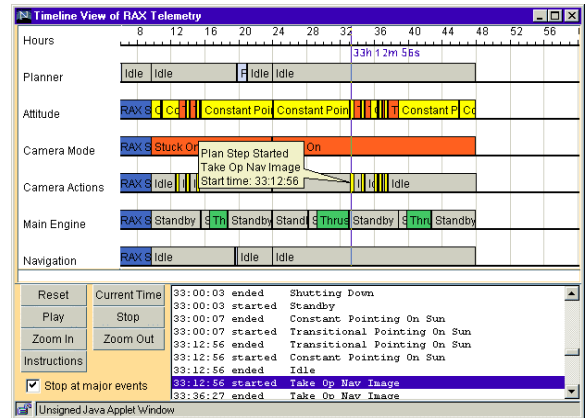


Figure 14. Timeline Applet

Additional details on the RAX Ground Tools can be found in [13].

2.7 Flight Test

RAX was scheduled to be performed on DS1 during a three-week period starting May 10, 1999. This period included time to retry the experiment in case of unexpected contingencies. On May 6, 1999, DS1 encountered an anomaly that led to spacecraft safing. Complete recovery from this anomaly took about a week of work by the DS1 team, both delaying the start of RAX as well as taking time away from their preparation for the asteroid encounter in July 1999. In order not to jeopardize the encounter, the DS1 project also decided to reclaim the third RAX week for encounter preparation, leaving only the week of May 17, 1999, for RAX. However, to maximize the time to try the more important 2-day experiment, they agreed to go ahead with the 2-day experiment without first doing the confidence-building 12-hour experiment. This decision was strong evidence that the DS1 project had already developed significant confidence in RAX during pre-flight testing.

2.7.1 Flight Test Part 1—The flight experiment started on Monday, May 17, 1999. At 11:04 am PDT, a telemetry packet was received confirming that the 2-day scenario had begun on DS1. Shortly thereafter, PS started generating the first plan. The first plan was generated correctly, but not before an unexpected circumstance created some apprehension among team members.

Figure 12 graphically shows the situation with the output of the PSGraph ground tool. The blue trajectory relates to a *Papabed* test that was run May 16, 1999 under identical condition to those of the flight test. The green trajectory describes what happened during flight. The deviation in the green trajectory from the 45° diagonal trajectory means that PS in flight backtracked significantly more than on *Papabed*. Since the conditions on the spacecraft were believed to be practically identical to those on the ground testbeds, there was no apparent reason for this discrepancy. Subsequently, the cause of this discrepancy was traced back to the spacecraft and *Papabed* differing on the contents of the AutoNAV file containing asteroid goals. Therefore, in flight PS was actually solving a slightly different problem than it had solved on the ground! Thus, this unexpected circumstance demonstrated that PS problem solving was robust to last-minute changes in the planning goals, increasing the credibility of the autonomy demonstration.

The 2-day scenario continued smoothly and uneventfully with the simulated MICAS switch failure, the resulting replan, long turns to point the camera at target asteroids, optical navigation imaging during which no communication with DS1 was possible, and the start of IPS thrusting.

However, around 7:00 am on Tuesday, May 18, 1999, it became apparent that RAX had not commanded termination of IPS thrusting as expected. Although plan execution appeared to be blocked, telemetry indicated that RAX was otherwise healthy. The spacecraft too was healthy and in no apparent danger. The decision was made to use EXEC's ability to handle low-level commands to obtain more information regarding the problem. Once enough information had been gathered, the decision was made to stop the experiment. By this time, an estimated 70% of the RAX validation objectives had already been achieved.

2.7.2 Troubleshooting and Recovery—By late Tuesday afternoon, the cause of the problem was identified as a missing critical section in the plan-execution code. This created a race condition between two EXEC threads. If the wrong thread won this race, a deadlock condition would occur in which each thread was waiting for an event from the other. This is exactly what happened in flight, though it had not occurred even once in thousands of previous races on the various-ground platforms. The occurrence of this problem at the worst possible time provides strong impetus for research on formal verification of flight-critical systems. Once the problem was identified, a patch was quickly generated for possible uplink.

Following the discovery of the problem, a 6-hour RAX scenario was generated to demonstrate the remaining 30% of the RAX validation objectives. This scenario included IPS thrusting, three failure scenarios, and back-to-back planning. This new scenario was designed, implemented,

and tested, together with the patch, on *Papabed* overnight within about 10 hours. This rapid turnaround allowed the team to propose a new experiment at the DS1 project meeting Wednesday. The DS1 project decided to proceed with the new scenario. However, they decided not to uplink the patch, citing insufficient testing to build adequate confidence. In addition, based on the experience on various ground testbeds, the likelihood of the problem recurring during the 6-hour test was very low. Nonetheless, a contingency procedure was developed and tested that would enable the team to achieve most of our validation objectives even if the problem recurred.

The DS1 project's decision not to uplink the patch is not surprising. What was remarkable was their ready acceptance of the new RAX scenario. This is yet more evidence that the DS1 project had developed a high level of confidence in RA and its ability to run new mission scenarios in response to changed circumstances. Hence, although caused by an unfortunate circumstance, this rapid mission redesign provided unexpected validation for RA.

2.7.3 RAX Flight Part 2—The 6-hour scenario was activated Friday morning, May 21. The scenario ran well until it was time to start up IPS. Unfortunately, an unexpected problem occurring somewhere between FSW and RAXM caused a critical monitor value to be lost before it reached RA. The cause of this message loss has not been determined. The problem of lost-monitor values could have been avoided with periodic refreshes of the monitor values. This was deemed out of scope for the purposes of the experiment, and RA was known to be vulnerable to message loss. This vulnerability led RA's estimation of the IPS state to diverge from the true state. Fortunately, the discrepancy proved to be benign. Hence, RA was able to continue executing the rest of the scenario to achieve the rest of its validation objectives.

By executing the two flight scenarios, RAX achieved 100% of its validation objectives.

2.8 Effectiveness of the Development and Test Process

Progress in development and testing during the RAX project can be analyzed through the Problem Reports (PRs) filed between April 1997 and April 1999 (see Table 9).

A developer or a tester could file a PR, usually reporting a bug or requesting a change in the software behavior. A few PRs were reminders of activities or checks to be performed. PRs remained open until the developers addressed them. When a resolution to the report was filed (e.g., a bug fix was provided), the originator of the report would check the validity of the resolution. If accepted, the resolution was included in a formal release. A few PRs were suspended. This meant that the risk of the problem was assessed and considered acceptable within the limits of RAX.

Table 9. Number of PRs by Subsystem

Subsystem	Number of PRs
Planner/Scheduler	233
Executive	100
MIR	85
RAX Manager	22
System	77
Communication	22
Simulator	30
Others	11
Total	580

Figure 15 gives an idea of the temporal distribution of new PRs filed over the duration of the project. The last four columns (from January 1999 to April 1999) relate to problems that were submitted to the CCB process. Notice that the number of PRs in this period is still quite high (91). This depended in part on the fact that integration with flight software started in earnest in December 1999, with RAX running on *Papabed*, and that until then RA had only been operating interacting with low-fidelity simulators.

PRs can be divided into three categories:

- Modeling PRs required by domain-specific knowledge changes relative to the DS1 spacecraft subsystems.
- Engine PRs effecting RA’s core reasoning engines.
- PRs related to other mechanisms such as the format of data file exchanged between RA components. This category also includes reminders and requests of change that were outside the scope of RAX.




Figure 15. Temporal Distribution of Problem Reports

Figures 16, 17, and 18 describe the distribution of problems by category for each individual engine. The most stable RA subsystem was MIR. This stability manifested itself both in

terms of the total number of Engine and Modeling PRs filed and in terms of the very few PRs of these categories filed in the last 4 months of the project. This was due both to the maturity of the MIR technology and to the fact that the problem addressed by MIR changed very little during the duration of the project.




Figure 16. Planner PRs by Category




Figure 17. Executive PRs by Category

The command language used by EXEC, ESL, was developed prior to the RAX project and caused a negligible number of PRs. The majority of the EXEC PRs fell into the Other category and were related to integrating the PS and MIR modules. The next-largest category of PRs was model related. These tended to manifest themselves each time RA was integrated on a higher-fidelity testbed. Models for EXEC were undergoing modifications quite late (February 1999 to April 1999). This was primarily due to the fact that these months covered a period of intense activity on *Papabed* with the interfaces with the details of how flight software operated being finally communicated to the RAX team. This resulted in some localized changes in interface functions and in task-decomposition procedures. The effects

of these changes were typically localized at the EXEC level and did not propagate up to PS models. This confirms the possibility of developing RA even on the basis of an accurate but abstract characterization of the modeled system, with much of the high-level behaviors remaining stable when further details on the behavior of the system are known.




Figure 18. MIR PRs by Category

Both in the case of MIR and EXEC, testing was very effective at validating models. EXEC and MIR models have many non-interacting or loosely interacting components that can be tested independently. This reduces the number of test cases that are needed. Testing small model components independently—like the team did in RAX—should scale-up for larger, future science-mission models.

In the case of PS, a larger overall percentage of PRs (about 45%) were model related. More importantly, a large number of new problems were discovered during the last four months of the project, after the formal testing process had ended. The vast majority of these problems consisted of PS operating correctly but unable to find a plan within the allocated time limit since its search was “thrashing.” These problems were particularly serious since they could easily arise in off-nominal situations during flight.

There were several reasons for this situation:

- The ranges of some parameters turned out to be different than those assumed by PS testing: e.g., PS testing assumed turn durations were at most 20 minutes, while actual turns could take over an hour. This created stress situations not considered by formal PS testing.
- Planning problems became more challenging in the transition from the 6-day to the 2-day scenario. The temporal compression led to the disappearance of slack time between activities. In the 6-day scenario, PS could exploit this slack to achieve subgoals without backtracking. In the 2-day scenario, backtracking

became necessary, revealing additional brittleness in the PS chronological backtracking search.

- A more fundamental issue was the independence between the PS-test generator and the structural characteristics of the domain model. This led to the test generator missing a number of stress cases. For example, one problem depended upon the specific values of three continuous parameters: the time to start up the IPS engine, the time to the next optical navigation window, and the duration of the turn from the IPS attitude to the first asteroid. An equation relating these parameters can crisply characterize the stress situations. Unfortunately, the automatically-generated test cases used for PS validation only covered pair-wise interactions. Therefore, they could not reliably detect such problems.

Given the late date at which these new problems were discovered, it was not feasible to modify the test suite to test extended variations around the new baseline. Instead, only the most crucial variation was focused on: the time at which re-plans might occur. The objective was to ensure that the planner could handle any re-planning contingency. Two steps were needed to accomplish this. First, the new 2-day scenario was designed to guarantee that the harmful constraint interactions of the PS domain model would be avoided under any hypothetical replanning contingency. The idea was to ensure that PS could always return a plan within the given time limit. Second, a new PS test suite was carefully designed and run to ensure that this was indeed the case.

The design methodology for this new PS test suite is instructive. Exhaustive generation of all possible plans was clearly impossible. Instead, using PS-model knowledge, boundary times at which the topology of the plans would change were manually identified. Twenty-five such boundary times were identified and generated a total of 88 test cases corresponding to plans starting at, near, or between boundary times. This led to the discovery of two new bugs. This number of tests is more than four orders of magnitude smaller than the total of 172,800 possible re-plan times. Furthermore, analysis of the test results showed that PS would fail to find a plan at only about 0.5% of all possible start times. Although the probability of this failure was extremely low, contingency procedures were developed to ensure that the experiment could be successfully continued even if this PS failure actually occurred.

The above test-suite-design methodology was used only toward the end of RAX, after the PS model and code had been frozen. However, this (currently manual) analysis method can be generalized and extended to provide an automatic PS testing procedure throughout the development process for new application domains.

Note that the number of PRs regarding the reasoning engines of PS, EXEC, and MIR was relatively small. For example, less than 10% of PS's PRs were Engine related and the last was filed in September 1998. However, the bug EXEC encountered during RAX shows that the engine validation methodology could have improved. In fact, the testing was primarily focused on validating the knowledge in the domain models. Tests were selected to exercise the domain models. By exercising RA on these test scenarios, the domain models and engines were effectively tested as a unit. However, especially for concurrent systems such as EXEC, a much better approach is to thoroughly, formally validate the logic of the engines through the use of formal methods [16]. Although expensive, this form of testing can give a high level of quality assurance on the core of the RA technology. Moreover, since the engines remain unchanged over a large number of applications, the cost of this testing can be amortized across several missions.

2.9 Costing

Figure 19 gives an overall view of the costing of RAX starting from October 1997, when tracking information was available. The figure describes costs based on development, testing, integration, and technical management activities. The Full Time Equivalence (FTE) exerted is shown on the Y-axis. Costing by FTEs is more appropriate in this case because of the differing accounting standards used at NASA's ARC and JPL.

The chart clearly shows the distinct development, testing, and integration efforts being partitioned in time; development efforts were clearly focused before the move to the high-fidelity testbeds. While testing and integrations efforts were ongoing activities, they came to dominate the latter part of the move to the testbeds. While the overall trend is a curve with diminishing figures, there are some features that need some explanation.

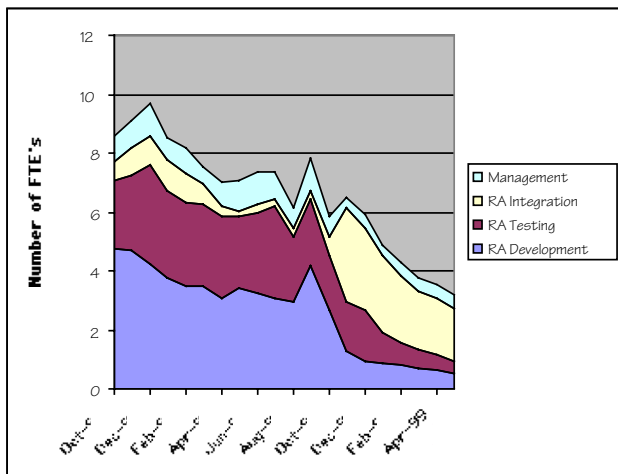


Figure 19. RAX Costing

The first peak in the October–December 1997 timeframe corresponds to the time when formal test plans were put together and UNIX testing began. In addition, RAXM was delivered to the flight team at this time. The peak, therefore, is categorized by these efforts and the resulting testing and bug fixing that took place.

The second peak in the August–October 1998 timeframe corresponds to a number of events. Primarily, this was dealing with new code deliveries to the planner engine to allow EXEC to deal more robustly with additional information in the plans. This increased effort highlights the extra individuals from outside the RA team who made these efforts possible. In addition, all team members were gearing up towards testing on the *Papabed*, the highest-fidelity testbed available at that time. Subsequent to that event, the curves show a deep decline, as expected, in the development efforts when the team focused more on integration and testing on the various testbeds available. Efforts dealing with integration, therefore, show a perceptible increase.

Lastly, the gap between the testing and integration efforts appears to be inverted in the December 1998 and May 1999 timeframe. The primary reason for this was our late arrival on the high-fidelity testbeds. This meant testbed integration had to be finished in half the normal time. It was also the case that working on these testbeds took time and effort beyond what that needed on the lower-fidelity testbeds (*Unix* and *Babybeds*) that were available early on. Configuration training and problem-detection also took substantial time and effort, causing a larger manpower effort for integration as shown.

The actual costs of the entire RA development effort was \$500K for the NewMAAP demonstration (May to November 1995), \$4.5 million during the DS1 autonomy FSW phase (December 1995 to March 1997), and \$3 million for RAX (April 1997 to June 1999), for a total cost of \$8 million.

2.10 Lessons Learned

The RA team learned valuable lessons in a number of areas including RA technology and processes, tools, and even autonomy benefits to missions.

2.10.1 Robustness of the Basic System—Model validation alone does not suffice; the rest of the system, including the underlying inference engines, the interfaces between the engines, and the ground tools, must all be robust. Given the resource constraints, the decision was made to focus our formal testing on model validation, with engine and interface testing happening as a side effect. This was a reasonable strategy: code that has been unchanged for years is likely to be very robust if it has been used with a variety of different models and scenarios. However, newer code does not come with the same quality assurance. Also, as the

deadlock bug in-flight showed, subtle-timing bugs can lay hidden for years before manifesting themselves.

Conclusion: The primary lesson is that the basic system must be thoroughly validated with a comprehensive test plan as well as formal methods, where appropriate, prior to model development and flight insertion. Interfaces between systems must be clean and well specified, with automatic code generation used to generate actual interface code, telemetry, model interfaces, and test cases; code generation proved enormously helpful in those cases where it was used.

2.10.2 Robustness of Model-Based Design—As mission development times become shorter and mission objectives more ambitious, it is less and less likely that an accurate model of each spacecraft component will be available early in the flight- and ground-software development cycle. Dealing with this uncertainty is a major problem facing future missions. By emphasizing qualitative and high-level models of behavior RA can help solve this dilemma. Qualitative, high-level models can be captured early in the mission lifetime and should need only minor adjustments when the hardware is better understood. Our experience on RAX essentially confirms this hypothesis. Initial spacecraft models used by PS, EXEC, and MIR were built early in the DS1 mission, before April 1997. During the following year and a half, EXEC and MIR models did not change and the PS model was only changed in order to support more efficient problem solving by the search engine, not in order to reflect new knowledge of the spacecraft behavior. In the last phase of the experiment preparation, when communications between the RAX team and the DS1 team resumed, adjustments were needed to finalize the interface between the low-level EXEC primitives and flight software.

Conclusion: Contrary to much concern, the type of qualitative, high-level models used by RA requires little tuning throughout the project. The usefulness of the models for software development has been validated.

2.10.3 Model Design, Development and Test—One of the biggest challenges was model validation. This was particularly true during validation testing, when even small changes in the models had to be carefully and laboriously analyzed and tested to ensure that there were no unexpected problems. In fact, in some cases it was decided to forego a model change and, instead, decided to institute flight rules that would preclude the situation that required the model change from arising. A related issue was that methods do not yet exist to characterize RA's expected behavior in novel situations. This made it difficult to precisely specify the boundaries within which RAX was guaranteed to act correctly. While the declarative nature of RA models was certainly very helpful in ensuring the correctness of models and model changes, the difficulty stemmed from unexpected

interactions between different parts of the model (e.g., different parts of the model may have been built under different, implicit, conflicting assumptions).

Conclusion: The central lesson learned here is the need for better model-validation tools. For example, the automated test running capability that was developed proved to be enormously helpful, as it allowed the team to quickly evaluate a large number of off-nominal scenarios. However, scenario generation and evaluation of test results were time consuming. In some cases, the laborious process followed to validate model changes has provided the team with concrete ideas for developing tools that would dramatically simplify certain aspects of model validation. Preliminary work in the area of formal methods for model validation is also very promising. Finally, there is a need to develop better methods for characterizing RA's behavior with a specific set of models, both as a way of validating those models and as a way of explaining the models to a flight team.

2.10.4 Onboard Planning—Since the beginning of RA, onboard planning has been the autonomy technology that most challenges the comfort level of mission operators. Commanding a spacecraft with high-level goals and letting it autonomously take detailed actions is very far from the traditional commanding approach with fixed-time sequences of low-level commands. During RAX, the flawless demonstration of onboard planning has provided powerful feasibility-of-approach proof. Discomfort with the discrepancy between tested behavior and in-flight PS behavior during RAX was a surprising mirror of the autonomy critics' objections.

Conclusion: It is difficult to move past the mindset of expecting complete predictability from the behavior of an autonomous system. However, RAX has demonstrated that the paradigm shift is indeed possible. In the case of PS behavior during RAX, the point is not that the combination of pictures requested by NAV had never been experienced before, but that the problem-solving behavior that the planner used to achieve each individual picture goal had indeed been tested. Confidence in complex autonomous behaviors can be built up from confidence in each individual component behavior.

2.10.5 Design for Testability—System-level testing is an essential step in flight preparation. Designing RA to simplify and streamline system-level testing and analysis can enable more extensive testing, thus improving robustness. In RAX, system-level testing proved to be cumbersome. The primary reason for this was the absence of efficient tools to generate new mission scenarios; therefore, all system tests had to be variations on the nominal scenarios. Hence, to test a particular variation, one was forced to run a nominal scenario up to the point of the

variation: e.g., testing thruster failures during turns required at least 6 hours, since the first turn occurred about 6 hours into the scenario.

Conclusion: The difficulty of generating new mission scenarios is easily addressed: a graphical tool allowing visual inspection and modification of mission profiles, as well as constraint checking to ensure consistency, can dramatically simplify the construction of new mission profiles. Such a tool is now being constructed. Nonetheless, overall RA validation is still necessary to ensure that RA will properly handle each new mission profile (see below).

2.10.6 Systems Engineering Tools—Coding the domain models required substantial knowledge acquisition, which is a common bottleneck in Artificial Intelligence systems. It is better to have the domain expert code the models directly.

Conclusion: Develop tools and simplify the modeling languages to enable spacecraft experts to encode models themselves. Employ tools and languages already familiar to the experts. Organize the models around the domain (attitude control, power, etc.) rather than around the RA technology (PS, EXEC, MIR).

2.10.7 Mission Profile Development—RA is commanded by goals specified in a mission profile. For the experiment, constructing the profile was a “black art” that only one or two people on the RA team could perform. The mission planners and operations personnel must be able to specify goals themselves.

Conclusion: Simplify the specification of goals. When possible, use approaches already familiar to mission planner, such as graphical-timeline displays and time-ordered listings. Provide automated consistency checking.

2.10.8 Adaptability to Late-Model Changes—The spacecraft requirements and operating procedures changed throughout development and after launch. It was not possible to encode late changes, due to the regression-testing overhead that each change required.

Conclusion: The validation cost of model changes must be reduced. Some possibilities include tools to evaluate the consequences of model changes on testing. The models already support localized changes. Procedures are needed to uplink and install just those changes.

2.10.9 Ground Tools—Ground tools ought to be developed well in advance of the actual flight and be used as a primary means to test and understand how to operate complex systems. Given the late date of development of most of the ground tools, a good many of them felt not well integrated.

As a result, only the tools displaying or interpreting data in the most obvious way were of high value.

2.10.10 Telemetry—In addition to an onboard textual log file downlinked at the end of the experiment or on request, RAX sent a stream of binary-telemetry packets, one for each significant event, that were displayed as color-coded text on the ground. Among other things, the telemetry let the team monitor all onboard communication among RAX modules and between RAX and FSW. This proved valuable in letting the team quickly diagnose the anomalies that occurred. It was immediately evident the reason RAX failed to turn off the ion engine was it had stopped executing the plan in some unanticipated manner; the team knew RAX was still running and could also rule out a plan abort or a failure to send just one command. Similarly, the upshot of the second anomaly was immediately narrowed down to a monitor message, which was either not sent or not received.

Conclusion: Ensuring sufficient visibility on all platforms, including in-flight, requires adequate information in telemetry. The best way to ensure this is to design the telemetry early and to use it as the primary, if not the only, way of debugging and understanding the behavior of the system during integration, test, and operations.

2.10.11 Team Structure for RA-Model Development—The RAX team was structured horizontally along engine boundaries. This meant that team members specialized in one of the PS, EXEC, and MIR engines and that each team was responsible for modeling all spacecraft subsystems for their engine. This horizontal organization was appropriate for RAX, since it was our first major experience in modeling spacecraft subsystems for flight. Hence, it made sense for engine experts to do all modeling for their engine. However, this organization has several shortcomings. Perhaps the most significant shortcoming was that knowledge of any one spacecraft subsystem (e.g., attitude control, ion propulsion, MICAS camera) was distributed across the three teams; one needed discussions with three individuals to get a complete understanding of how a subsystem was commanded by RA.

Conclusion: These shortcomings suggest an alternate structuring for a future SW team. Instead of a horizontal structure, teams might be organized vertically along spacecraft subsystem or domain-unit boundaries (e.g., a single team would be responsible for developing all models for ACS). This would ensure internal coherence of the resulting model. Furthermore, since modelers would need to understand how to use all three engines, they can make effective decisions on how best to model a subsystem to exploit the strengths of each engine and avoid information duplication.

While a vertical-team organization has its benefits, certain aspects of model development intrinsically involve managing and reasoning about global constraints: e.g., power allocation strategies, system-level fault protection. Hence, it is important to involve systems engineers to develop these global strategies.

2.11 Answers to a Project Manager's Questions

In August 1997, after a meeting between the RA team and DS1's project management, David Lehman, DS1 project manager, was asked how the RA team could convince a future science mission's project manager to use RA.

Lehman responded with a series of questions that a project manager would ask if he or she was just starting a new science-mission development.

Answering these questions after RAX would be a good way to summarize our current understanding of the technology. Also, the reader should keep in mind that these answers apply as well to other software frameworks comparable to RA in functionality and approach.

What does RA do to make my life easier?

It makes life easier because:

- It is possible to operate with a high level of autonomy during more phases of a mission outside of critical sequences.
- It provides a framework that facilitates the translation of system engineering requirements into operational code during the development phase of a mission.
- The RAX experience shows that RA can indeed operate autonomously and respond robustly to likely anomalies without intervention from a ground team and the associated delay due to round trip light time, diagnosing the problem, creating a command sequence, and validating it. This can translate into lower-operational costs and improved science.
- RA can reduce the need for communication with ground. This means less time on the highly subscribed DSN and further cost reductions.

Is RA a new technology?

RA is a novel integration of three technologies; their application to spacecraft is also new. Each of the component technologies in RA is an AI technology with a long history. Theoretical papers exist that demonstrate strong formal properties of some RA components [8][10]. Significant applications exist for each of the technology components. The most significant risk that was addressed by the overall RA development was the integration of the three technologies into a highly autonomous agent. The team believes that RAX demonstrates that successful integration.

Why is RA the best thing to do in order to make the spacecraft have autonomous operations?

Other systems exist that are comparable with some subset of the capabilities provided by RA; however, the other systems typically do not integrate all aspects of RA. For example, the team is not aware of any operational software for autonomous agents that contains an onboard planning and scheduling system.

One of the problems with FSW is that the FSW team is at the end of the "requirements food chain." Late requirements to FSW in turn results in increased costs and wasted efforts in the beginning of the project. With RA, more stuff must be put into the code, like the models of the hardware and how users want the spacecraft to operate. Therefore, this requirement should further exacerbate the standard FSW problem of the past. How can that be fixed?

Indeed, coding RA models requires a substantial up-front system engineering effort. The advantage of the declarative approach is that the impact of late model changes is lower compared to conventional flight software. Because of their abstract nature, the vast majority of RA models remained completely valid and operational throughout the project.

FSW is hard to test. FSW + RA should be even harder to test? How is that fixed?

The RAX experience confirms that testing FSW is hard. The bug that was found during flight shows that more attention and effort needs to be spent validating the basic engines. The validation cost is well worth the effort because the engines are components that can be re-used over many missions.

With respect to the capabilities provided by the domain models, our experience shows that testing of RA can be successfully layered. Testing RA can be separated from testing FSW. Also, internal to RA, different capability levels of abstraction can be separated, taking advantage of the each layer's different requirements. For example, low-level, real-time capabilities require testing on the slower real-time testbeds, while the higher-level functionalities can undergo extensive testing on readily-available, cheaper workstations.

With respect to coverage of possible RA behaviors, the experience is that the larger the space of possible combination of parameters and the higher the number of possible interactions between subsystems, the harder it is to guarantee that RA will work nominally under all circumstances. This is not surprising. Restricting harmful interaction by design is the standard problem that needs to be addressed by system and fault-protection engineers in a mission. RA does not make the problem go away. However,

during development RA can be used in simulation to provide a useful tool to explore system behavior under stress situations. This could help detect and fix potential problems with constraint interactions that are difficult to identify otherwise.

Mission operation is a big deal. However, most projects do not think about it until late in the development. Does RA offer any benefits here?

RA operates a spacecraft by generating plans that meet goals and flight rules. The development of goals and flight rules is intrinsic to RA development and forces issues to be worked hand in hand with flight software. The result is a tighter integration between mission operations and flight software, which is a good thing.

What parts of the operations phase is RA best suited for? Normal cruise phase when nothing is happening, flying around something when DSN is not in sight? Is RA needed during the whole mission, or only during some critical mission phases, like an orbit insertion?

In principle, RA can support all phases of a mission. This does not mean that all of its component technologies are suited for all phases. For example, the performance of the current implementation of PS³ makes it unsuitable for closed-loop use with tight response times (seconds to a few minutes). However, PS could be very valuable for scheduling competing observation of different levels of for a long-term, observatory mission. In these situations, the conditions are more similar than those demonstrated in RAX, where the next plan can be generated while the current one is executing.

RAX demonstrated that RA is viable during mission-cruise phase. Although this was done on a reduced model of the spacecraft, the team believes that the scaling up factors in this case should be linear and within current RA-technology's reach. With respect to the potential use during a critical phase, EXEC's event-driven, conditional execution and MIR's model-based fault protection are best suited for onboard use. Also, even within its current performance characteristics, PS could be useful during the design phase of the scripts to be executed by RA. RAX gives some evidence that this is possible. However, the ultimate demonstration of these capabilities will require more work.

3.0 FUTURE APPLICATIONS

Future work regarding Remote Agent can be divided into three categories: fundamental improvements in the

³ The RAX plans consisted of 15–25 executable activities, 50–80 tokens and 90–134 constraints. They took 50–90 minutes to generate using about 25% of the RAD6000 CPU.

capabilities of its components, improvements in usability or deployability, and upcoming demonstrations or applications. Since the experiment, a significant effort has gone into basic research to improve future iterations of Remote Agent. For example, a more capable version of Livingstone has been developed that better handles ambiguity when tracking the state of the spacecraft. Livingstone now tracks a number of most-likely states the spacecraft could be in, given the observations it has received thus far. If new observations invalidate the possible states MIR considered most likely, it re-analyzes the commands that have been given and the possible failures in order to determine which previously unlikely states now explain the unexpected observations.

PS has a number of efforts underway to improve the underlying software implementation: it now has a new modular-software architecture that allows plugging various search techniques into the engine. Work is underway in model analysis that will allow early detection of domain-model inconsistencies. Analysis of static models is also being undertaken to automatically generate search-control instrumentation. The latter approaches will allow rapid prototype development of planner models by non-technologists using incremental model development via “what-if” analysis to vastly reduce development costs. It will also provide mission staff with a better understanding of how autonomy architectures will fit into the overall design of FSW.

Other efforts are also in place to redesign the system architecture to allow EXEC access to the planner temporal database and algorithms. A unified-modeling language is being developed with cleaner semantics to allow EXEC to respond to exogenous events more rapidly.

The architectural themes pioneered in Remote Agent are gaining more general acceptance in the flight-software and mission-operations communities [15]. Applying RA to the DS1 spacecraft provided a wealth of practical lessons about what was needed to create a sustainable autonomy-engineering process and make this technology usable for main-line mission development and operations. PS and MIR have been re-architected, modularized, and implemented in C++ rather than Lisp. These next-generation versions are in alpha testing at the date of this report. EXEC is expected to be re-architected and implemented in C by the end of calendar year 2000. RA team is now developing tools for graphically creating and debugging models, for automating much of the integration of RA with traditional flight software, and for allowing humans and autonomous software to interact more easily. The team is collaborating with software-verification researchers at NASA's ARC and at Carnegie-Mellon University to allow certain Remote Agent models to be analyzed to prove they cannot recommend undesired behavior. In short, these research and development efforts are designed to make RA and similar

technologies more capable, easier to use, and easier to test and validate.

RA technology is successfully being transferred beyond the original team, and several groups are currently building prototypes with RA to evaluate it. At NASA's Kennedy Space Center, Remote Agent applications are being developed to evaluate RA for missions involving in-situ-propellant production on the Mars 2003 lander or a future piloted mission. Applications for shuttle operations are being pursued as well.

At the Jet Propulsion Laboratory, RA is being evaluated as the baseline-autonomy architecture for the Origins Program interferometry instruments and is being used in the JPL-interferometry testbed. One early customer of this development may be New Millennium Program's Deep Space Three, a space-based interferometry mission that includes two or three spacecraft cooperating to make science observations. At Johnson Space Center, components of Remote Agent are being integrated into an ecological life-support testbed for human missions beyond Earth orbit. At Ames, Remote Agent technology is being incorporated into software for more robustly controlling planetary rovers. Along with Orbital Sciences Corporation, Ames is working to demonstrate Remote Agent as it applies to streamlining the checkout and operation of a reusable launch vehicle. This demonstration will fly on the X-34 vehicle. In collaboration with Boeing, a similar experiment will be flown on the X-37 vehicle.

4.0 ACKNOWLEDGMENTS

This report describes work performed at NASA's Ames Research Center and at the Jet Propulsion Laboratory, California Institute of Technology, under contract to the National Aeronautics and Space Administration.

The Remote Agent Experiment would not have been possible without the efforts of the DS1 flight team and Harlequin Inc. In addition, the direct contribution of the following individuals is gratefully acknowledged:

Steve Chien, Micah Clark, Scott Davis, Julia Dunphy, Chuck Fry, Erann Gat, Ari Jonsson, Ron Keesing, Guy Man, Sunil Mohan, Paul Morris, Barney Pell, Chris Plaunt, Greg Rabideau, Scott Sawyer, Reid Simmons, Mike Wagner, Brian Williams, Greg Whelan, and David Yan.

In addition, the following individuals were instrumental in providing valuable support and encouragement throughout the duration of the development and flight of the Remote Agent:

Abdullah Aljabri, Martha Del Alto, Ralph Basilio, Magdy Bareh, Kane Casani, Rich Doyle, Dan Dvorak, Dan Eldred,

Julio Fernandez, Peter Gluck, Jack Hansen, Ricardo Hassan, Lorraine Fesq, Ken Ford, Sandy Krasner, David Lehman, Frank Leang, Mike Lowry, Nicole Masjedizadeh, Maurine Miller, Alex Moncada, Mel Montemerlo, Peter Norvig, Keyur Patel, Bob Rasmussen, Marc Rayman, Mark Shirley, Martin Simmons, Helen Stewart, Gregg Swietek, Hans Thomas, Phil Varghese, and Udo Wehmeier.

The team also gratefully acknowledges Caelum Research Corp. and Recom Technologies' participation and support.

5.0 LIST OF REFERENCES

- [1] E. Gat, "ESL: A language for supporting robust plan execution in embedded autonomous agents," in *Proc. 1997 IEEE Aerospace Conference*, 1997, pp. 319–324.
- [2] E. Gat and B. Pell, "Abstract resource management in an unconstrained plan execution system," in *Proc. 1998 IEEE Aerospace Conference*, 1998 [CD-ROM].
- [3] B. Pell, E. Gat, R. Keesing, N. Muscettola, and B. Smith, "Robust periodic planning and execution for autonomous spacecraft," in *Proc. IJCAI-97*, 1997, pp. 1234–1239.
- [4] B. C. Williams and P. P. Nayak, "A model-based approach to reactive self-configuring systems," in *Proc. AAAI-96*, 1996, pp. 971–978.
- [5] N. Muscettola, "HSTS: Integrating planning and scheduling," in *Intelligent Scheduling*, M. Fox and M. Zweben, Eds. San Francisco: Morgan Kaufman, 1994, pp. 169–212.
- [6] N. Muscettola, B. Smith, S. Chien, C. Fry, G. Rabideau, K. Rajan, and D. Yan, "Onboard planning for autonomous spacecraft," in *Proc. Fourth International Symposium on Artificial Intelligence, Robotics and Automation for Space (iSAIRAS-97)*, 1997, pp. 229–234.
- [7] Muscettola N., P. P. Nayak, B. Pell, and B. C. Williams, "Remote Agent: To boldly go where no AI system has gone before," *Artificial Intelligence*, vol. 103, nos. 1–2, pp. 5–48, Aug. 1998.
- [8] J. de Kleer and B. C. Williams, "Diagnosing multiple faults," *Artificial Intelligence*, vol. 32, no. 1, pp. 97–130, 1987.
- [9] D. S. Weld and J. de Kleer, *Readings in Qualitative Reasoning about Physical Systems*. San Francisco, CA: Morgan Kaufmann, 1990.
- [10] A. Jonsson, P. Morris, N. Muscettola, K. Rajan, and B. Smith, "Planning in interplanetary space: Theory and practice," *Proc. Fifth International AI Planning Systems (AIPS 2000)*, 2000, pp. 177–186.
- [11] B. Smith, K. Rajan, and N. Muscettola, "Knowledge acquisition for the onboard planner of an autonomous spacecraft," in *Knowledge Acquisition, Modeling and Management*, E. Plaza and R. Benjamins, Eds. New York, NY: Springer, 1997, pp. 253–268.

- [12] N. Masjedizadeh, “Remote Agent: 1999 Co-Winner of NASA's Software of the Year Award” [Online document], 1999 May 16, [cited 2000 Jul. 6], Available HTTP: <http://rax.arc.nasa.gov>
- [13] K. Rajan, M. Shirley, W. Taylor and B. Kanefsky, “Ground tools for autonomy in the 21st century, to appear in *Proc. IEEE Aerospace Conference*, 2000.
- [14] D. Smith, J. Frank, and A. Jonsson, “Bridging the gap between planning and scheduling,” to appear in *Knowledge Engineering Review*, 2000.
- [15] D. Dvorak, R. Rasmussen, G. Reeves, and A. Sacks, “Software architecture themes in JPL’s Mission Data System,” in *Spaceflight Mechanics 1999: Proc. AIAA-99*, 1999.
- [16] K. Havelund, M. Lowry, and J. Penix, “Formal analysis of a spacecraft controller using SPIN,” presented at the Fourth International SPIN Workshop, Paris, France, Nov. 1998; also NASA’s Ames Technical Report, Nov. 1997.
- [17] B. Pell, D. Bernard, S. A. Chien, E. Gat, N. Muscettola, P. P. Nayak, M. D. Wagner, and B. C. Williams, “An autonomous spacecraft agent prototype,” *Autonomous Robots*, vol. 5, no. 1, Mar., pp. 29–52, 1998.

Appendix A. List of Telemetry Channels and Names

The bulk of RAX monitoring and validation during the experiment was from the RAX telemetry on APID 9 & 10, channels W-500 to W-570, and the downlinked log files.

Channel	Mnemonic
W-500 through W-570	(RAX channels)
P-0300	LPE_PASM_mgr
APID 9 and APID 10	Monitored RAX behavior. Packets were in a RAX-specific format.
APID 45	Log files downlinked after the experiment (plan files and detailed execution trace).

The following channels were also activated for RAX:


Channel	Mnemonic
F-1048	FaultEnaStat
F-1052	BusSCstatus
F-1055	IPS_SCstatus
F-1057	PDS_SCstatus
F-1058	ACS_SCstatus
F-1060	RAX_SCstatus
F-1063	BusGDstatus
F-1066	IPS_GDstatus
F-1068	PDU_GDstatus
F-1069	ACS_GDstatus
F-1071	RAX_GDstatus
D-0149	buf_pkt_09
D-0150	sent_pkt_09
D-0165	buf_pkt_10
D-0166	sent_pkt_10
F-0716 through F-0727	

Appendix B. Date of Turn-on/off and Frequency of Data Capture

The Remote Agent Experiment first ran from May 17, 1999, 5 am PST to Wed May 19, 1999, 7 pm PST. It ran again from May 21, 1999, 7:15 am PST to 1:30 pm PST (RAX_STOP). The log files were downlinked by May 21, 1999 4:00 p.m. PST.

Appendix C: List of Acronyms and Abbreviations

AutoNAV	Autonomous Navigation subsystem of FSW	MIR	Remote Agent Mode Identification and Recovery module (Livingstone)
ACS	Attitude Control Subsystem of FSW	MR	Mode Recovery component of MIR
APE	Attitude Planning Expert subsystem of FSW	MM	Remote Agent Mission Manager module
ARC	Ames Research Center	NASA	National Aeronautics and Space Administration
CCB	Change Control Board	NewMAAP	New Millennium Autonomy Architecture rapid Prototype
CPU	Central Processing Unit (computer)	OD	Orbit Determination
DDL	PS Domain Description Language	OPNAV	Optical Navigation Module subsystem FSW
DS1	Deep Space 1 spacecraft	PASM	Power Actuation and Switching Module
DSN	Deep Space Network	PEF	Predicted Events File
ESL	Executive Support Language	PR	Problem Report
EXEC	Remote Agent Smart Executive	PS	Remote Agent Planner/Scheduler
FTE	Full Time Equivalent	RA	Remote Agent
FSW	DS1 Flight Software	RAX	Remote Agent Experiment
GMT	Greenwich Mean Time	RAXM	RAX Manager
HGA	High Gain Antenna	RCS	Reaction Control System
HSTS	Heuristic Scheduling Testbed System	RT	Remote Terminal
IPS	Ion Propulsion System	TDB	HSTS Temporal Database
JPL	Jet Propulsion Laboratory	TVC	Thrust Vector Control
MICAS	Miniature Integrated Camera And Spectrometer		
MI	Mode Identification component of MIR		




The Scarlet Solar Array: Technology Validation and Flight Results

David M. Murphy
dmurphy@aec-able.com

AEC-Able Engineering Co., Inc
www.aec-able.com

Tel: 805.685.2262



AECABLE



JPL
Schafer

Deep Space 1 Technology Validation Report

Table of Contents

Section	Description	Page
1.	EXTENDED ABSTRACT	1
2.	INTRODUCTION	3
3.	TECHNOLOGY DESCRIPTION	3
3.1	Overview	3
3.2	Key Validation Objectives at Launch	3
3.3	Expected Performance Envelope	4
3.4	Detailed Design Description	5
3.4.1	Overview	5
3.4.2	Mechanical Description	5
3.4.3	Electrical Description	5
3.4.4	Optical Description	6
3.4.5	Thermal Design	7
3.5	Technology Interdependencies	8
3.6	Test Program	8
3.6.1	Ground Test Validation	8
3.6.2	Flight Test Validation	10
3.7	Analysis of Flight Test Results	10
3.7.1	Deployment	10
3.7.2	Pointing	12
3.7.3	Temperature	13
3.7.4	Power	15
4.	TECHNOLOGY VALIDATION SUMMARY	18
5.	APPLICATION FOR FUTURE MISSIONS	19
6.	ACKNOWLEDGMENTS	19
7.	REFERENCES	19

Figures

Figure 1.	One wing of <i>Scarlet</i> for DS1 showing module level details: Lens, Frame, and Photovoltaic Receiver ...	1
Figure 2.	DS1 <i>Scarlet</i> (Wing 1 of 2) on Deploy Rail	3
Figure 3.	DS1 Mission Timeline vs Sun Distance	4
Figure 4.	<i>Scarlet</i> Module: Lens and Receiver	5
Figure 5.	DS1 <i>Scarlet</i> Wing	5
Figure 6.	Temperature Profile Across Cell	7
Figure 7.	Temperature Profile Across Panel (From cell center to centerline between module rows)	7
Figure 9.	<i>Scarlet</i> String I-V Curves	9
Figure 10.	Ascent to Deployment Temp. Transient	11
Figure 11.	Deployment Duration vs. Temperature	12
Figure 12.	Light Collection versus Alpha Angle	12
Figure 13.	Pointing Validation Summary	12
Figure 14.	RTDs: 4 on a Module (in 2 Locations)	13
Figure 15.	Thermal Modeling Results for 1.0174 AU	13
Figure 16.	Steady State Temperature vs. Power Draw	14
Figure 17.	Short Circuit Current Degradation	15
Figure 18.	Historical Sunspot Activity	16
Figure 19.	Tap Module - Wing 1, Panel 1	16
Figure 20.	SPEAK Sequence Flight Data	16
Figure 21.	SPEAK Sequence, Incrementing Voltage Near Peak Power	16
Figure 22.	SPEAK Data, Array Output at 1.1185 AU	17

Table of Contents

Section	Description	Page
Figure 23.	Mini-SPeak 1, Mission Day 160	17
Figure 24.	Mini-SPeak 1, Peak IV Data	17
Figure 25.	Original Forecasts and Flight Data	18

Tables

Table 1.	Optical Losses in Lens Assy.....	7
Table 2.	Random Vibration Test Levels.....	9
Table 3.	Protoqual Deploy Timing Ranges.....	11
Table 4.	Flight Release and Deploy Durations.....	11

Appendices

Appendix A	Telemetry Channels Related to Array Technology Validation.....	A-1
Appendix B	In-Flight Array Validation Activities and Data Summary	B-1

The Scarlet Solar Array: Technology Validation and Flight Results

1. EXTENDED ABSTRACT

The Solar Concentrator Arrays with Refractive Linear Element Technology (SCARLET) system used on the Deep Space 1 (DS1) spacecraft has been validated through successful performance in flight. *Scarlet*[™] is the first successful concentrator array ever used as primary power for a spacecraft.

Flight results to date show that performance projections were within 1% of measured results, making *Scarlet* one of the highest performance solar arrays ever used in space. The *Scarlet* array uses linear, arched Fresnel lens concentrators to focus sunlight onto narrow rows of multiple band gap solar cells to produce 2.5-kW of power. This paper describes the array technology, development process, array assembly and qualification, and flight operations of this novel system.

DS1, the first of the NASA New Millennium series of exploratory spacecraft, was launched in October 1998 and completed its primary mission in July 1999. The primary objective for DS1 was to test advanced technologies that can reduce the cost or risk of future missions. Although part of the advanced technology validation study, the array is also the power source for the spacecraft and its NASA Solar electric propulsion Technology Application Readiness (NSTAR) electric propulsion system. The array continues providing power to DS1 and its NSTAR ion electric propulsion system on the way to the next encounter object.

Sponsored by the Ballistic Missile Defense Organization (BMDO), the *Scarlet* concentrator solar array is the first space application of a refractive lens concentrator and the first to use both dual and triple junction solar cells. As part of the DS1 validation process, the amount of diagnostics data acquired was more extensive than would be the norm for a more conventional solar array.

These data included temperature measurements at numerous locations on the 2-wing, 4-panel per wing, solar array. For each panel, one 5-cell module in one of the circuit strings was wired to obtain complete I-V curves. The data was used to verify sun pointing accuracy and array output performance. In addition, the spacecraft power load could be varied in a number of discrete steps, from a small fraction of the array total power capability, up to maximum power. For each of the power loads, array operating voltage could be measured along with the current output from each wing.

The performance of *Scarlet* on DS1 substantially validated all aspects of the novel structural platform, Fresnel optics, multi-junction cell module performance, and electrical design. The major features of safe stowage through launch, deployment, and sun acquisition were clearly demonstrated on the first day of the mission. Stability of the array system, in particular, the ability to maintain the relatively tight pointing, has been verified over more than a year. The array performance has continued to achieve design specifications without imposing any onerous requirements on the spacecraft.

The main feature of the technology is that for a given power level, the *Scarlet* optical system reduces the required solar cell area by approximately a factor of seven. The decreased cell area can significantly reduce solar array cost while at the same time providing state-of-the-art performance. *Scarlet* allows the cost-effective implementation of advanced cell technologies, as demonstrated on DS1, especially when early production may limit availability or greatly elevate costs.

Another particular advantage of *Scarlet* is for applications with severe radiation environments where a thick cell coverglass is needed. The low cell area fraction means that thick glass won't significantly increase wing mass. This can be a mission-enabling feature for MEO orbits or for exploration at or near large planets with high trapped radiation levels. Often interplanetary missions must contend with the debilitating effect on cell performance that low light intensity and low temperature (LILT) can cause. The concentrating optics of *Scarlet* can be utilized to overcome these performance losses as well.





Figure 1. One wing of *Scarlet* for DS1 showing module level details: Lens, Frame, and Photovoltaic Receiver

To obtain further information on the development and commercialization of the *Scarlet* technology please contact the author, Dave Murphy, or ABLE's marketing technical director, Brian Spence, at (805) 685-2262. Learn about our past history and recent developments, such as *Scarlet* advancements, on the web at www.aec-able.com.

SCARLET™



LOW COST
LOW WEIGHT
RADIATION HARD




The Ballistic Missile Defense Organization sponsored Scarlet™ for use on DS1. Scarlet™ technology was developed with ENTECH, Inc, who supply the Fresnel optics, and the NASA Glenn Research Center at Lewis Field

ABLE provided the Scarlet™ solar concentrator arrays for the Deep Space 1 (DS1) spacecraft. Launched October 24, 1998, the solar arrays deployed flawlessly and have operated according to pre-flight predictions ever since. DS1 is the first spacecraft primarily powered by Scarlet™ solar arrays and will rely on them to energize the electric propulsion and other systems during the full course of the mission.

The revolutionary Scarlet™ arrays employ a patented refractive Fresnel lens system, which concentrates sunlight onto the solar cells. Because of this, less solar cell area is required, providing tremendous weight and cost savings. Additionally, using fewer and smaller cells allowed the cost-effective implementation of high efficiency multi-junction GaInP2/GaAs/Ge photovoltaic cells aboard DS1.

The DS1 mission is part of The New Millennium Program, NASA's most aggressive technology demonstration program. According to Ray Garza, ABLE's DS1 Scarlet™ Program Manager, "This high technology mission challenged us to build the most advanced solar array in the world." Such strides were made with this new technology that ABLE was recognized in 1999 with a NASA Group Achievement Award and by the University of New Mexico's Institute for Space and Nuclear Power with their Schreiber-Spence Achievement Award.



The performance of the Scarlet™ arrays on DS1 validated all aspects of the novel structural platform, optics, and electrical design as well as the analytical models used to characterize the array. The Scarlet™ technology proven on DS1 will continue to be refined to benefit future science mission as well as commercial endeavors such as mid-level orbit satellites, and communication constellations.

DS1 SCARLET™

SPECIFICATIONS

Wing Dimensions: 206 in. x 64 in.
Panel Dimensions: 45 in. x 63 in. (4 panels per wing)
Array Power: 2500 W (1 AM0)
Wing Mass: 27.7kg (with tiedowns)

Specific Power: 45 W/kg
Stowed Stiffness: 92 Hz
Deployed Stiffness: 0.37 Hz
Deployed Strength: 0.015 g's

AEC-ABLE ENGINEERING COMPANY, INC
93 CASTILIAN DRIVE
GOLETA CA 93117

HTTP://WWW.AEC-ABLE.COM/SOLAR
E-MAIL: SOLARARRAYS@AEC-ABLE.COM
TEL: 805.685.2262 • FAX: 805.685.1369

The Scarlet Solar Array: Technology Validation and Flight Results

2. INTRODUCTION

This report summarizes the design, development, and test of the Solar Concentrator Array with Refractive Linear Element Technology (SCARLET) system and details the flight validation on the New Millennium Deep Space 1 (DS1) mission. Array deployment, system pointing, thermal performance, and power production are analyzed and discussed in comparison to ground results and mission predictions. In summary, the solar array has operated flawlessly and all aspects of the technology were successfully validated in pre-launch and mission activities.

The flight of *Scarlet* on DS1 has enabled the development and validation of a low-recurring cost technology which offers significant advantages for radiation applications, such as MEO orbits, or LILT applications. For example, missions to large outer planetary bodies can benefit from the weight efficient radiation hardness and the LILT advantages of the *Scarlet* technology. This novel flight-validated solar array is a cost-effective and mission-enabling technology.

3. TECHNOLOGY DESCRIPTION

3.1 Overview

Scarlet is a concentrator solar array for space applications which uses linear refractive Fresnel lenses to focus sunlight onto spaced rows of solar cells. For a given power level, the *Scarlet* optical system reduces the required solar cell area by approximately a factor of 7. The decreased cell area can significantly reduce solar array system cost and weight, especially in high radiation environments where thick cell coverglass is required.

The DS1 array is derived from and scaled up from the prototype *Scarlet* wing that was built for the METEOR satellite in 1995. Due to the failure of the Conestoga launch vehicle, DS1 was the first flight of *Scarlet* technology. This second-generation *Scarlet* solar array incorporated many additional advanced technologies such as multi-junction solar cells and a new mechanization and structural design.

AEC-Able Engineering Company, Inc. (ABLE), designed, assembled, and tested the 2.5 kW concentrator solar array for the DS1 mission, which launched on October 24th of 1998. The Ballistic Missile Defense Organization (BMDO) Innovative Science and

Technology Directorate has sponsored development of *Scarlet* through the first New Millennium Space flight on the Deep Space 1 spacecraft. Substantial funding support and technical aid for this application was also provided by the Jet Propulsion Laboratory (JPL).

The DS1 *Scarlet* solar array has made significant advances in the state-of-the-art of space-demonstrated solar cells, concentrators, lightweight packaging, and deployment techniques.




Figure 2. DS1 *Scarlet* (Wing 1 of 2) on Deploy Rail

The pioneering success of the *Scarlet* Array was recognized with the 1999 Schreiber-Spence Award for Significant Technology Advances. This year's award was presented to the NSTAR and *Scarlet* teams in recognition of the first use of solar-powered ion propulsion as primary propulsion and of the first use of a multi-band-gap, concentrator array for a robotic, deep-space mission, thereby helping to open the solar system to frequent, low-cost exploration. *Scarlet*, with its radiation hardness capability and ion propulsion employed together, also provides an excellent combination for the cost-saving concept of orbit raising from LEO.

3.2 Key Validation Objectives at Launch

Data collection objectives for technology validation were formulated for two phases of the mission. Within the first month of the mission the data desired was:

- Initial power telemetry data collection on earliest day possible using all tap circuits and temperature sensors: 8 taps and 10 Resistive Temperature Detectors (RTDs).

This data verifies initial performance prior to on-orbit calibration. RTDs were used to extrapolate cell

temperatures and to validate the array thermal design by measuring gradients in the structure around the focal line of the tap modules.

- On-orbit calibration to maximize power output of the array prior to beginning cruise phase.

This data was used to evaluate and validate the accuracy of the initial alignment of the array/spacecraft with respect to perceived attitude.

- Power telemetry data sets taken nominally every week to validate performance vs. AU, temperature, and environmental degradations.

It was important to record data often at the start of the mission to capture early degradation effects such as spacecraft or array outgassing contamination and UV darkening.

For beyond the first month the data desired was:

- Power telemetry data sets nominally every two weeks (every month as a minimum when or if mission events conflict) to validate performance vs. AU, temperature, and environmental degradations.

Regular data sets form the basis for validating and correlating power output and modeling.

Criteria for incremental success in flight validation were developed and documented in the New Millennium Program - Deep Space One Project Technology Validation Agreement between BMDO, JPL, and ABLE. Those criteria were:

- 50% Successful zero-g deployment of both wings
- 60% Successful acquisition of launch and deployment activities data, Spacecraft orientation from spin-down to start of deployment event, Time history of the states of telemetry switches and RTDs
- 75% Successful acquisition of initial power telemetry data set, which includes: Current and voltage (IV) curve data points for each of the 8 tap circuits, Temperature readings from each of 10 RTDs, day and time, Heliocentric distance, Wing orientations: alpha for each wing and beta of spacecraft, Spacecraft orientation reference data and alpha/beta offsets
- 80% Successful acquisition of initial on-orbit alpha-beta calibration data
- 90% Produce power in excess of 2400 W at BOL

100% Successful acquisition of periodic power production data

In summary, each of these goals was fully met, and thus, 100% success was attained, with one caveat: The "periodic" power telemetry data sets were not as numerous as planned. The spacecraft has had a series of anomalies unrelated to the solar array that caused delay and postponed *Scarlet* validation activities by several months.

As the mission has progressed, opportunities to refine the understanding of detailed modeling factors have become available. By March of 2000 the spacecraft will be back to 1.1 AU after traveling out to almost 1.35, and after nine more months will have reached 1.35 again. This mission profile, shown in Figure 3, will allow analysis to separate time and distance effects.




Figure 3. DS1 Mission Timeline vs Sun Distance

3.3 Expected Performance Envelope

The performance envelope of any future *Scarlet* array is well understood as many detailed point designs have been generated from the DS1 baseline. It has been configured into arrays from 500 to 25,000 W, for LEO, MEO, GEO, and interplanetary missions. The key metrics for array evaluation are specific power (W/kg) and cost. The cost advantage with concentration is self-evident, and the remaining questions for planning or selection are whether the array will meet all mission requirements and at what level of performance.

Scarlet technology has now been proven out with the flight of DS1 and the completion of subsystem qualifications tests which complete the verification of thermal cycle capability in various environments. The key measure of performance, specific power, demonstrated on DS1 is at the state-of-the-art (45 W/kg) and can easily be increased with now proven design enhancements and/or size increase. The specific power increases as wing size grows because the mechanism overhead of tiedowns and yoke diminish relative to the photovoltaic content. For example, with larger arrays the specific power approaches 70 W/kg for a 15 year GEO environment.

Generally speaking, *Scarlet* technology versus standard array technology is transparent to the spacecraft user, with one exception: pointing. DS1 has shown that the accuracy level required is not a significant engineering or operations burden. The required accuracy of the spacecraft knowledge and pointing was set at 0.5° maximum. Typically communications spacecraft point to better than 0.1° , so this was not a novel challenge.

However, if needed, the technology can also be configured for wider pointing acceptance angles. The optics are not sensitive to errors in one axis and thus GEO orbit seasonal off-pointing ($\pm 24^\circ$) can be accommodated with only a small reduction in performance.

3.4 Detailed Design Description

3.4.1 Overview

The basis of the technology is to use a linear refractive Fresnel lens to focus sunlight onto a 1 cm wide strip of solar cells as shown in Figure 4.




Figure 4. *Scarlet* Module: Lens and Receiver

The *Scarlet* array for DS1 is based on the prototype *Scarlet* array that was built for the METEOR satellite in 1995. [ref. 1] After the failure of the Conestoga launch vehicle, the BMDO Innovative Science and Technology Directorate sponsored the development of this second-generation *Scarlet* solar array - which incorporates advanced technologies such as multi-junction solar cells and an improved structural design - for use and validation on DS1.

The first generation *Scarlet* array was a melding of ABLE's standard planar array structure, PUMA [ref. 2], with concentrator optics. That structural baseline was reassessed for the DS1 *Scarlet* design to improve the union between the cell substrates and lens panels. The result is a simple cable-synchronized structure which deploys flat. The major advantages are fewer piece parts,

simplified pointing control analysis, reduced stowed volume, and simplified yoke structure. Additionally, the lens panels are held securely between power panels in the stowed condition.

Basic proven mechanisms such as release assemblies, tiedown cup-cones, cable pullers, and hinges of *Scarlet* I were utilized again for *Scarlet* II, but optimized to minimize weight. [ref. 2]

3.4.2 Mechanical Description

The DS1 *Scarlet* solar array consists of two wings of four panels each. The wings are delivered fully integrated with tiedowns, gimbal drive assembly, and spacecraft interface plate as shown below.




Figure 5. DS1 *Scarlet* Wing

Deployment of a wing is initiated when power is applied to the high output paraffin (HOP) linear actuators in each of two tiedown assemblies. A resistive load causes the paraffin to heat and change phase, which forces a pin forward releasing the restraint arm on a torsion tube. Tiedown cables are wrapped and captured in fittings on either end of the tube, so when a torsion spring revolves the tube, both cables are released. When the second release mechanism has actuated, the wing unfolds, driven by double-wound torsion springs distributed on each hingeline.

The hingelines are synchronized by a system of cables which are wound over static pulley cams. The synchronization transfers the deploy torque to the root where redundant rotary viscous dampers retard the deployment rate.

3.4.3 Electrical Description

The cells used by DS1 *Scarlet* are about 1 cm wide and 4 cm long and are spaced in rows 8.6 cm apart. The low cell area per watt needed beneath the concentrator greatly lowers cost and also eases the risk of utilizing emerging, high-performance cell technologies. For this reason,

BMDO elected to specify the procurement of an entirely multibandgap-cell-based solar array.

In 1996 production quantities of GaInP₂/GaAs/Ge dual-junction cells were not yet available. Tecstar was the only cell vendor willing to participate, and the DS1 team was cautious about the difficulties of bringing new cell technology into production. So to mitigate risk, and to set performance criteria for the flight build, an engineering build quantity of 100 cells was procured.

The III-V cell design, termed “Cascade” by Tecstar, had previously been qualified in the standard series of environments for space applications. The only modification required was gridline sizing for the high flux profiles of the concentrator.

The engineering evaluation result was very encouraging, with the average efficiency result coming in at 24.25% at 7.5X air mass zero (AM0). The performance criteria for the flight build was set at 23.25%, - partially because of losses anticipated for over-glassing, but mostly as insurance against the uncertainties of a larger build.

During the flight production phase Tecstar experienced a series of setbacks in producing the flight cells. The most persistent problem was shunting (Reviewed in [ref. 3]) which reduced the performance of many of the cells to as low as 16% at 1 sun intensity. Fortunately, the high current injection levels created by the lens overrides the fixed magnitude shunts and the performance at concentration is only slightly degraded.

After intensive effort by Tecstar, and aided by the synergism of the early dual-junction-cell manufacturing technology (ManTech) development program, remarkable improvement in yield and performance were achieved. But schedule delays eventually forced the acceptance of cells with a minimum lot average - under 7.5X concentration - of 22.6 %.

The cells were covered by Tecstar with 0.004-inch-thick coverglass with an anti-reflection coating with blue/red filtering (BRR). The reflection of the near infrared lowers the operating temperature of the cell by 11°C.

The cell receiver module consists of five series cells, each with bypass diodes, affixed to a circuit on a high thermal conductivity substrate as depicted in Figure 4. The modules are joined using overlapping redundant tabs with reflowed solder to form 50 cell strings that generate 40 watts at an operating voltage of 90 volts at 1 AU.

Cells in the module are interconnected along both their long edges. Given the long aspect ratio (4:1), the most probable crack direction will never leave a section of the cell isolated. Dual ohmics also provide balanced off-track performance and lower gridline resistance losses.

Cell interconnect reliability is also greatly improved over standard CIC construction because 120 interconnects (in parallel per cell) connect the cell to the circuit board carrier. The automated wire bonder, which stitches at a rate of three cells per minute, results in large cost savings by eliminating hand labor.

Engineering modules underwent thermal cycling from -160°C to +110°C for 100 cycles to assure a high margin of compatibility with the single thermal cycle experienced on the DS1 mission at the start of its interplanetary mission. All modules experienced no visible degradation and comparison of pre- and post-IV curves under the X25 solar simulator at NASA Glenn Research Center (GRC) showed no measurable electrical degradation.

To demonstrate the orbital applicability of the *Scarlet* technology, flight modules and lenses were later successfully cycled for GEO thermal extremes for 1350 cycles and to MEO extremes for 40,000 cycles.

3.4.4 Optical Description

The Fresnel lens is comprised of precisely formed individual ridges which refract incident light from a 3.22-inch aperture down to a strip of light focused in the middle of the 0.40-inch-wide cell strip to leave margin for pointing error.

The average optical efficiency of the DS1 lenses, which have no anti-reflective (AR) coatings, (used with the Cascade cell described above) has been measured at 89%. The effective concentration ratio, 7.14 (= 0.89 x 3.22/.40), was selected to provide for reasonable pointing error. The purpose was to create a cost-effective system to manufacture and assemble which is compatible with standard gimbal and spacecraft ACS architectures.

The linear Fresnel pattern is molded in a continuous roll process using space-grade silicone. Individual lenses are machined-trimmed and bonded to glass superstrates which have been thermally formed into cylindrical sections. The materials chosen for the lens, the bondline, and the glass are well understood: DC 93-500 silicone and ceria-doped borosilicate glass (Corning 0213). The glass protects the lens from particle radiation and with an AR/ITO (Indium Tin Oxide) coating, planned for future programs, the optical efficiency is enhanced and charge buildup is minimized.

The space between lenses must be minimized to maximize packing factor (the ratio of the area of light which passes through the lens to the total panel area). To demonstrate the survivability of the thin glass lens mounted in this minimal structure frame, five lens-in-frame components were tested - successfully - to conservative local acoustic/random levels (29 G_{rms} out-of-

plane, 9 in-plane). In qualification testing of the completed wings less than 2% of lenses had any cracking.

The efficiency of the lens overall is a function of the refractive index matching of the lens, superstrate, and optical coating used, as well as the surface finishes and sharpness of the lens teeth. The manufacturing of the lens produces smooth and sharp prisms with small root radii. The close match between the refractive index of the silicone and the glass (1.523 and 1.409 respectively) causes a slight loss of 0.3%. The losses are summarized in Table 1.

Well established optical coatings would reduce the large loss the lens outer surface transmittance, but the thermal forming of the lens superstrates occurs at a temperature which is higher than the survival temperature of typical coatings. Application of AR coatings to the curved surface of the glass superstrate was developed at OCLI, but not in time to coat all flight lenses.

Table 1. Optical Losses in Lens Assy
(without AR coating)

Component	Material	Interface Reflection	Absorptance Scattering
Space	Vacuum		0.0%
Cover	Glass	4.5%	0.5%
Lens	Silicone	0.3%	3.1%
Space	Vacuum	3.0%	0.0%
Totals		7.7%	3.6%
		Combined Loss:	11.0%
		(Multiplicative along light path)	

Two coated lenses flew on the DS1 array in positions where their contribution to module efficiency could be measured and compared to non-coated lenses. The coating developed combined AR performance with electrical conductivity, using ITO to dissipate charge to the grounded lens frame structure. Component testing of the coated lenses demonstrated a 2% efficiency gain.

3.4.5 Thermal Design

The thermal design challenge is to spread the absorbed but unconverted solar energy (heat) from the cell modules out across the panel to engage the full area and high emissivity of the graphite panel to radiate efficiently. The cell, circuit layers, and panel were analyzed with a detailed finite difference model so that material choices and thicknesses could be optimized to reduce the cell temperature. The calculation results, for 1 AU illumination, are shown in Figure 6.




Figure 6. Temperature Profile Across Cell
(Edge at D = 1)

The largest temperature rise is caused by the Kapton. The thermally-conductive silicone adhesives and the cell carrier present very little resistance to thermal conduction. The carrier, which is used during cell laydown to stabilize the Kapton circuit and thus protect the cells from damage, is made of a high-conductivity composite to match the thermal expansion coefficient (nearly zero) of the panel. This creates minimal strain along the long bondline, and in addition, the material is very light and stiff.

As can be seen in Figure 7, once conducted to the panel the heat spreads rapidly out through the facesheets and core. This is because the facesheets are constructed of an ultra-high-conductivity fiber (which also possesses good compressive strength) which has a conductivity of 384 W/m-K (for unidirectional layup, 60% fiber volume), which is 60% higher than pure aluminum.




Figure 7. Temperature Profile Across Panel
(From cell center to centerline between module rows)

Typically, *Scarlet* cells operate about 20°C hotter than a planar GaAs design, mostly due to front surface radiation blockage by the lenses and the temperature gradients associated with heat spreading. The thermal modeling was first correlated to a balloon flight module [ref. 4] then later, with a vacuum thermal balance test at NASA GRC. The excellent agreement with the flight results, discussed later, was aided by these early validations.

3.5 Technology Interdependencies

The performance of many subsystems is often critical to the survival of a spacecraft. Of the dozen new technology experiment on DS1, several also functioned as bus subsystems. Perhaps the most critical to the mission is the power system - the *Scarlet* array. Without a functional solar array, a spacecraft cannot long survive. If the array had suffered even slightly degraded performance the ion engine may not have been able to thrust at a level sufficient to meet the first encounter target.

But the critical nature and unique features of the *Scarlet* array were generally transparent to the other spacecraft subsystems. It was mated to the spacecraft with no more interfaces or complexity than a generic solar array.

One important exception is that in operation the array requires much tighter pointing control than standard arrays. The pointing abilities of the spacecraft were not taxed by this need. The array included Moog-SMI gimbal assemblies with excellent orientation capability. The largest error sources were in the sub-assemblies of the wing itself. At one point in the mission a software, potentially a single event upset, caused one wing to off-point significantly. The inherent redundancy of the two wing system prevented the temporary power loss from threatening the mission.

3.6 Test Program

Due to limitations in evaluations that can be feasibly executed on-orbit, a thorough validation of technology such as *Scarlet* is highly dependant on ground testing. Of course, a systematic qualification test program is also essential to minimize flight risk. The ground test program (reviewed in detail in ref. 5) is discussed below as a prelude to the flight observations.

3.6.1 Ground Test Validation

Wing qualification testing was initiated once assembly of each wing was completed. No spare hardware was fabricated and a protoflight test approach on the flight hardware was used. The levels for the protoflight testing were defined by the DS1 Component Verification Specification (CVS).

Figure 8 shows the individual tests that were included in the program as well as the order that the tests were performed. Following each of the major tests (thermal cycle, acoustic, random vibration), the wing was deployed and inspected to verify that no critical damage occurred during the test. Before and after the full testing sequence, a full electrical functional test was performed on the wing to verify that the power level and electrical functionality of the wing was not degraded by the exposure to the test environments.

Deployment Tests: As part of the initial deployment tests, the array was deployed at thermal extremes, including a 10°C margin, based on modeling of the possible conditions in space at the time of deployment. The array was successfully deployed at -66°C, 30°C, and at ambient temperature.

Thermal Cycle Test: Thermal cycling was not a major issue for the DS1 mission because after the spacecraft leaves the earth's shadow following launch, it is in the sun for the rest of the mission. However, the CVS required a limited number of thermal cycles to ensure that the initial cycle from ambient (launch) to cold (umbra) to hot (in the sun) would not be a problem.

The wings were cycled three times between -123°C and +113°C. The temperatures were determined by mission analysis of the hottest and coldest possible conditions, plus margin. The tests were conducted in a dry nitrogen environment. Prior to array assembly, all of the flight array components had also gone through at least three thermal cycles.

Following the test, the arrays showed no measurable power reduction and no structural damage. Nearly 4% of the glass concentrator lens superstrates developed small cracks during the testing. However, the shape of the lens was maintained because the curved shape was formed in a zero stress state. The silicone adhesive also helped to maintain the configuration of the lens. No deformation or optical degradation was observed in the cracked lenses. All of the cracked lenses subsequently survived the acoustic and random vibration environments. Thus, the program decided not to replace the cracked lenses.

Acoustic Test: The wings were exposed to acoustic environments between 105 and 135-dB over a frequency range of 30 to 10,000 Hz during a 1-minute test. The arrays experienced no measurable power reduction or structural damage due to the test. As with the thermal cycle test, a small number of concentrator lens glass superstrates were cracked during the test, less than 2% of superstrates in this case. Again, there was no deformation or optical degradation and all of these lenses subsequently

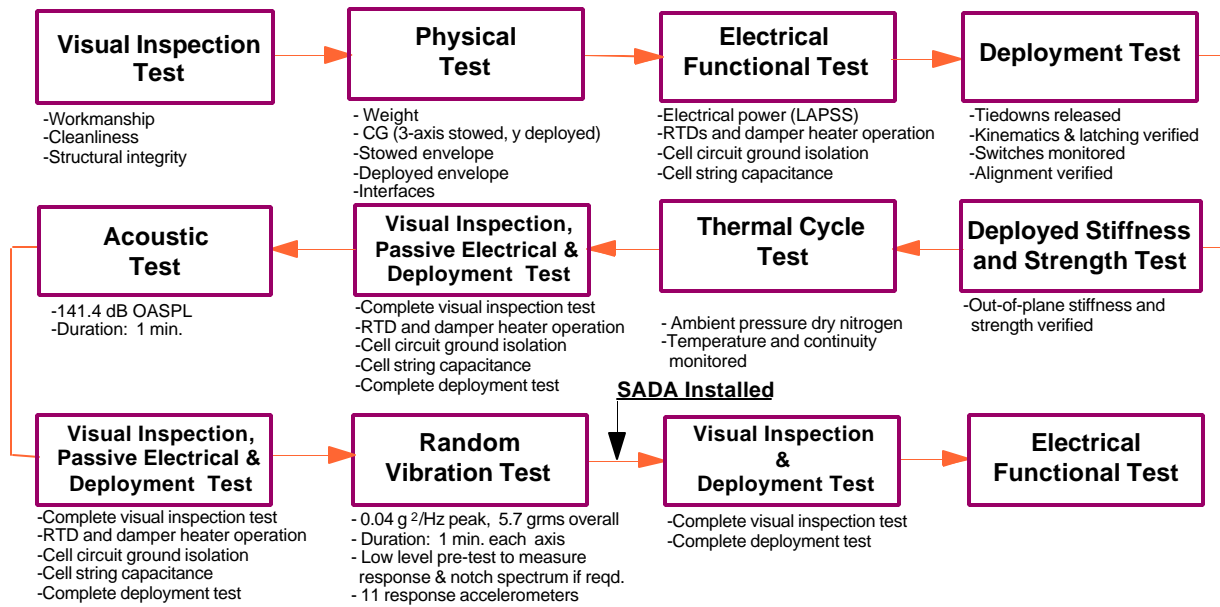


Figure 8. Protflight Test Sequence

survived the random vibration test. No other problems were observed during this test.

Random Vibration Test: The wings were random vibration tested in all three axes to the levels shown in Table 2. The test duration was one minute in each axis. As with the previous protflight test, the arrays experienced no measurable power reduction or structural damage due to the test. Again, a small number of concentrator lens glass superstrates were cracked during the test; in this case, less than 1% of superstrates. As in previous tests, there was no deformation or optical degradation.

Frequency (Hz)	PSD (g^2/Hz)
20	0.0016
50	0.04
500	0.04

Table 2. Random Vibration Test Levels

Power Measurement: The power was measured using a Large Area Pulse Solar Simulator (LAPSS). Due to the concentrator lenses, the light must be collimated perpendicular to the wing for the power to be measured. This required that wing power be measured one string at a time. The configuration of the strings results in an area of 17.3-cm across by 110-cm wide for each string. Testing showed that the light from the LAPSS was sufficiently collimated over this area. Special filtering was required to improve the LAPSS spectral balance for accurate multi-junction cell measurement (ref. 6).

All of the measurements performed as expected during these tests. Figure 9 shows the results of testing a string with and without the concentrator lenses. The string I_{sc} for a long series string was consistent with the measured average lens concentration ratio of 7.14. The string V_{oc} boost of 8% was consistent with individual module results. The string P_{max} increased by 4% more than expected (increased fill factor). This was determined to be due to the presence of shunts in some of the cells. Under concentration, the cells generate more current, so the shunts become less significant.




Figure 9. Scarlet String I-V Curves

3.6.2 Flight Test Validation

A large portion of the validation of *Scarlet* in flight was accomplished with the initial survival of the launch environment and successful deployment of the mechanism. After many months of data accumulation the power capabilities and robustness of the wing in the space environmental was convincingly demonstrated.

To reach our first goal at a 50% validation level, required the array - with its novel cable-linked synchronization, hinges, root articulation, four-bar lens panel kinematics, and latches - to perform as expected. 60% was obtained when the telemetry data was obtained and evaluated. A detailed review of the results is included in "Analysis of Flight Test Results, Deployment," Section 3.7.1.

Receiving the first power telemetry data set obtained 75% validation as even one data set gives flight confirmation of the capability of the technology to produce power in the space environment.

The original Power Telemetry data set was comprised of:

- Current and voltage (IV) curve data points for each of the 8 tap circuits
- Temperature readings from each of 10 RTDs
- Day, time, and heliocentric distance
- Wing orientations angles
- Spacecraft orientation reference data and wing angle null offsets

The information for wing and spacecraft orientation turned out to be superfluous. A measurement of the available power versus gimbal position was performed early in the mission as planned. The experiment demonstrated that the wings were so well aligned that power roll-off was not a factor. These results, which pushed successful validation past 80%, are discussed in "Analysis of Flight Test Results, Pointing," Section 3.7.2.

The operating temperature is a key element of the flight validation. The wings were well instrumented to measure cell temperature and gradients in the structure around the focal line. An analysis of the performance and a comparison to pre-flight expectations is discussed in "Analysis of Flight Test Results, Temperature," Section 3.7.3.

The sequence of commands executed by the spacecraft to collect and store the power telemetry data set is termed SIVPerf, for solar array IV performance. Because SIVPerf measures the IV curves of a single module within a string of 10 modules on each panel, whereon there are 9 strings, the calculation of total wing power production is a large extrapolation.

While the SIVPerf data is useful for certain investigations, there is another data set which gives much better confidence for full array power.

Successful execution of that sequence, termed SPEak for solar array peak power, affirmed the validation at 90%: Power production in excess of 2400 W. The various power telemetry from both SIVPerf and SPEak have been analyzed, combined, and contrasted to track the array power versus heliocentric distance and time.

The system power production has followed the original model very closely, thus completing 100% validation of the *Scarlet* technology in-flight, from mechanism structure and kinematics through thermal/optical/electrical modeling and power performance. The analyses of power performance are discussed in "Analysis of Flight Test Results, Power," Section 3.7.4.

3.7 Analysis of Flight Test Results

Analysis of the in-flight validation of *Scarlet* is grouped into four areas of performance: Deployment, pointing, temperature, and power. The following sections address the flight activities and data gathered and compare the data to pre-flight analyses and forecasts.

3.7.1 Deployment

The first activity required from the *Scarlet* system in space, deployment, occurred 1 hour after launch. The sequence plan for deployment of the wings was as follows:

0. Power to damper heaters (ON since launch)
1. Disable attitude control system (to prevent reaction during wing motion)
2. Power HOP primary heaters for 180 seconds
3. Power HOP secondary heaters for 180 seconds
4. Power HOP primary heaters for 180 seconds
5. Power HOP secondary heaters for 180 seconds

During steps 3-5, if all 8 release indications OR all 4 deployment indications become true, wait 10 sec then turn off the HOP heaters and go to step 6.

6. Wait 240 seconds for extension of wings ("deployment")

During steps 3-6, if all 4 deployment indications become true, turn off the damper heaters.

7. Turn OFF the damper heaters
8. Enable reaction control system
9. Index solar array gimbals

Through considerable discussion this conservative and straightforward algorithm was developed by JPL with ABLE input. In flight none of the redundancies proved necessary as the deployment was nominal. The logic of this sequence is as follows:

- Requiring all 8 tiedown release or all 4 deployment indications was to protect against a false positive, while still allowing quick re-enabling of attitude control and positioning of the array to the sun.
- Waiting 180 seconds before looking at the switch status meant that a primary heater failure could be tolerated too. This double failure protection was felt to be justified by JPL because of problems with microswitches in the past.
- Broken wires or open connectors can cause a false positive, because in all cases the microswitch configuration desired for positive indication was “open.” False negatives from the switches were prevented from being a detrimental factor, by OR-ing the tiedown set with the deployment set, and because the sequence would run regardless of microswitch problems.
- Waiting 10 seconds after getting all 8 release or 4 deployment indications was added to prevent any possibility of indications coming before actual release.

Activation of the paraffin actuators activates the mechanism to release the panel tiedown cables. After the tiedown cables are released, the wing deploys powered by torsion springs and rate-limited by viscous dampers. The nominal timing and range in protoqual testing for the release activation and the unfolding of the wing are shown in Table 3.

Table 3. Protoqual Deploy Timing Ranges

HOP Temp (°C)	Release Time (sec.)	Damper Temp (°C)	Deploy Time (sec.)
30	60 ± 10	30	42 ± 10
18	77 ± 7	18	62 ± 5
0	110 ± 15	10	88 ± 10
-10	130 ± 30	0	160 ± 20

Power to the tiedown mechanisms was autonomously commanded at about 6:08 am PDT on October 24, 1 hour after launch. The spacecraft was in eclipse. Telemetry

was being recorded at 5-second intervals. Forty minutes later, when the real time link was reestablished at JPL and the monitors filled with data, it was evident the array was deployed: The indicator switch states were all in agreement and power was being produced.

Later analysis of the recorded data allowed the duration in seconds of the deployment events of HOP heating to SATM release and damped wing motion to latching at full deployment to be determined, as listed in Table 4.

Table 4. Flight Release and Deploy Durations

Event	Time in seconds	Wing 1	Wing 2
1 st Tiedown Release	75	70	
2 nd Tiedown Release	85	80	
Deployed & Latched	85	80	
Total Time	170	160	

The HOP heating durations on each of the four actuators ranged from 70 to 85 seconds. These values agree well with the duration predicted, 77 seconds, for the HOP temperature of 18°C. The thermal modeling of the temperature transients from fairing jettison after launch to deployment in eclipse is shown in Figure 10. The model predicted a HOP temperature of 15°C.




Figure 10. Ascent to Deployment Temp. Transient

The duration of the wing deployment depends on the damper temperature, which was forecast to be near zero. The damper body and silicone fluid and thermostat were modeled by a single node, as the primary intent was to determine how early the thermostat could turn on causing the damper heater to draw power and if the 10 watt heater was sufficient to maintain the damper above 0°C.

The actual fluid temperature would certainly lag behind the cooling of the casing exterior and the thermostat body.

So the flight temperature was probably between 0 and 15°C. Placing the average wing deployment time on the curve of predicted time versus temperature, Figure 11, suggests the fluid temperature was near 11°C.




Figure 11. Deployment Duration vs. Temperature

In summary, all telemetry indicates the deployment occurred precisely as designed. This was a significant milestone for technology validation because - although the highlight of the technology is the optical/thermal/electrical performance of the concentrator/cell module - the kinematic control and joint mechanisms were also making their debut.

3.7.2 Pointing

The criticality of proper alignment of a concentrator system is plain. System performance is dependent on all elements (cells, modules, lens frames, panels, hinges, yoke,...) being assembled accurately, being deployed reliably, and being resistant to thermal distortion. Numerous industry efforts to build concentrator systems have failed at various stages prior to launch due to the inherent design and manufacturing difficulties. The industry has had limited success of late with low concentration ratios, for example the STEX trough concentrator at 2X. The *Scarlet* system is the first system on-orbit to provide significant concentration benefits.

The Fresnel optics provide an advantage in tolerance to shape error that reflective systems lack. This technology, when properly integrated at a 7X concentration level with ± 2 degree error tolerance, creates a system with significant cost and performance benefits.

While it was demonstrated that manufacture of the piece parts and assembly of *Scarlet* was straightforward, the proof of success - power production - required on orbit data. Would each and every lens be pointed accurately to the sun within the accumulated errors of piece part fabrication, sub-assembly, system assembly, thermal distortion, spacecraft knowledge and pointing control? The eighth day of the mission provided an opportunity to find out.

On-Orbit Calibration: The sequence was termed SCal, for solar array calibration. To determine if each wing was positioned by the spacecraft (for beta) and the wing gimbals (for alpha) at the angles which provided maximum power, the wings were steered to various positions over a range of $\pm 4^\circ$ in alpha and $\pm 8^\circ$ in beta. The alignment of the system was judged by the short circuit current (I_{sc}) output, a direct indication of the light flux on the cells, of the tap modules on each panel.

The sequence ran for about 6 hours, where each beta position was selected and the various alpha positions were steered through. The drift of the spacecraft causes some random variation in the current output results, but a parabolic curve fit, see example in Figure 12, can be used to estimate whether the alignment of a module is centered or offset.




Figure 12. Light Collection versus Alpha Angle

When the offsets are compared against the design specifications, as in Figure 13, the success of the system in achieving far better alignment than required is clearly evident.




Figure 13. Pointing Validation Summary

In the event the SCal data had shown any significant difference in current between modules or wings, a topology study was planned to determine the best beta correction for the spacecraft and the best alpha correction for each gimbal. Actually, it was hoped that an adjustment of pointing to achieve maximum power would not be required. Calibration is an activity that places a burden on spacecraft operations and it was important to demonstrate to future *Scarlet* users that it isn't required.

The results of SCal demonstrated that *Scarlet* achieved the pointing accuracy goals not only for the design and assembly of the wings, but of integration with gimbals and the spacecraft structure, and for integrated performance with the spacecraft issues of position knowledge, pointing control, and drift.

3.7.3 Temperature

A critical validation of the power model is the operating temperature of the array. Each wing was equipped with resistance temperature devices (RTDs): Four on the inboard panel and four on the outboard panel. As shown in Figure 14, the cluster of four RTD's were located:

- Next to cell on the front facesheet (as close as Module Base width allowed)
- On module-to-module centerline on front facesheet
- Directly behind the cell on the back facesheet
- On module-to-module centerline on back facesheet

On the second wing only the RTDs nearest the cells were recorded, due to channel restrictions. Therefore the flight data set consists of 8 RTDs on Wing 1 and 2 on Wing 2.




Figure 14. RTDs: 4 on a Module (in 2 Locations)

A fairly detailed finite difference model was developed, starting in 1996, to analyze the complex heat balance and thermal gradients beneath the lens. The line focus and regular module-to-module spacing allows for accurate temperature predictions using a half-symmetry 2-D model. The model was refined in 1997 based on thermal balance testing performed in a 1-sun vacuum environment at NASA Glenn Research Center.

The modeling results for each node of the model are shown in Figure 15, for the module at maximum power (minimum waste heat in cell). In this case the sun distance is set at 1.0174 AU, to correspond to the first applicable flight data set, presented next. The model is shown *without* the 10°C margin used in the power prediction, because the flight data demonstrates it was not needed, as will be shown below.




Figure 15. Thermal Modeling Results for 1.0174 AU

On the 37th day of the mission the ion engine was, for the first time, commanded to thrust at increasing increments up to maximum available power. The data available on all the RTDs at the intermediate levels between zero and near full power are plotted in Figure 16.

The model prediction curve is also plotted for comparison. The general agreement is excellent. Several observations about the data can be made:

- The agreement is fairly precise, on average, near maximum power for the RTDs nearest the cells.
- The flight data shows the gradients of heat spreading across and through the panel were slightly larger than the model forecast.

Since the gradients are larger than expected, but the near-cell RTD temperatures were accurate, it can be surmised that optimistic and conservative simplifications in the model were offsetting. Two known corrections which would reduce model conservatism are:







Figure 16. Steady State Temperature vs. Power Draw

- Recent analysis of the various optical filtering effects - both gray and wavelength dependent absorptance, scattering and reflectance - of the lens superstrate, silicone lens, cell cover, and cell have shown that the fraction of sunlight which reaches the cell is less than previously assumed.
- The carbon-carbon carrier beneath the cell is 50% wider than the cell, but this was conservatively not represented in the model.

The most likely effects which would reduce the efficiency of thermal spreading are: The core to facesheet conduction is limited by the joining adhesive, and/or the facesheet conductivity is lower due to thickness or resin fraction.

The net effect of incorporating corrections for these optimisms and conservatisms would likely be of little benefit to the power modeling correlation since the thermal modeling predicted the near-cell temperature precisely and is only off by 2.5°C on the back of the panel between modules.

3.7.4 Power

The validation of power production relies on two sequences: SIVPerf and SPeak, which stand for solar array IV performance and peak power, respectively. SIVPerf measures the full IV curve of a 5cell module within a string of 10 modules on each panel, for a total of eight module level curves. SPeak produces a partial IV curve for each wing, as a byproduct of a sequence intended to utilize the maximum power available for thrusting.

Because SIVPerf measures the IV curves of a single module within a string of 10 modules on each panel, whereon there are 9 strings, the calculation of total wing power production is a large extrapolation. SPeak is a better measurement of array performance and many more data sets have been collected leading to better correlating analysis and performance projections. The SIVPerf results will be discussed first.

SIVPerf: This validation sequence was run on the earliest day possible after launch using all tap circuits and temperature sensors (8 taps and 10 RTDs) to verify initial performance prior to on-orbit calibration. The first SIVPerf was run on mission day 7, October 31, 1998, the second on mission day 18, November 11, 1998.

The SIVPerf sequence was not run again until May 25, 1999. The 28-week testing hiatus was necessary to identify and correct a power distribution unit failure mode using software, upload and verify the new software, and to schedule the activity into the busy validation/operations planning.

The SIVPerf sequence provides certain types of data the SPeak test cannot. Namely, data to the left of peak power on the IV curve. The short circuit current is of particular interest as it validates the lens optical efficiency and functions as a monitor of the combined effects of UV and radiation darkening, outgassing contamination, and micrometeoroids.

The I_{sc} values for the modules shows significant change over time as the spacecraft has traveled out to 1.3 AU. For purposes of comparison the data, shown in Figure 17, has been corrected for insolation changes due to heliocentric distance and the effect of temperature as the cooling of nearly 50°C produces a significant current reduction.

The average of the early data shows that no significant darkening took place in the initial weeks on orbit. However, contamination of the lenses from spacecraft (or array) outgassing may have occurred prior to the first readings.




Figure 17. Short Circuit Current Degradation

The initial indications were that the sensitivity to short term UV was small as expected. The longer-term effects of UV darkening and radiation darkening are less than expected. The trend says the degradations over time have not been as severe as expected, indicating that the expected values for all or most of the degradation factors (S/C outgassing, optical losses in lens from UV degradation, UV degradation of cell/cover, and cell radiation degradation) were conservative.

It was the intent to be conservative in each of these factors. For example UV light is refracted over the cell coverglass. Degradation of the cover adhesive would only occur during times of off-pointing.

The radiation degradation was based on an analysis by the Aerospace Corporation performed in August of 1997. The EOL degradation was forecast to be 0.965. A linear degradation with time was assumed although actual degradation from constant fluence follows a high order polynomial. In truth, a large solar flare could occur at any point in the mission. Use of a linear fit compensates partially for the possibility of a large solar flare early in the mission.

Since the October 1998 launch solar sunspot activity has been well below historical averages, as can be seen in Figure 18, below. Since flare activity and particle radiation have been correlated to sunspot activity it is plausible that the radiation degradation is proceeding at well below the modeled rate.

It also is possible that the some other factor in the model was conservative or the data itself is optimistic (meaning actually lower) due to systematic inaccuracies.




Figure 18. Historical Sunspot Activity

The entire IV curves for the eight modules showed no surprises. In general, the power output is much lower due to greater sun distance for the later two curves. Between day 213 and 282 the spacecraft traveled from 1.330 AU to 1.341 AU and back to 1.325 AU again. So the insolation level was essentially equal - as are the IV curves recorded - although 69 days have passed. The degradation of the lens appears smaller than is measurable.





Figure 19. Tap Module - Wing 1, Panel 1

SPEAK: The spacecraft electrical power system was designed with the capability to utilize the batteries as a buffer to allow maximum thruster output without collapse. A software routine, termed SPEAK, finds the solar array peak power for use as an input to controlling operation at a maximized thrusting level. This routine provides the best information on array level performance.

The first SPEAK sequence, on mission day 90, began by incrementally increasing the ion engine power level,

stopping at two intermediate points between nominal bus loads and maximum power, as shown in Figure 20. This was done to allow intermediate power level data to be recorded and to let the array cool to near the full power



operating temperature before moving to the full power load voltage.

Figure 20. SPEAK Sequence Flight Data

The last setting was chosen to be about 100 W in excess of the expected maximum (preflight) predicted power of the array. The battery was relied upon to supply the differential power. The spacecraft software will step down the thrust level if the battery discharge exceeds a predetermined level.

During this test run, the load voltage set point was intentionally stepped lower at 0.3 V increments to obtain the most detailed data. As can be seen in the closeup of the data near max power, Figure 21, when the array peak power was approached, the increments of voltage resulted in negligible array power output change.




Figure 21. SPEAK Sequence, Incrementing Voltage Near Peak Power

Plotting the wing currents against voltage rather than time, yields the more familiar "IV" curve. A detail of the data near the knee is shown in Figure 22.




Figure 22. SPeak Data, Array Output at 1.1185 AU

The power output has leveled off mostly, to 2084 watts, as the voltage was reduced to the last recorded value of 93.7 V. The flight values from the best fit curve are compared to the model predictions in Table 5.

Table 5. SPeak Results Comparison

Value	V_{MP}	P_{MAX}
Prediction	93.5	2094
Flight Results	93.7	2084
FR/P Ratio	1.003	0.995

The excellent agreement between the forecast and the first flight results is a clear validation of the *Scarlet* technology.

The full SPeak test was found to be too time-consuming and a shorter version evolved, termed mini-SPeak. Three mini-Speaks were performed in April of 1999 on the 2nd, 8th, and 22th. During each test run, the load voltage set point was stepped at 0.6 V increments and held for a number of hours to record a significant number of data points at a low data rate. An example of the data set is shown in Figure 23.




Figure 23. Mini-SPeak 1, Mission Day 160

Plotting the wing currents against voltage rather than time, puts the data into the more familiar "IV" curve presentation. In Figure 24 is the IV curve near the knee as this is the point of interest (and the limit of the data set).




Figure 24. Mini-SPeak 1, Peak IV Data

The flight values for Pmax and the associated voltage, determined from graphs like the previous example, are compared to the model predictions in Table 6.

No.	Date	AU	Predicted Power (W)	Predicted Voltage (V_{max})	Observed Power (W)	Observed Voltage (V_{max})
1	2-Apr-99	1.2676	1677	97.7	1628	99.16
2	8-Apr-99	1.2774	1653	98.0	1595	99.16
3	22-Apr-99	1.2978	1604	98.4	1538	99.76

Table 5. Mini-SPeak Results Comparison

A summary depiction of the flight data against the forecasts for power, peak power voltage, and temperature is shown in Figure 25. The flight data is summarized by month in charts in Appendix B. Hourly data is depicted for voltage, wing currents, and RTD temperatures (averaged for each of the four positions in module).




Figure 25. Original Forecasts and Flight Data

4. TECHNOLOGY VALIDATION SUMMARY

On DS1, the first mission of NASA’s New Millennium Program, the *Scarlet* array has provided power as designed, over 2.5-kW at nominally 90 volts, to power the NSTAR ion propulsion engine, plus spacecraft bus loads, on the interplanetary mission.

The accomplishments of the entire DS1 *Scarlet* program, in summary, were:

- Development, manufacture and successful spacecraft integration of a novel advanced, multi-junction-cell-based, high voltage concentrator solar array;
- Flight validation of new structural, mechanical, optical and electrical systems;
- Flight validation of the modeling and predictions for power output;
- Successful operation of the concentrator system on the DS1 spacecraft.

These accomplishments were achieved over 3 phases: Pre-flight: design, fabrication, test and integration, Launch and initial deployment, and the initial cruise phase of the mission. The specific successes of the *Scarlet* technology achieved on DS1 were:

- Demonstration of novel technology

- ✓ *The first successful concentrator array system in space*
- Utilizing the most advanced solar cell technology
 - ✓ *The first successful array system in space to use dual and triple junction solar cells*
- Demonstrating high specific power
 - ✓ *The first successful array system in space to use dual and triple junction solar cells*
- Demonstrating structural robustness
 - ✓ *The system stowed and deployed stiffness performance is excellent*
- Being designed for producibility
 - ✓ *This will allow for low recurring cost on future applications*
- Supporting all DS1 mission requirements
 - ✓ *The array performance provided the power needed to reach the encounter targets*

The risks that are inherent in the creation of such new technology are numerous and complex. A large portion of risk retirement was completed with thorough ground tests, but many elements of the technology simply couldn't be validated without the success of the launch, deployment, and the subsequent flight operations performed by JPL. The areas of highest remaining risk prior to launch were:

- **Structural or electrical damage during launch**
Broken lenses, loose wires or cracked cells
Insufficient tiedown preload
- **Deploying successfully in zero gravity**
Imbalance from thermal loading
Cable loads
Hingeline binding
Mechanism torque margin
Array jump-out loading
Damper failure due to deadband
Panel insert strength margin
Failure of the lens frame to deploy
- **Proper alignment to the sun**
Shifts in mountings from launch loads
SC to array mounting accuracy
Gimbal pointing accuracy
Hinge stop adjustment
Thermal warping
Warping of the power panels and/or lens panels
Moisture outgas shape change in composites
Alignment of module elements
Cells on carrier, carrier on substrate
Lens panel over substrate

Lens to lens spacing, lens shape

- **Providing power as expected**
 - AMO vs. LAPSS performance
 - Spectrum and collimation*
 - Failed circuits
 - Mechanical failure*
 - ESD induced failure*
 - Operating temperature
 - Panel conductivity*
 - Radiation exchange*
 - Lens absorption and transmission*
 - Environmental degradation
 - Particle radiation darkening of the lenses*
 - UV radiation darkening of the lenses*
 - Contamination from the ion engine or SC*
 - Combined effects*
 - Performance vs. AU
 - Temperature change*
 - Insolation effects*

While the telemetry for critical metrics such as power, temperature, deploy time, etc were documented to be as predicted, many of the risks listed above were impossible to cost-effectively monitor. From the proper performance of the system on mission the elimination of all these risks can be safely inferred.

The flight validation of *Scarlet* will allow a multitude of future users, particularly deep space science missions, mid-level orbit satellite applications, and communication constellations, to confidently baseline ABLE's product offering and garner the benefits of state of the art performance at a fraction of the cost of standard technology.

5. APPLICATION FOR FUTURE MISSIONS

The hardware demonstrated on DS1 is not the limit of this technology's promise. During the design and fabrication of the DS1 *Scarlet* array there was a concurrent review for potential improvements. A number of advancements were developed as part of ABLE's internal R&D effort. The most significant improvements which have been demonstrated to date are the elimination of the carrier used in cell laydown and of the lens clip stampings used in the lens panel assemblies. These improvements contribute equally to a 10% increase in specific power while reducing fabrication costs.

Additionally, a demonstration model has been built which shows the stowed height can be significantly reduced by nesting the lens panels into the substrate panels.

A common mission requirement, not required by the DS1 mission, is extended thermal cycling. The element of the hardware, peculiar to *Scarlet*, which is vulnerable to

thermal cycling stress - the lenses - was thermal cycled from -180 to + 110°C for a number of cycles equivalent to a 15-year GEO mission and from -160 to + 110°C for 40,000 cycles to represent a LEO mission.

NASA and BMDO also continue to develop technologies that will radically improve the performance of *Scarlet*. For example, the future development of advanced multiple band gap cells that will deliver 30 to 35% efficiency under concentration. Small cell size and low total PV area allow *Scarlet* to take advantage of such cells early in their production cycle when quantities are low and cost is high.

In addition, monolithic polymer concentrator lens materials that can survive both radiation and UV exposure are being demonstrated by an Entech/ABLE team. This will allow construction of much lighter foldable lenses and will lead to a huge gain in specific power and reduction of stowed volume.

For future applications the cost and performance advantages of *Scarlet* will be maximized for missions with high radiation or LILT conditions. As new exotic cells start to become available that promise large efficiency gains, but come with high initial production costs, the *Scarlet* technology can be effectively utilized. With a small size cell and a low total PV area a high power *Scarlet* array can be assembled.

6. ACKNOWLEDGMENTS

The triumph of DS1 and the *Scarlet* technology honors the efforts of dedicated engineers and technicians who produced a working system out of many novel technologies. Patience and commitment to problem solving of the team members at ABLE, BMDO, and JPL were all critical to the achievement. The effort to design and develop the *Scarlet* technology as described in this paper was performed under the sponsorship of the Ballistic Missile Defense Organization. Contributions by the Jet Propulsion Laboratory and the National Aeronautics and Space Administration were essential in mitigating risk and thereby producing a successful mission.

7. REFERENCES

1. Allen, D., Jones, P., Murphy, D., and Piszczor, M., "The SCARLET Light Concentrating Solar Array," Conference Record of the Twenty Fifth IEEE Photovoltaic Specialists Conference, IEEE, 1996.
2. Wachholz, J., and Murphy, D., "SCARLET I: Mechanization Solutions for Deployable Concentrator

Optics Integrated with Rigid Array Technology,” Proceedings of the Aerospace Mechanisms Conf., 1996.

3. Eskenazi, M., Murphy, D., Ralph, G., and Yoo, H., “Testing of Dual Junction SCARLET Modules and Cells Plus Lessons Learned,” Conference Record of the Twenty Sixth IEEE Photovoltaic Specialists Conf., IEEE 1998.
4. Eskenazi, M., Murphy, D., Anspaugh, B., Brinker, D., and O’Neill, M., “Balloon and Lear Jet Testing of SCARLET Modules and Cells,” Proceedings of the 15th Space Photovoltaic Research and Technology Conference (SPRAT XV), 1997.
5. Allen, Douglas M., and David M. Murphy, “An Update on the Deep Space 1 Power System: SCARLET Integration and Test Results”, 33rd IECEC Proceedings, Colorado Springs, CO, 2-6 August 1998, IECEC-98-403.
6. Krut, Dimitri D., Lovelady, James N., and Cavicchi, B. Terence, “Spectrally Balanced Light IV Testing of Production Dual Junction Space Cells”, 2nd World Conference and Exhibition on Photovoltaic Solar Energy Conversion, Vienna, Austria, 6-10 July 1998.

Appendix A

In the table below is shown the telemetry channels for each of the two technology validation activities conducted on the spacecraft for *Scarlet*. The sequence SIVPerf characterizes the IV curves of 8 individual modules on the array. The sequence SPeak characterizes the peak power point of the array by stepping down the solar array regulation voltage setpoint while the ion engine is thrusting at a level high enough to induce array regulation.

Channel	Mnemonic	SIVPerf	SPeak
P-2030	SA_V	X	X
P-2040	SA1_I	X	X
P-2050	SA2_I	X	X
P-2061	ESS_BUS_V		X
P-2060	ESS_BUS_I		X
P-2062	NON_BUS1_I		X
P-2063	NON_BUS1S_I		X
P-2064	NON_BUS2_I		X
P-2065	NON_BUS3_I		X
P-2011	BAT1_I		X
P-2021	BAT2_I		X
P-3170	SA_MOD_LDSEL	X	
P-3171	SA_MOD_SEL	X	
P-3172	ARR_OPV_SL_C		X
P-0020	BAT1_SOC		X
P-0022	BAT2_SOC		X
P-4041	SA1_VAL_TMP1	X	X
P-4042	SA1_VAL_TMP2	X	X
P-4043	SA1_VAL_TMP3	X	X
P-4044	SA1_VAL_TMP4	X	X
P-4045	SA1_VAL_TMP5	X	X
P-4046	SA1_VAL_TMP6	X	X
P-4047	SA1_VAL_TMP7	X	X
P-4048	SA1_VAL_TMP8	X	X
P-4051	SA2_VAL_TMP1	X	X
P-4052	SA2_VAL_TMP2	X	X
P-2031	SA_MOD_I_TLM	X	X
P-2032	SA_MOD_V_TLM	X	X
P-3081	SA_MOD_I	X	X
P-3082	SA_MOD_V	X	X
P-3170	SA_MOD_LDSEL	X	X
P-3171	SA_MOD_SEL	X	X
B-3101	SA_OP_PT_LSB		X
P-0006	ARR_OPV_SEL		X
B-2082	HCDRC0_STA		X
V-0141	XLINECUR		X
V-0142	XLINEVOL		X
V-3421	XTHRTLVL		X
V-3421	XTHRTLVL	X	
A-1401	ACMSUNBODY0	X	
A-1402	ACMSUNBODY1	X	
A-1403	ACMSUNBODY2	X	
A-1711	SADA_ANGLE_0	X	
A-1712	SADA_ANGLE_1	X	

APPENDIX B

Table B-1. Array Validation Activities Timeline

Mission Day	Activity	Date	DOY
7	SIV Perf Seq	31-Oct-98	304
8	SCal Seq	1-Nov-98	305
18	SIV Perf Seq	11-Nov-98	315
90	SPEak Seq	22-Jan-99	022
160	Mini SPEak	2-Apr-99	091
166	Mini SPEak	8-Apr-99	098
180	Mini SPEak	22-Apr-99	121
213	SIV Perf Seq	25-May-99	145
279	Mini SPEak	30-Jul-99	211
282	SIV Perf Seq	2-Aug-99	214
285	Mini SPEak	5-Aug-99	217
285	Mini SPEak	6-Aug-99	218
325	Mini SPEak	14-Aug-99	257
346	Mini SPEak	5-Oct-99	278

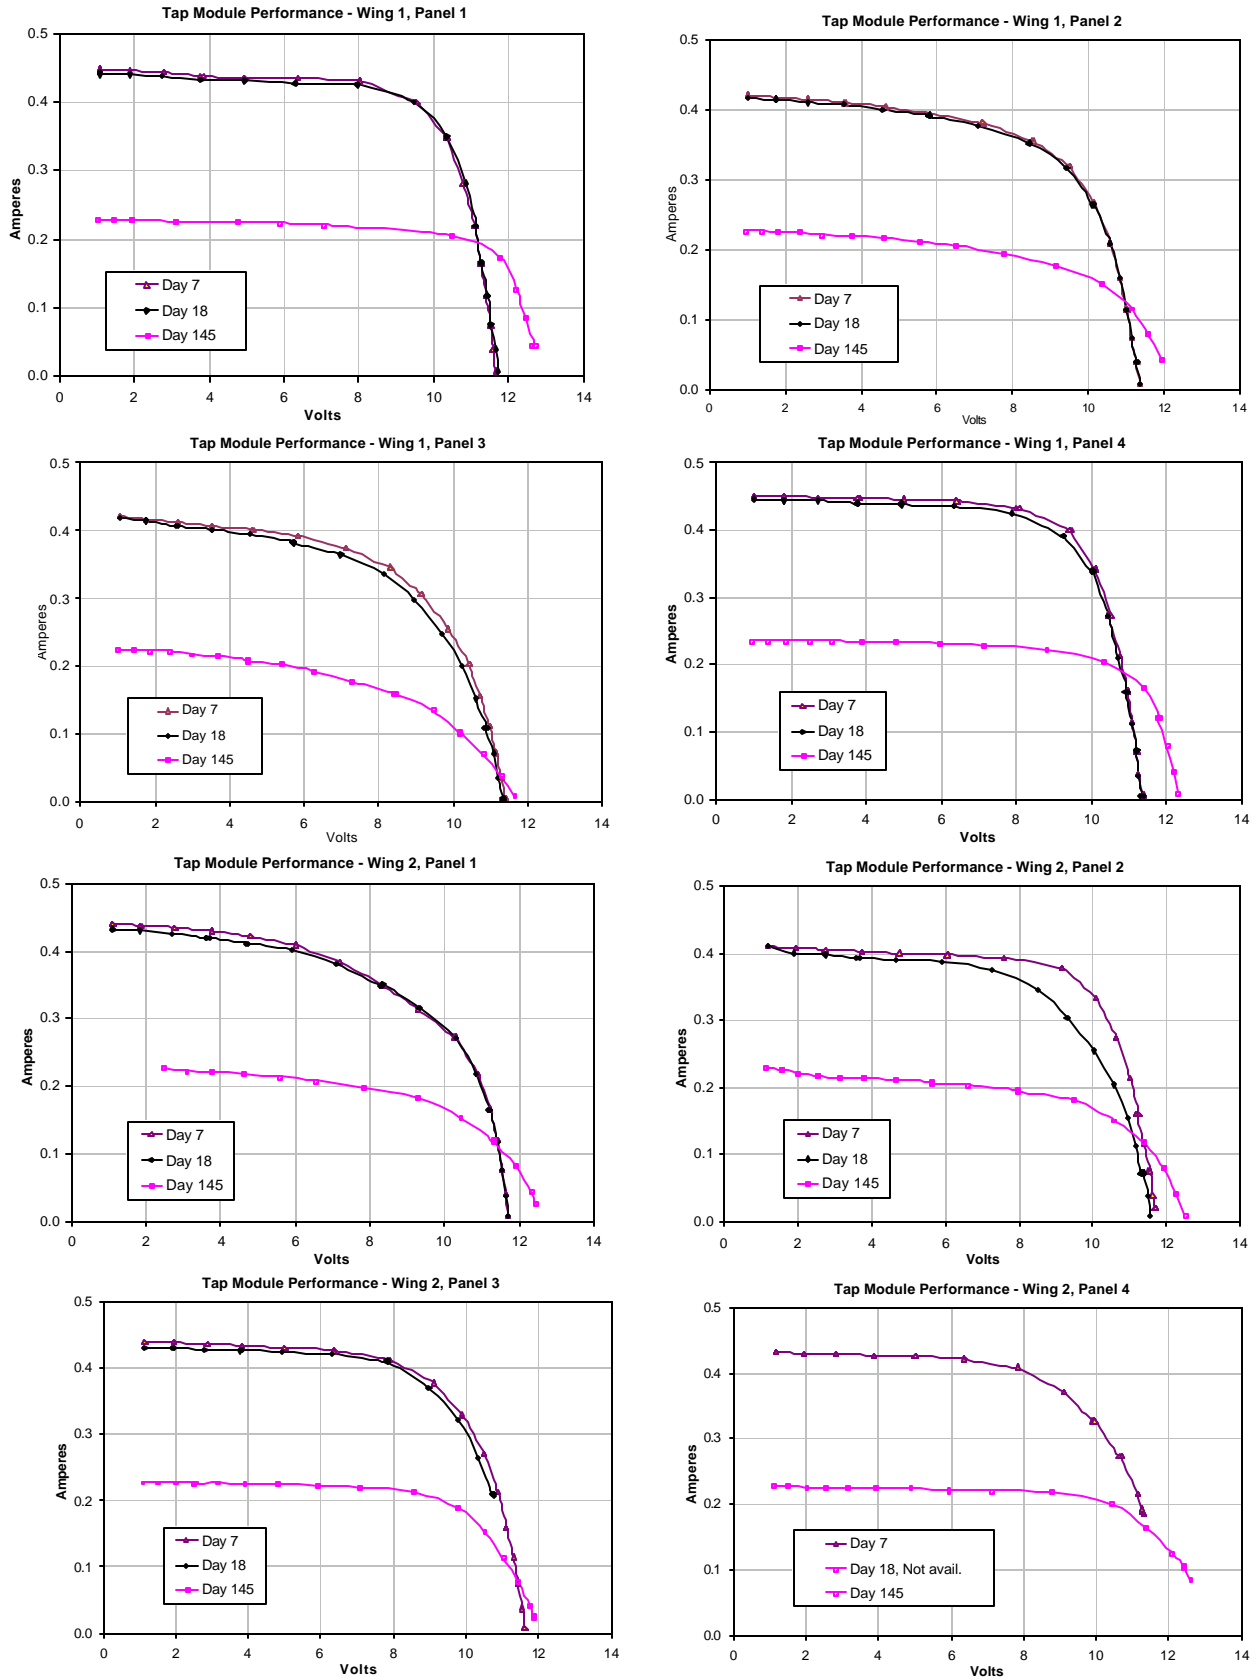


Figure B-1. SIVPerf Tap Module Data for Mission Days 7, 18 and 213 (DOY 145)




Figure B-2. DS1 SCARLETT Array Data Hourly: First Month of Mission

Figure B-3. DSI SCARLET Array Data Hourly: Second Month of Mission





Figure B-4. DSI SCARLET Array Data Hourly: Third Month of Mission






Figure B-5. DS1 SCARLET Array Data Hourly: Fourth Month of Mission

DS1 SCARLET Array Data Hourly: Fifth Month of Mission




Figure B-6. DS1 SCARLET Array Data Hourly: Fifth Month of Mission

DS1 SCARLET Array Data Hourly: Sixth Month of Mission




Figure B-7. DS1 SCARLET Array Data Hourly: Sixth Month of Mission




Figure B-8. DS1 SCARLET Array Data Hourly: Seventh Month of Mission

DS1 SCARLET Array Data Hourly: Ninth Month of Mission




Figure B-9. DS1 SCARLET Array Data Hourly: Ninth Month of Mission

DS1 SCARLET Array Data Hourly: Tenth Month of Mission





Figure B-10. DS1 SCARLET Array Data Hourly: Tenth Month of Mission



Small Deep Space Transponder (SDST) DS1 Technology Validation Report

Chien-Chung Chen, Shervin Shambayati, Andrew Makovsky,
F. H. Taylor, Martin I. Herman, Samuel H. Zingales
*Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California 91109*

Carl Nuckolls, Keith Siemsen
*Motorola
Government Space Systems
8201 E. McDowell Road
MD: H2242
P.O. Box 1417*



Table of Contents

<u>Section</u>	<u>Page</u>
Extended Abstract	iv
Small Deep Space Transponder (SDST) DS1 Technology Validation Report	1
Abstract.....	1
1.0 Introduction	1
2.0 Key SDST Functions.....	2
3.0 SDST Validation Objectives	3
3.1 Uplink Functions	3
3.2 Downlink Functions	3
3.3 Radio Metrics	4
3.4 Analog Engineering Telemetry Collection.....	4
3.5 Ka-band Readiness Demonstration	4
4.0 SDST Validation Process	4
4.1 Receiver Best Lock Frequency.....	5
4.2 Signal Acquisition Range and Rate.....	5
4.3 Self and False Lock.....	5
4.4 Uplink Command Reception.....	7
4.5 Uplink Power Measurements	7
4.6 Telemetry Encoding and Modulation.....	8
4.7 Noncoherent Mode Operation.....	9
4.8 Nonlinear Phase Modulator Performance	10
4.9 Coherent Doppler Tracking.....	11
4.10 Noncoherent Downlink Frequency Stability	12
4.11 Ranging Functional Verification	12
4.12 Beacon Tone Generation and Tone Tracking.....	16
4.13 External Telemetry Sampling Functions	16
5.0 DSN Ka-band Readiness Verification	16
5.1 34-m Antenna Pointing Performance	19
5.2 Ka-band System Noise Temperature and Link Capacity Projection.....	19
5.3 Ka-band Performance Threshold.....	22
5.4 X/Ka-band Radio Science	23
5.5 Ka-band Link Threshold at Low Bit Rate	23
5.6 Ka-band Antenna Pointing and Gravity Compensation at 70 m	23
6.0 Summary and Conclusion.....	24
7.0 List of References	25
8.0 Acknowledgments	25
Appendix A. List of Telemetry Channels and Names.....	26
Appendix B. Date of Turn-on/off and Frequency of Data Capture.....	26

Figures

<u>Figure</u>	<u>Page</u>
Figure 1. (a) Receiver Best Lock Frequency (BLF) Variation as a Function of Voltage-Controlled Oscillator (VCO) and Baseplate Temperature During Ground Testing, and (b) Measured In Flight.....	6
Figure 2. Measured SDST Sweep Rate at KSC Testing.....	7
Figure 3. Measured X- and Ka-band Carrier Power as a Function of Time During DOY 1998-344, Showing Various Ranging Carrier Suppression Values (Note: ERT = Earth-Received Time).....	11
Figure 4. (a) Measured Coherent Frequency Stability at X-band and (b) at Ka-band; and (c) Relative Stability of X/Ka-bands	13
Figure 5. (a) Measured Allan Deviations for X-band and (b) Ka-band Downlinks, and (c) Measured X/Ka-band Relative Stability	14
Figure 6. (a) Measured Noncoherent Downlink Carrier Stability at X-band and (b) at Ka-band.....	15
Figure 7. Predicted Versus Actual Pr/No for DOY 1999-035.....	16
Figure 8 (a–c) Ranging Residuals for DOY 1999-096 (4/6/99): (a) DSS 25, (b) DSS 34, and (c) DSS 46	17
Figure 8 (d) Ka-band Ranging Residuals for DOY 1999-096 (4/6/99).....	18
Figure 9. Ranging Residuals (Measurement Errors) After Accounting for Station and Spacecraft Motions.....	18
Figure 10. DSS-25 Conscan Pointing Residuals, Showing that Pointing Error is Generally Less than 4 millirads.....	20
Figure 11. DOY 1998-344 Ka-band Measured SNT and 50% Weather-Predicted SNT	21
Figure 12. Ka-band G/T Advantage Over X-band, DOY 1998-344, DS1 Track at DSS-25.....	21
Figure 13. Ka-band G/T Advantage Over X-band, DS1 Track, DOY 1999-035, DSS-25	22
Figure 14. Ka-band Threshold Measurement Data	23

Tables

<u>Page</u>	<u>Page</u>
Table 1. Comparison of SDST Mass and Power Consumption with those of Mars Pathfinder (MPF) Telecom Components with Equivalent Functions.....	1
Table 2. SDST Technologies and their Design Heritage.....	1
Table 3. SDST Validation Objectives	5
Table 4. SDST Command Rates Verified In Flight to Date.....	7
Table 5. Command Threshold Table as Predicted	8
Table 6. Uplink Residuals as Measured In Flight	8
Table 7. Telemetry Data Rates verified In Flight.....	9
Table 8. Measured Downlink Telemetry Residuals In Flight	10
Table 9. Encoding Modes/Data Rates Verified in Noncoherent Downlink Mode	10
Table 10. X-band Ranging Suppression Due to Nonlinear Phase Modulator, Measured Versus Predicted on the Basis of Pre-flight Data	12
Table 11. Measured Ranging Suppression at Ka-band	12

EXTENDED ABSTRACT

The small deep space transponder (SDST) is a Level-1 technology validation objective of the New Millennium Deep Space 1 (DS1) mission. The SDST was developed as a replacement for the Cassini deep space transponder (DST) and supports the radio frequency transmit, receive, and radio metric functions, as did previous transponders. Additionally, the SDST provides significantly greater functional integration by combining the command detection unit (CDU) and telemetry modulation unit (TMU) in one assembly. The integrated design allows for smaller size, mass, and power consumption of the telecom subsystem compared to the previous generation of hardware. Furthermore, the SDST is the first Ka-band capable deep space transponder. Previous Ka-band capable missions, such as Mars Observer (MO), Mars Global Surveyor (MGS), and Cassini, have relied on either an external frequency translator or a frequency multiplier to provide the Ka-band downlink. The SDST provides full support of Ka-band downlink functions, including telemetry modulation, and radio metrics (coherent Doppler, ranging, and differential one-way ranging [DOR]).

The development of the SDST was performed by Motorola Inc., Scottsdale, AZ, under funding from a JPL multimission consortium. Developed over a 3-year span at a cost of \$10.4 million (including non-recurring engineering and flight unit costs), the SDST development process is a model for the better-faster-cheaper development paradigm. Key technologies enabling the SDST design include: radio frequency integrated circuit (RFIC), advanced high frequency multichip modules (MCMs), and 70,000-gate complimentary metal oxide semiconductor (CMOS) application specific integrated circuits (ASICs), that implement the bulk of the receiver and telemetry modulation functions. Some of the design (down-conversion frequency scheme, dielectric resonator oscillators [DROs]) were derived directly from the Cassini DST, while others, such as the MCMs and ASICs, were new developments. The mixture of inherited technology and new development shortened the design cycle and lowered the development cost.

A high firmware content was implemented in the SDST's digital signal processing module, which was designed to work in X-band deep space, S-band Spaceflight Tracking and Data Network (STDN) facilities, and S-band Space Ground Link System (SGLS) transponders. The high firmware content enables many optional capabilities to be provided with only firmware changes, and allows specific tailoring for each mission. Particular attention was paid during development to ensure that the SDST provides flexible control in software. This feature was important for the multimission consortium, where different spacecraft

designs may dictate slightly different control interfaces. Transponder modes, such as the telemetry and ranging modulation indices, telemetry subcarrier frequency, and convolutional coding type, are user-controllable during mission operation. Other functions, such as the carrier-tracking loop bandwidth and automatic uplink acquisition, are firmware options. Furthermore, the SDST design accommodates interface with the spacecraft avionics via either a MIL-STD-1553, MIL-STD-1773, or RS422 serial bus, using the 1553 protocol. This design allows future flight users maximum flexibility in selecting the system architecture.

This report summarizes the results of DS1's in-flight technology validation activities related to the SDST. These activities were designed to show that the intended functions of the transponder can be achieved under the operating environment in space. Specific in-flight checkout activities were designed to exercise the transponder through different operating modes. Relevant performance data were collected both onboard by the flight system and on the ground by monitoring Deep Space Network (DSN) stations. Additional validation data were obtained through routine operations of the spacecraft by thoroughly monitoring the telecom-link performance and relevant SDST performance data. All SDST functions for uplink, downlink, and radio metric measurements were successfully validated, including the optional Ka-band downlink. In some cases, such as frequency stability measurements, the in-flight checkout activity also provided measurements of SDST performance in the actual operating environment not achievable with ground-based testing. Specifically, the in-flight technology validation activities focused on the following performance criteria:

Uplink:

- Uplink carrier receiver acquisition.
- Command data rate and command threshold.
- Carrier-tracking and uplink power measurements.

Downlink:

- Verification of telemetry encoding and carrier modulation.
- Verification of the transition between two-way coherent and one-way modes.
- Validation of the phase-modulator performance model.
- Validation of the Ka-band exciter technology and its associated performance characteristics.
- Validation of beacon tone generation.

Radio metrics:

- Measurement of the frequency stability of the DS1 auxiliary oscillator under in-flight temperature conditions.
- Verification of coherent carrier-tracking performance.

- Verification of the X/Ka-band relative carrier-tracking performance.
- Verification of the X/Ka-band ranging functions.

Although not strictly an SDST validation objective, the availability of a stable Ka-band downlink signal from DS1 permitted a direct verification of the Deep Space Network's operational readiness at Ka-band. The DS1 Ka-band downlink was used to:

- Demonstrate dual-band (X/Ka) end-to-end telemetry flow from a spacecraft to the DS1 Mission Support Area (MSA).
- Demonstrate the capability to generate necessary station predicts for Ka-band tracking.
- Demonstrate station capability to perform radio metric tracking (Doppler and ranging) on the Ka-band downlink.
- Verify X/Ka-band radio metrics performance.
- Measure Ka-band system noise temperature, which compares favorably with the model.
- Demonstrate DSS-25 capability to accurately point the 34-m antenna using blind pointing.

The in-flight checkout activities and ongoing flight validation of the SDST provided confidence in the transponder design. With successful flight validation and experience gained through mission operations, the risk of using the transponder design for future missions has been substantially reduced.

Subsequent to a successful DS1 flight validation, the design of the SDST has been enhanced to remove some of the operational idiosyncrasies due to the nonlinearity of the phase modulator and the changes in the receiver best-lock frequency. The current generation of SDST, scheduled to be flown on the Mars 01 and Space Infrared Telescope Facility (SIRTF) missions, has incorporated these changes. Furthermore, unlike the DS1 SDST, which functioned only with single-string command and data handing (C&DH), the Mars 01 SDST supports dual-string cross strapping with the C&DH subsystem. These performance improvements and this added functionality, together with DS1's in-flight validation, make use of the SDST truly low-risk for future flights.

Small Deep Space Transponder *Fact Sheet*

Key Features

- Deep Space Network Compatible
- X-band Receiver, X-band and Ka-band Exciters
- 2.5 dB Noise Figure (Nominal @25 o C)
- -156 dBm Receiver Threshold
- Temperature Compensated Receiver VCO
- Low Exciter Spurious, Phase Noise and Allan Deviation
- Radio Science Mode (USO Input Available)
- 40 ns Maximum Ranging Delay Vairation
- 3 ns Maximum Carrier Delay Variation
- Bus Interface - Mil-Std 1553/1773 Options
- External Power Converter Synchronization Capability
- Operates Under Launch Environments
- Radiation and SEU Resistant
- Internal Telemetry Modulation Encoder
- Internal Command Detector
- Mounting in Either of Two Axes

Performance Characteristics

Transponder

X-band Uplink Frequency Range	7.145–7.235 GHz
X-band Downlink Frequency Range	8.400–8.450 GHz
X-band Tx/Rx Ratio	880/749
Ka-band Downlink Frequency Range	31.800–32.300 GHz
Ka-band Tx/Rx Ratio	3360/749
Carrier Delay Variation	< 3ns p-p
Ranging Delay Variation	< 40 ns p-p

X-band Receiver

Noise Figure	<2.5 dB @ 25° C
Carrier Tracking Signal Range	-70 to -156 dBm
Carrier Loop BW (2-sided)	20 Hz nom. At threshold, expands to 200 Hz strong signal
Carrier Loop Damping Factor	0.5 @ 0 dB loop S/N
Tracking Range	>200 kHz about f0
Cmd Subcarrier Frequency	16 kHz
Cmd Subcarrier Mod Index	0.2–1.3 rad pk.
Ranging Filter Type	3-pole Chebychev
Ranging Filter BW	1700 kHz nominal
Temperature Stability	+/- 6.5 ppm (-40 to +50° C)

Exciters (X- and Ka-band)

X-band Output Power	+12 dBm @ 25° C
X-band Residual Phase Noise	-20 dBc/Hz at 1 Hz offset -80 dBc/Hz at 100-100 kHz
Ka-band Output Power	+4.0 dBm @ 25° C
Frequency Stability, 0 to +50° C	5.0 ppm
Spurious and Harmonic Outputs	<-50 dBc
Phase Mod Linearity	10% to 2.0 rad pk.
Tim Format	NRZ-L
Tim Convolutional Encoding	15-1/2, 15-1/4, 15-1/6, 7-1/2
Tim Subcarrier	Programmable, 2kHz to 4 MHz sq wave.
Tim Phase Deviation	0 to 90° peak
Ranging Modulation Index	Selectable, 2.1875, 4.375, 8.75, 17.5, 35° pk.
Differential One-way Ranging Tones	19.2 MHz, Coherent with carrier
Direct Modulation Mode	Available
Bi-φ -L Coding	Available



	SDST	Mars Pathfinder Equivalent
Mass	3 kg	TMU: 0.435 kg DST: 4.000 kg CDU: 0.365 kg Ka-band Exciter: N/A
Power	12.9 W	TMU: 1.4 W DST+CDU: 13.1 W

Small Deep Space Transponder (SDST) DS1 Technology Validation Report

*Chien-Chung Chen, Andrew Makovsky, Shervin Shambayati, F. H. Taylor, Martin I. Herman,
Sam H. Zingales, Keith Siemsen (Motorola), Carl Nuckolls (Motorola)
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California*

ABSTRACT

This report summarizes the in-flight technology validation results for the small deep space transponder (SDST). Specific in-flight checkout activities were designed to exercise the transponder through different operating modes; relevant performance data were collected both onboard by the flight system and on the ground by monitoring Deep Space Network (DSN) stations. Additional validation data were obtained through routine operations of the spacecraft by thoroughly monitoring the telecom-link performance and relevant SDST performance data. All SDST functions for uplink, downlink, and radio metric measurements were successfully validated under the intended operating environment, including the optional Ka-band downlink.

1.0 INTRODUCTION

The small deep space transponder (SDST) is a Level-1 technology validation objective of the New Millennium Deep Space 1 mission (DS1). The SDST was developed as a replacement for the Cassini deep space transponder (DST) and supports the radio frequency transmit, receive, and radiometric functions, as did previous transponders. Additionally, the SDST provides a significantly greater functional integration by combining the command detection unit (CDU) and telemetry modulation unit (TMU) in one assembly. The integrated design allows for smaller size, mass, and power consumption of the telecom subsystem compared to the previous generation of hardware. A comparison of mass and power consumption of the SDST with the Mars Pathfinder (MPF) telecom subsystem is shown in Table 1.

Table 1. Comparison of SDST Mass and Power Consumption with those of Mars Pathfinder (MPF) Telecom Components with Equivalent Functions

	DS1	Mars Pathfinder (equivalent function)
Mass	3 kg	TMU: 0.435 kg DST: 4.000 kg CDU: 0.365 kg
Power	12.9 W	TMU: 1.4 W DST+CDU: 13.1 W

The development of the SDST was performed by Motorola Inc., Scottsdale, AZ, under funding from a JPL multimission consortium. Developed over a three-year span at a cost of less than \$10.4 million (including nonrecurring engineering (NRE) and flight unit costs), the SDST development process is a model for the better-faster-cheaper development paradigm.

Key technologies enabling the SDST design include the radio frequency integrated circuit (RFIC), advanced high-frequency multichip modules (MCMs), and 70,000-gate complimentary metal-oxide semiconductor (CMOS) application specific integrated circuits (ASICs), that implement the bulk of the receiver and telemetry modulation functions. Some of the designs (downconversion frequency scheme, dielectric resonator oscillators (DROs)) were derived directly from the Cassini DST, while others, such as the MCMs and ASICs, were new developments. This mixture of inherited technologies and new developments shortened the design cycle and lowered development costs. A summary of key SDST technologies and their design heritage is shown in Table 2.

A high firmware content was implemented in the SDST's digital signal processing module, which was designed to work in X-band deep space, S-band NASA Spaceflight Tracking and Data Network (STDN) facilities, and S-band USAF Space Ground Link System (SGLS) transponders. The high firmware content enables many optional capabilities to be provided with only firmware changes, and allows specific tailoring for each mission. Particular attention was paid during development to ensure that the SDST provides flexible control in software. This feature was important for the multimission consortium, where different spacecraft designs may dictate slightly different control interfaces. Transponder modes, such as the telemetry and ranging modulation indices, telemetry subcarrier frequency, and convolutional coding type, are user-controllable during mission operation. Other functions, such as the carrier-tracking loop bandwidth and automatic uplink acquisition, are firmware options. Furthermore, the SDST design accommodates interface with the spacecraft avionics via either a MIL-STD-1553, MIL-STD-1773, or RS422 serial bus, using the 1553 protocol. This design allows future flight users maximum flexibility in selecting the system architecture.

Table 2. SDST Technologies and their Design Heritage

Key Technologies	SDST Heritage
Frequency scheme	Cassini
Dielectric resonator oscillators (DROs)	Cassini (smaller)
DRO lock technique	New (sampling phase detectors (SPDs))
Ceramic first intermediate frequency filter	Cassini
Preselector	Cassini
Voltage-controlled oscillator (VCO) and Auxiliary Oscillator (AuxOsc)	Cassini
Ka-band multiplier	JPL heritage
Low-noise RFIC	Motorola heritage
Power supply design	Cassini
RFIC phase modulator	New (JPL small business innovative research (SBIR))
RF board manufacturing technique	Duroid boards bonded (not fused)
Low-temp cofired ceramic MCMs	Motorola heritage
Command and control interface	1553, 422, 1773
Uplink/downlink interface	Cassini
Mechanical packaging	Cassini

The capabilities of the SDST include:

- X-band receiver/downconverter capable of carrier tracking at or below -156 dBm.
- Command detector unit function.
- Telemetry modulation function.
- X- and Ka-band exciters.
- Beacon mode operation.
- Coherent and noncoherent operation choice.
- X- and Ka-band ranging.
- Differential one-way ranging (DOR) for both X-band and Ka-band.
- Command and Data Handling (C&DH) communication via 1553.
- Data interface via RS422.
- External ports for temperature sensors.
- External port for an analog signal.

All SDST functional capabilities were verified on the DS1 mission, including the optional Ka-band downlink. This report summarizes the results of DS1's technology validation activities related to the SDST. With successful flight validation and experience gained through mission operations, the risk of using the transponder design for future missions has been substantially reduced. Indeed, the

SDST is currently in full production for the Mars 2001 and Space Infrared Telescope Facility (SIRTF) missions.

2.0 KEY SDST FUNCTIONS

The SDST is the first deep space transponder using digital receiver technology. The use of digital technology allows for tighter integration of functions and more flexibility in their control. Additionally, the SDST is the first Ka-band-capable deep space transponder. Previous Ka-band-capable missions, such as Mars Observer (MO), Mars Global Surveyor (MGS) and Cassini, rely on either an external frequency translator or frequency multiplier to provide the Ka-band downlink. The SDST provides full support of Ka-band downlink functions, including telemetry modulation and radio metrics (coherent Doppler, ranging, and DOR).

The design of SDST supports the following functions:

1. Uplink-related functions:
 - Receive and demodulate the X-band uplink carrier.
 - Monitor for self or false lock.
 - Provide an uplink automatic gain control (AGC) for receiver power measurement.
 - Receive and demodulate the command subcarrier and data stream.
2. Downlink-related functions:
 - Provide the capability of a noncoherent downlink with auxiliary oscillator or ultrastable oscillator (USO).
 - Perform convolutional encoding and subcarrier modulation of downlink telemetry.
 - Perform X- and Ka-band carrier modulation of downlink with variable modulation indices.
 - Provide independent control of X- and Ka-band downlinks.
 - Provide differential one-way ranging (DOR) modulation on downlink.
 - Generate a beacon tone.
3. Radio metrics:
 - Provide stable one-way downlink for use when the transponder is not in lock with the uplink.
 - Support two-way coherent operations by phase locking downlink with the uplink signal carrier.
 - Demodulate uplink ranging modulation and remodulate ranging signals on the downlink.
4. Collect analog engineering status within the subsystem

A summary of SDST functions and relevant requirements can be found in the SDST detailed functional specifications [1].

3.0 SDST VALIDATION OBJECTIVES

The SDST design has been subjected to a series of verification and validation tests before and after launch. Before launch, the SDST was subjected to a series of functional verification tests. These tests were intended to verify functional specifications and performance requirements. Additionally, continuous checkout and monitoring of transponder performance throughout the integration and test (I&T) process ensured that the performance and functional specifications of the transponder were met.

In contrast to the verification tests, the technology validation activities were designed to ensure that the intended functions of the transponder could be achieved by the design. This was achieved through a series of Deep Space Network (DSN) compatibility tests on the ground, several planned in-flight checkout (IC) activities, and monitoring of transponder/downlink performance throughout normal mission operations. The DSN compatibility tests were conducted using the Compatibility Test Trailer (CTT) at Motorola and at the Kennedy Space Center (KSC). Additional compatibility tests were conducted at JPL using the DSN Development and Test Facility. These tests validated the Level-3 system requirements to ensure flight-ground compatibility and key functions of the telecommunications subsystem. The results of the testing are summarized in a DSN compatibility test report [2].

After launch, the technology validation activities were designed to show that the intended functions of the transponder could be achieved by the design under a relevant operating environment. To that end, several in-flight checkout (IC) activities were planned specifically to verify and validate that the SDST reliably performed its required uplink, downlink, and radio metric functions with the tracking stations. In some cases, such as frequency stability measurements, the in-flight checkout activity also provided measurements of SDST performance in the actual operating environment, measurements not obtainable through ground-based testing.

The objectives of flight validation tests for each of the uplink, downlink, and radio metric functions are summarized below.

3.1 Uplink Functions

The receiver receives and demodulates X-band uplink. The SDST implements a hybrid analog-digital receiver. The uplink signal is first passed through the downconverter stages. The receiver also performs the wide-band AGC function. The downconverted intermediate-frequency signal is then digitized at a rate of $4/3 F_1$ (approximately 12.6 megahertz). The rest of the receiver functions are implemented in the digital ASIC, which includes the

narrow-band AGC, the carrier demodulation, and the command data demodulation functions. The digital receiver also derives the phase error between the receiver voltage-controlled oscillator (VCO) and the incoming radio frequency (RF) carrier. This error signal is then filtered and used to drive the VCO to close the carrier phase-tracking loop.

Once the carrier signal is demodulated, the command subcarrier synchronization and demodulation is performed by the command detector unit (CDU) within the digital ASIC. The SDST CDU uses a digital implementation similar to the Cassini/Mars Observer CDU. The CDU outputs the command data, clock, and a lock-detect indicator to allow for subsequent decoding of the command uplink by the spacecraft avionics.

In-flight validation objectives related to the uplink functions include validation of the following functions:

- Uplink carrier receiver acquisition.
- Command data rate and command threshold.
- Carrier-tracking and uplink power measurements.

3.2 Downlink Functions

The SDST contains two independently controllable exciters: one for X-band downlink and one for Ka-band downlink. These two downlinks are provided with independent subcarrier generator and convolutional encoder and can be configured to transmit independent downlinks. For the DS1 SDST, the X-band and Ka-band share common telemetry and clock inputs (since there is only a single-string avionics) and the two streams are configured for the same encoding rate.

In-flight validation objectives for the downlink functions include:

- Verification of telemetry encoding and carrier modulation.
- Verification of the transition between two-way coherent and one-way modes.
- Validation of the phase modulator performance model.
- Validation of the Ka-band exciter technology and its associated performance characteristics.
- Validation of beacon tone generation.

The phase modulator performance model is particular to the DS1 SDST, which exhibited nonlinear phase modulation characteristics under test. The nonlinearity results in a large intermodulation loss when both ranging and telemetry modulations are applied. A nonlinear loss model was constructed prior to launch using ground-test data; in-flight validation of the phase modulator performance model verified the validity of the performance model.

3.3 Radio Metrics

The SDST supports radio metric functions by providing two-way coherent transponding of the uplink carrier and by providing the turn-around ranging capability. These functions are similar to previous deep space transponders except, of course, that the SDST supports radio metric measurements in both the X-band and the Ka-band.

In-flight validation objectives for the radio metric functions include:

- Measurement of the frequency stability of the DS1 auxiliary oscillator under in-flight temperature conditions.
- Verification of coherent carrier-tracking performance.
- Verification of the X/Ka-band relative carrier-tracking performance.
- Verification of the X/Ka-band ranging functions.

3.4 Analog Engineering Telemetry Collection

In addition to reporting the internal status of the SDST, the external analog telemetry interface built into the SDST is also used to collect external analog engineering status from the telecom subsystem. Four (4) analog voltages and four (4) external temperature sensor interfaces are provided by the SDST. For the DS1 radio frequency subsystem (RFS), these input channels are mapped to the following SDST analog measurement channel assignments:

<u>Ext channel</u>	<u>Measurements</u>
1	Ka-band power amplifier (KAPA) input power monitor
2	KAPA output power monitor
3	X-band power amplifier (XPA) input power monitor
4	Detector amplifier module (DAM) secondary voltage

<u>Ext temp sensor</u>	<u>Location</u>
1	DAM temperature
2	SDST sidewall temperature
3	KAPA input detector temp
4	KAPA output detector temp

Collection of these engineering telemetry values, especially those related to the Ka-band power amplifier, were intended to support the KAPA technology validation activity, which will be described in a separate report.

3.5 Ka-band Readiness Demonstration

Although not strictly an SDST validation objective, the availability of a stable Ka-band downlink signal from DS1 permitted a direct verification of the Deep Space Network's operational readiness at the Ka-band. The DS1 Ka-band downlink was used to:

- Demonstrate dual-band (X/Ka), end-to-end telemetry flow from a spacecraft to the DS1 Mission Support Area (MSA) (DS1-g).

- Demonstrate the capability to generate necessary station predicts for Ka-band tracking.
- Demonstrate the station capability to perform radio metric tracking (Doppler and ranging) on the Ka-band downlink.
- Verify X/Ka-band radio metrics performance.
- Demonstrate the Deep Space Network Station 25 (DSS-25) capability to accurately point the 34-m antenna using blind pointing.
- Measure the Ka-band system noise temperature, which compares favorably with the model.

Additionally, the DS1 Ka-band downlink was used to support characterization of 70-m antenna pointing accuracy at the Ka-band.

4.0 SDST VALIDATION PROCESS

For in-flight checkout of the transponder, a series of validation objectives were identified. Each SDST validation objective, summarized individually in Table 3, requires the active participation of the ground tracking station. Depending on the purpose of the activity, the station provided an X-band uplink carrier and received either an X-band downlink or simultaneous X-band and Ka-band signals from the spacecraft.

In most SDST flight-validation activities, the power level of the carrier or the signal-to-noise ratio (SNR) of the command, telemetry, or ranging data was collected and compared with the values predicted on the basis of the DS1 communications link models. Some of the data are available from SDST (for example, uplink-related measurements) while others are available through DSN station monitoring.

Since DS1 supports both Ka-band and X-band downlinks, a significant portion of the validation activities need to be performed at both the Ka-band and the X-band. However, the Ka-band horn antenna on DS1 is directive and must be pointed to Earth in order to conduct Ka-band related validation activities. The pointing constraints of the spacecraft, therefore, limit the times at which Ka-band activities can take place. During the initial checkout phase, the Miniature Imaging Camera and Spectrometer (MICAS) pointing constraints resulted in delaying the Ka-band related activities until 25 days after launch (L+25D). In contrast, X-band validation activities can be conducted using the X-band low-gain antenna (LGA) and are not constrained by spacecraft attitude. A second operational constraint on Ka-band activities is that the Goldstone tracking complex has the only DSN station (DSS-25) capable of receiving Ka-band transmissions. Therefore, technology validation tests involving Ka-band downlinks were conducted over Goldstone sites only.

Table 3. SDST Validation Objectives

Objectives	Pre-launch	In-flight Checkout	Tests
Receiver best lock frequency	Measure	Validate	Routine ops
Signal acquisition range and rate	Measure	Validate	Routine ops
Self/false lock characterization	Measure	Validate	Routine ops
Uplink command reception	Measure	Validate	Routine ops
Uplink power measurements	Characterize	Validate	Routine ops
Telemetry encoding and modulation	Test	Validate	Routine ops, Xtlm
Noncoherent mode operation	Test	Validate	Routine ops
Phase modulator performance	Characterize	Validate	Routine ops, Xrange
Noncoherent carrier frequency stability	Test	Measure	Xstable
Coherent Doppler tracking performance	Test	Validate	Routine ops
Ranging functional verification	Test	Validate	Xrange Krange
Beacon mode (a separate experiment)	Test	Validate	Xtone
Analog engineering telemetry sampling	Test	Validate	Routine ops

4.1 Receiver Best Lock Frequency

A predictable best lock frequency (BLF) is important for deep space mission operations. An accurate receiver BLF predict would allow the ground station to provide an uplink acquisition sweep over a sufficiently narrow range to rapidly acquire the SDST. During ground testing, it was discovered that the transponder’s best lock frequency is a sensitive function of temperature. Over the in-flight allowable range, the receiver BLF can vary by as much as ± 25 kHz from its predicted frequency profile. This fact was verified with in-flight measurement (see Figure 1). Subsequent development of the SDST for Mars 2001 missions has significantly reduced the amount of BLF drift compared to that of the DS1 SDST.

4.2 Signal Acquisition Range and Rate

Even though ground/flight testing of the transponder showed significant BLF variation, ground testing of the SDST also showed that the transponder can be acquired at a much higher rate than could previous transponders. Shown in Figure 2 is a plot of the acquisition rate as a function of uplink power level measured during the final DSN compatibility test at KSC prior to launch.

Based on the test data, it was recommended that a frequency sweep range of ± 30 kHz be used. A sweep rate of 900 Hz/s was recommended at power levels above -130 dBm and a sweep rate of 300 Hz/s was recommended at power levels less than -130 dBm. The combined sweep rate/range resulted in a sweep-acquisition time of less than 400 seconds in the worst case, and 130 seconds at higher power levels.


The SDST acquisition performance was validated on every track where coherent downlink is required. Over the mission lifetime of a year, there has been no uplink acquisition failure due to the transponder.

4.3 Self and False Lock

It is important that the receiver exhibit no self or false lock events. The absence of such events is critical for successful mission operations because false/self lock can prevent the receiver (SDST) from receiving a valid command uplink, rendering the spacecraft uncommandable.


The SDST provides frequent updates of its internal state. This status information is available to the spacecraft through the 1553-bus transactions for health monitoring purposes. Additionally, the SDST provides an event counter that registers every change in state of the SDST. Should an unexpected change of state occur, the event counter will advance incrementally. An unexpected lock-up and subsequent drop-lock of the carrier, for example, will advance the receiver event counter by 3 increments.

The SDST event counter was closely monitored throughout the IC period. During that period, an attempt was made to correlate incremental changes in the event counter with identifiable state changes. No self/false lock events or unexpected state changes were detected for the SDST during either ground testing or in-flight operations.



(a)

DS1 sdst receiver spe vs VCO temperature, as of April 28, 1999



(b)

Figure 1. (A) Receiver Best Lock Frequency (BLF) Variation as a Function of Voltage-controlled Oscillator (VCO) and Baseplate Temperature During Ground Testing (A), and (b) Measured In Flight




Figure 2. Measured SDST Sweep Rate at KSC Testing

4.4 Uplink Command Reception

The SDST was required to be commandable at each of the following data rates: 2000, 1000, 500, 250, 125, 62.5, 31.25, 15.625, and 7.8125 b/s. Uplink command reception at these rates was to be demonstrated by X-band uplink (XUPL) testing, as well as by routine commanding during operations. Due to project time constraints, XUPL was not completed. Instead, routine commanding was performed at all data rates except 31.25 b/s. The data rate during routine mission operations was selected based on the supportable uplink rate during a particular day (based on the spacecraft orientation, the antenna selected, the Earth-spacecraft range, and the ground transmitter power). Since DS1 mission operations require the spacecraft to employ different antennas at different spacecraft orientations, most of the required SDST data rates were verified in flight (see Table 4).

In-flight verification of command threshold was not performed because of lingering project concern about deliberately sending below-threshold command data to the flight software. The possibility of a resulting spacecraft safing could not be ruled out and the mission timeline had no margin to recover from safing. Although no command tests have been done, these command thresholds have been verified indirectly: during flight operations, the predicted command rate was always successful. Shown in Table 5 are the predicted thresholds that are used routinely to determine what command rate can be used on a particular DSN pass. The successful uplink activities provide indirect

confirmation that the receiver noise floor had not been degraded.

Table 4. SDST Command Rates Verified In Flight to Date

Command Bit Rate, b/s	Verified in Flight Operations?	Operational Signal Level, Pt, dB (1 mW)
2000	Yes	-114
1000	Yes	-124
500	Yes	-120
250	Yes	-131
125	Yes	-128
62.5	Yes	-132
31.25	Not yet	N/A
15.625	Yes	N/A
7.8125	Yes, used when recovering from fault protection	-140

4.5 Uplink Power Measurements

Uplink carrier threshold was indirectly verified using the uplink residuals measurements: the SDST measures the uplink signal-to-noise ratio and telemeters the measured data (carrier lock accumulator). This data is then compared to the predicted uplink carrier power and predicted system noise temperature of the receiver. The results provide an indirect confirmation of the receiver sensitivity and carrier threshold.

Table 5. Command Threshold Table as Predicted

Command Bit Rate, b/s	Mod Index, Radians	Uplink Carrier Suppression, dB	Threshold Pt/No (Ranging OFF), dB-Hz	Threshold Pt/No (3 dB Uplink Ranging Suppression), dB-Hz
2000	1.2	-3.5	47.55	50.6
1000	1.2	-3.5	44.3	47.3
500	1.2	-3.5	41.2	44.2
250	1.2	-3.5	38.5	41.5
125	1	-2.3	36	39
62.5	1	-2.3	32.7	35.7
31.25	1	-2.3	30	33
15.625	0.9	-1.9	27.5	30.5
7.8125	0.8	-1.4	26.2	29.2

Link residuals may be due to a modeling error (antenna gain or system noise temperature), operating conditions (spacecraft deadbanding), or changes in system performance. Shown in Table 6 are the uplink residual data compiled for passes when the high-gain antenna (HGA) was in use and when the spacecraft was Earth-pointed (in order to eliminate uncertainties due to spacecraft attitude). It is seen that the uplink residual is in reasonable agreement with the prediction. The larger standard deviation (two sigma is 1.2 dB) shows that the project should plan its link capability based on a link margin of at least 2 dB (3 sigma).

The uplink residuals served to provide only indirect verification of the SDST uplink threshold. Ongoing activities to monitor the uplink residuals will be required to monitor for long-term trend.

4.6 Telemetry Encoding and Modulation

The SDST is designed to support multiple telemetry encoding modes using an externally supplied data stream and clock signal up to 4.4 megasymbols per seconds. The external clock signal supplied needs to be coherent with the data stream and at a rate equal to the symbol coding rate selected (e.g., at multiples of the data rate). Additionally, the SDST supports both subcarrier modulation and direct carrier modulation (see SDST specifications [1]).

Full validation of telemetry encoding and modulation mode was not performed due to configuration limits of the spacecraft and DS1 downlink strategy. The available clock rate from avionics (the Reed Solomon downlink (RSDL) ASIC) supports only clock rates that are 1x, 2x, and 6x the data rate. Additionally, DS1's flight avionics system (hardware plus software) supports a maximum telemetry data rate of only 19908 b/s. The downlink strategy for DS1 requires that (7,1/2) and (15,1/6) codes be supported for subcarrier modulation mode only (no direct carrier modulation). The (7,1/2) code was used during initial acquisition (2100 b/s) and when the spacecraft was in one of the several standby modes. Most of the mission was conducted using the (15,1/6) code.

Table 6. Uplink Residuals as Measured In Flight

Time	DSS	Uplink Residual (Actual-Predict)	Spacecraft Antenna
1999-009 02:25-07:59	25	+0.8	HGA
1999-009 17:00-010 02:09	65	+0.8	HGA
1999-012 17:40-013 00:09	65	+0.7	HGA
1999-013 16:55-23:36	65	+0.7	HGA
1999-014 16:55-015 03:44	65	-0.2	HGA
1999-016 3:25-07:39	15	+0.7	HGA
1999-016 16:40-017 00:29	65	+0.3	HGA
1999-017 00:25-05:44	15	+0.7	HGA
1999-017 16:40-018 00:44	65	+0.3	HGA
1999-018 00:25-09:29	15	+0.3	HGA
1999-019 02:55-07:44	25	-1.1	HGA
1999-021 03:10-07:34	25	-0.6	HGA
1999-022 01:10-09:59	15	+0.7	HGA
1999-022 10:55-15:44	34	-0.3	HGA
1999-022 19:40-023 00:24	54	+1.7	HGA
1999-023 16:25-23:44	65	0.	HGA
1999-024 02:55-11:44	25	+0.03	HGA
1999-024 19:40-025 00:09	54	+0.13	HGA
1999-025 03:25-11:14	25	+0.15	HGA
Average		0.3 dB	
Standard deviation		0.6 dB	

4.6.1 *Telemetry Data Rate Verification (X- and Ka-band, Coherent Mode)*—Telemetry encoding and modulation was verified for SDST using the 19 planned data rates for both X-band and Ka-band downlinks at the planned encoding modes (Table 7). The activity (Xtlm) was conducted when the SDST was operating in the two-way coherent mode. The SDST provided a convolutionally encoded telemetry data stream at the symbol rate (either (7,1/2) or (15,1/6)) and modulated the symbol stream onto the required subcarrier (either a 25-kHz or 375-kHz square wave). Finally, the SDST modulated this subcarrier plus data onto the RF carrier (either X-band or Ka-band) at the desired modulation index.

All the data rates and both convolutional codes have been validated at X band, not only during Xtlm, but during routine operations at many data rates. At each planned operating rate, the ground station successfully locked onto and decoded the telemetry data stream (at both the X-band and Ka-band) and transmitted the decoded telemetry stream to the DS1 MSA.

Table 7. Telemetry Data Rates Verified In Flight

Data Rate	X Mod Index (DN), Ranging ON	Ka (DN), Ranging ON	Convolutional Code
19908	38	54	(15,1/6)
13272	38	53	(15,1/6)
9480	38	Not used	(15,1/6)
6636	38	Not used	(15,1/6)
4424	38	Not used	(15,1/6)
3150	38	Not used	(15,1/6)
2100	38	40	(15,1/6)
1422	38	45	(15,1/6)
1050	38	44	(15,1/6)
790	38	43	(15,1/6)
600	38	42	(15,1/6)
420	38	41	(15,1/6)
300	37	39	(15,1/6)
200	36	38	(15,1/6)
150	36	37	(15,1/6)
79	33	33	(15,1/6)
40	30	27	(7,1/2)
10	23	16	(7,1/2)

4.6.2 *X-band Telemetry Link Performance (Link Residuals)*—In addition to verifying that the SDST can effectively modulate the downlink, the X-band downlink performance has also been verified by tracking the link residuals over multiple passes. Shown in Table 8 are the X-band downlink performance values versus the expected downlink signal values (carrier power and symbol SNR) measured using the block-V receiver (BVR). Spacecraft deadbanding (an attitude control error that varies between ±1 degree) can result in a significant degradation of the

downlink (as much as several tenths of dB). This deadband effect has been removed from the data by using the peak signal level for each pass. However, bad weather—system noise temperature variation—has not been taken into account.

The average symbol SNR (SSNR) residual is comparable to the carrier power residual (+0.5 dB). Furthermore, when adjusted for system noise temperature (SNT), the residual is only 0.1 dB. This indicates that the link model (total power, modulation index, as well as downlink signal quality) is sufficiently accurate. The measured residual spread (0.4 dB, one sigma) with SNT and spacecraft deadband effects removed provides a measurement of the uncertainty in link performance. These data are useful for future missions and can be used to estimate the effective link margins required.

4.7 *Noncoherent Mode Operation*

The DS1 SDST typically operates in the coherency-enabled mode with downlink driven by a VCO. When no uplink signal is detected (no receiver lock) or when the SDST is configured for coherency-inhibited mode (two-way noncoherent mode), the downlink is driven by an auxiliary oscillator (AuxOsc). Validation of noncoherent mode operation must:

- a. Validate that the SDST can successfully generate a noncoherent downlink signal driven by the AuxOsc. The SDST was commanded to the noncoherent mode during initial acquisition and standby modes and during certain technical validation activities (like Xstable and beacon mode testing).
- b. Validate that the SDST can generate a noncoherent downlink with telemetry modulation. This is the standard operating mode at the beginning of any station pass that does not overlap a previous pass. The station is usually able to acquire one-way downlink telemetry before it locks to the two-way downlink a round trip light time later. Data rates verified during IC activities are shown in Table 9.
- c. Validate that the transponder can be successfully commanded out of a coherency-inhibited (two-way non-coherent (TWNC)) mode. Although DS1’s standard operating mode is coherency-enabled, the transponder was intentionally set to operate in TWNC mode during launch and when the spacecraft enters standby mode. This is so that there will be a detectable downlink signal even if there is a problem with the uplink. Since launch, the spacecraft has entered standby mode at least six times, and every time the spacecraft was successfully commanded to return to the normal (coherency-enabled) mode.

Table 8. Measured Downlink Telemetry Residuals In Flight

Day of Year and Time	DSS	SSNR	SSNR	SSNR Delta	Pc	Pc	Pc Delta	SNT	SNT	SNT	Adjusted for System Noise Temp (SNT)		
		Actual	Predicted	Actual-Pred	Actual	Pred	Actual-Pred	Actual	Pred	dB Delta	SSNR Delta	Pc Delta	Spacecraft Antenna
1999-009 02:15–07:59	25	7.85	7.15	0.7	-131.2	-131.7	0.5	30	32.5	0.3	0.4	0.2	HGA
1999-009 17:00–010 02:09	65				-131.8	-132.2	0.4	25	32	1.1	-1.1	-0.7	HGA
1999-012 17:30–013 00:14	65	7.2	6.4	0.8	-132.8	-132.75	-0.1	24	30.5	1.0	-0.2	-1.1	HGA
1999-013 16:45–23:41	65	6.75	6.2	0.6	-132.4	-133	0.6	26.5	33.5	1.0	-0.5	-0.4	HGA
1999-014 03:15–07:49	25	6.5	6	0.5	-132.5	-132.85	0.3	30.4	32.8	0.3	0.2	0.0	HGA
1999-015 03:15–08:02	25	6.3	5.8	0.5	-132.6	-133.05	0.5	30	32.5	0.3	0.2	0.1	HGA
1999-016 03:15–07:39	15	6.4	5.85	0.6	-132.75	-133.6	0.8	29	29.3	0.0	0.5	0.8	HGA
1999-016 16:30–017 00:29	65	6.3	5.5	0.8	-133.4	-133.6	0.2	26	30.5	0.7	0.1	-0.5	HGA
1999-017 00:15–05:44	15	5.5	5.65	-0.2	-133.2	-133.65	0.5	29	30	0.1	-0.3	0.3	HGA
1999-017 16:30–018 00:44	65	5.3	5.3	0.0	-133.6	-133.85	0.3	25	30.5	0.9	-0.9	-0.6	HGA
1999-018 00:15–09:29	15	5.9	5.45	0.5	-133.25	-133.8	0.6	29	29.5	0.1	0.4	0.5	HGA
1999-019 02:45–07:39	25	5.25	5	0.3	-133.45	-133.85	0.4	30.7	32.9	0.3	-0.1	0.1	HGA
1999-022 01:00–09:59	15	5.5	4.7	0.8	-133.5	-134.7	1.2	29.25	30	0.1	0.7	1.1	HGA
1999-022 19:30–023 00:24	54	4.7	4.1	0.6	-134.2	-134.7	0.5	31	33.1	0.3	0.3	0.2	HGA
1999-023 16:15–23:45	65	5	4.15	0.9	-134	-135	1.0	30	30	0.0	0.9	1.0	HGA
1999-024 02:45–11:44	25	4.6	4	0.6	-134.25	-135	0.8	30	32.9	0.4	0.2	0.3	HGA
1999-24 19:30–025 00:09	54	4.2	3.7	0.5	-134.7	-135.1	0.4	30	33	0.4	0.1	0.0	HGA
1999-025 03:15–11:14	25	4.45	3.75	0.7	-134.4	-135.05	0.7	30	33	0.4	0.3	0.2	HGA
AVERAGE				0.5			0.5			0.4	0.1	0.1	
MIN				-0.2			-0.1			0.0	-1.1	-1.1	
MAX				0.9			1.2			1.1	0.9	1.1	
Variance, assuming Gaussian				0.0			0.0			0.0	0.1	0.1	
Sigma				0.2			0.2			0.2	0.3	0.4	

Table 9. Encoding Modes/Data Rates Verified in Noncoherent Downlink Mode

Data Rate	Convolutional Code
40	(7,1/2)
2100	(7,1/2)
3150	(15,1/6)
13272	(15,1/6)
19908	(15,1/6)

4.8 Nonlinear Phase Modulator Performance

Because of the intermodulation effect, the SDST’s ranging and telemetry-carrier suppression deviates significantly from what the established theory of linear phase modulation would predict. For this reason, DS1’s telecom team constructed a special nonlinear phase-modulation-loss model, which was used to predict ranging-induced carrier suppression for the SDST.

The validity of this model was tested on day of year (DOY) 1998-344 when the ranging modulation was applied to the downlink with and without telemetry modulation as part of

the Xrange test. The plot of downlink carrier power (Pc) versus time (see Figure 3) shows the carrier suppression at low ranging mod index (17.5°) to be approximately 1.0 dB (Pc= -124.4 dBm for ranging OFF, -125.4 dBm for ranging low at 38 data number [DN]). At 35 degrees ranging mod index and a telemetry mod index setting of 32 DN, the carrier suppression was measured to be 7.6 dB. The contribution from telemetry modulation at 32 DN is 5 dB, based on ground-test data. The ranging induced carrier suppression is, therefore, approximately 2.6 dB at 35 ° ranging modulation setting, which agrees well with pre-flight test data (see Table 10) and shows that the X-band phase modulator has not deviated in performance since pre-launch tests. The pre-flight measurement data are contained in section 2.7.2 of the flyable engineering model (FEM) test report dated 12/18/97.

At Ka-band (see Table 11), the phase modulator is linear, the suppression due to telemetry modulation is modeled as 20*log(cos(telemetry mod index)), and the suppression due to ranging is 20*log(J₀(ranging mod index)). A plot of Ka-band carrier power as a function of time at different ranging mod index settings is shown in Figure 3.

M-0727 (AA5 PC) vs ERT




Figure 3. Measured X- and Ka-band Carrier Power as a Function of Time During DOY 1998-344, Showing Various Ranging Carrier Suppression Values (Note: ERT = Earth-Received Time)

4.9 Coherent Doppler Tracking

Coherent Doppler tracking was conducted as part of the Xstable test. The intent was to validate the coherent frequency stability of the SDST both for Doppler tracking and for future radio science applications. During the test, the SDST generated both X-band and Ka-band downlinks that were coherent to the X-band uplink, with coherency ratios of 880/749 and 3360/749, respectively. The X-band and Ka-band downlink signals were received at DSS-25 using the BVR and at DSS-13 using the experiment tone tracker (ETT). The frequency measurements from both the BVR and the ETT were then used to measure the phase deviation and Allan deviations of the X-band and Ka-band downlinks. The stability of the downlink carrier as received at the tracking station should not be affected by the presence of command or ranging modulation on the uplink, or telemetry modulation on the downlink.

Shown in Figure 4 are the X-band and Ka-band frequency residuals taken at DSS-13 using the ETT, after correcting for Earth rotation and spacecraft Doppler effects. It is seen that periodic frequency variations of ± 5 millihertz at X-band

(Figure 4a) and ± 20 millihertz at Ka band (Figure 4b) were visible in the X-band and Ka-band downlink-frequency residuals. These variations are common to both the X-band and Ka-band, and are believed to be due to deadbanding of the spacecraft. When the common mode is removed by subtracting the X-band frequency residual and a Ka-band residual scaled down by a factor of 3360/880, no periodic variation is visible in the data (see Figure 4c). The 0.1-hour (6-minute) period shown in Figure 4 (a–b) was similar to the deadband cycle frequency of the spacecraft.

The two-way Allan deviation performance of the SDST is illustrated in Figure 5. Both X-band and Ka-band downlinks showed an Allan deviation of better than 1 part in 10^{13} with 10 seconds integration time. This translates to a Doppler measurement accuracy of 0.8 millihertz at 10 seconds integration time (or 0.015 mm/s). When the common mode variation was removed, the X/Ka-band downlinks showed a delta frequency stability of better than 1 part in 10^{14} with 10 seconds integration time, or 0.0015 mm/s in Doppler measurement.

4.10 Noncoherent Downlink Frequency Stability

This test verified that the SDST generates downlink frequencies (X-band and Ka-band) from its auxiliary oscillator that have sufficient stability as downlink carriers to be received by the tracking station. The stability of the X-band downlink in the SDST’s noncoherent mode was not expected to be affected by reception of an uplink carrier. The test (Xstable) was conducted on DOY 1998-344, when DS1 pointed the +X axis to Earth and transmitted both X-band and Ka-band downlinks. The test measured the frequency of the X-band and Ka-band downlinks over a period of two hours. Shown in Figure 6 is a plot of downlink frequency as a function of time for both the X-band and Ka-band. It is seen from this plot that, under nominal operating conditions (including spacecraft deadbanding), the X-band downlink varies by approximately 75 Hz, whereas the Ka-band downlink varies by a corresponding ratio of (3360/880) and has a maximum frequency deviation of approximately 300 Hz. The close resemblance of the X-band and Ka-band downlinks is expected since they are coherent with the same multiplication ratio. A check of pre-flight temperature data showed that the SDST has a frequency rate of change of over 200 Hz/°C. Therefore, the perceived frequency change can be due to small thermal variations at the spacecraft.

4.11 Ranging Functional Verification

The SDST is designed to provide turnaround ranging simultaneously with uplink command and downlink telemetry. Since the SDST has a nonlinear phase modulator, which effectively causes excessive inter-modulation losses when ranging modulation is applied simultaneously with telemetry modulation at high mod indices, ranging performance validation is limited to the modulation indices planned for routine mission operations. That is, X-band telemetry modulation is limited to 38 DN (approximately

65 degrees) at 17.5 degrees ranging mod index setting and to 32 DN (approximately 58 degrees) at 35 degrees ranging mod index setting. At Ka-band, the phase modulator behaves linearly; it is expected that ranging modulation effects can be correctly predicted using well-established equations.

Both X-band uplink/X-band downlink (X/X) and X/Ka ranging have been successfully demonstrated. Downlink ranging mod indices of 17.5° and 35° have been used at both the X-band and the Ka-band. The actual ranging signal-to-noise ratio (Pr/No) agrees with the predicted within 1 dB at X-band and 1.5 dB at Ka-band. The ranging residuals are larger than one-way (downlink or uplink) residuals because ranging is a two-way link: both a stronger than predicted uplink (typically, 0.7 dB) and a stronger than predicted downlink (typically 0.7 dB) contribute to a larger residual.

Shown in Figure 7 are typical examples of ranging Pr/No values at the X-band and Ka-band; predicts are also shown for comparison to actuals for the X/Ka-band track on 2/4/99. It is seen that the actual downlink Pr/No values are in good agreement with the predicted values. Similar X/Ka-band data, collected for DOY 1999-096 (4/6/99), are shown in Figure 8. It is seen that the X-band Pr/No is within 2 dB of the predicted value, whereas the Ka-band Pr/No is within 3 dB of the predicted value.


Ranging residuals, after accounting for the spacecraft’s trajectory, are shown in Figure 9. The ranging residuals (measurement errors when corrected for spacecraft trajectory effects) are typically on the order of 0.5 m when the HGA is used. Larger residuals are seen when the LGA is selected and when the spacecraft is pointed away from Earth. This fact correlates with the weaker uplink and downlink signals.

Table 10. X-band Ranging Suppression Due to Nonlinear Phase Modulator, Measured Versus Predicted on the Basis of Pre-flight Data


Carrier Suppression	Measured In Flight	Pre-Flight Test Data	Delta, In-Flight to Pre-Flight	Link Analysis, Model	Delta, Measured-Model
17.5° ranging (telemetry 38 DN)	1.0	0.9	0.1	-0.79	-1.0
35° ranging	2.6	2.7	-0.1	-2.94	-2.6

Table 11. Measured Ranging Suppression at Ka-band


Telemetry Modulation	Ranging Modulation	Pc, Measured	Pc, Predicted	Delta, Measured-Predicted
0 DN (0°)	0°	-122 dB (1 mW)		
54 DN	35°	-138 dB (1 mW)	-137.8 dB (1 mW)	0.2 dB



(a)



(b)



(c)

Figure 4. (a) Measured Coherent Frequency Stability at X-band and (b) at Ka-band; and (c) Relative Stability of X/Ka-bands






Figure 5. (a) Measured Allan Deviations for X-band and (b) Ka-band Downlinks, and (c) Measured X/Ka-band Relative Stability



(a)



(b)

Figure 6. (a) Measured Noncoherent Downlink Carrier Stability at X-band and (b) at Ka-band

4.12 Beacon Tone Generation and Tone Tracking

The SDST was designed to provide a flexible selection of downlink telemetry subcarrier. This feature was used to provide the beacon tone for noncoherent signaling: the SDST would provide one of four selectable subcarrier frequencies (20, 25, 30 and 35 kHz) at a near-90° modulation index (complete downlink carrier suppression). Detection of the tone frequency can be used to signal one of four possible spacecraft states.

The Xtone activity was designed to show that the beacon downlinks signal from the SDST could be detected effectively, even at low signal power level. The results show that the SDST was able to generate and transmit the four required beacon tones (frequencies of 20, 25, 30, and 35 kHz) at X-band. No beacon experiment was performed at Ka-band.

At the planned modulation index of 54 DN, more than 99% of the power was in the subcarrier sidebands. The expected beacon tones were successfully detected. In order to test the detection of beacon tones on the ground at weaker signal levels, Xtone was successfully performed at much lower modulation indices (1.7, 3.4, 5.1, and 6.8 degrees), a procedure that allowed a much lower signal to be detected. The tone-detection system successfully detected signals as low as SNR=4.5 dB.

4.13 External Telemetry Sampling Functions

The SDST samples external analog and temperature telemetry signals. These channels served, among other

functions, to provide the necessary engineering data for KAPA performance validation.

5.0 DSN KA-BAND READINESS VERIFICATION

Since the SDST is the first Ka-band capable deep space transponder, a significant portion of the technology validation activity was conducted for both the X-band and the Ka-band. In addition to the technology validation objectives cited previously, a side benefit of the DS1 Ka-band downlink is direct verification of the operational readiness of the DSN Ka-band subnet and of the performance advantages of the Ka-band relative to the X-band. Although only one of the three subnet stations (DSS-25) was ready in time to support DS1, the performance data gathered using DS1's Ka-band downlink were useful in evaluating the projected Ka-band performance at other subnets in the future.

DS1 powered on Ka-band during December 9–10, 1998 and again after January 10, 1999 prior to the first thrusting cruise arc. The initial characterization tests (December 1998) supported the following SDST technology validation objectives:

- Demonstrate SDST capability to support simultaneous X-band and Ka-band downlinks at various data rates and modulation indices.
- Measure one-way and two-way frequency stability and X/Ka-band relative frequency stability of Ka-band downlink.







Figure 7. Predicted Versus Actual Pr/No for DOY 1999-035



(a)



(b)



(c)

Figure 8 (a–c). Ranging Residuals for DOY 1999-096 (4/6/99): (a) DSS 25, (b) DSS 34, and (c) DSS 46




Figure 8 (d). Ka-band Ranging Residuals for DOY 1999-096 (4/6/99)




Figure 9. Ranging Residuals (Measurement Errors) after Accounting for Station and Spacecraft Motions

- Verify X/Ka-band radio metrics performance.
- Demonstrate operation of 3-W (2.5-W) Ka-band solid-state power amplifier (SSPA) in space.
- Collect operating data for the Ka-band SSPA (gate current and drain voltage telemetries and operating temperature data) for future analysis.

Additionally, the Ka-band downlink supported the following readiness demonstration objectives:

- Demonstrate DSN readiness to support Ka-band missions.
- Demonstrate dual-band (X/Ka), end-to-end telemetry flow from a spacecraft to the DS1 MSA.
- Demonstrate the capability to generate necessary station predicts for Ka-band tracking.
- Demonstrate the station capability to perform radio metric tracking (Doppler and ranging) on the Ka-band downlink.
- Measure Ka-band system noise temperature and link performance advantages relative to X-band.
- Demonstrate DSS-25 capability to accurately point the 34-m antenna using blind pointing.

Ongoing characterization tests (since January 1999) also demonstrated continuous operation of the Ka-band downlink and supported characterization of the Ka-band downlink threshold and verified link-margin calculation. The DS1 Ka-band downlink was also used to provide a stable signal for measurement/characterization of 70-m DSN station pointing and receiving capability using different techniques (array feed, deformable flat mirror, etc.)

Future plans for the Ka-band downlink from DS1 include:

- Ka-band beacon tone experiment.
- Long-term monitoring of X/Ka-band propagation data with spacecraft Ka-band downlink.
- Characterization of Ka-band performance under very low downlink power.
- Characterization of Ka-band performance during solar conjunction.
- Possible radio science during solar conjunction.

Additionally, the Ka-band downlink is relatively insensitive to solar plasma-induced scintillation. Since DS1's next encounter (with comet Wilson-Herrington) occurs at an Sun-Earth-Probe (SEP) angle of 2 degrees, the availability of the Ka-band downlink can be very valuable as it can be the only direct confirmation of the link at a low SEP angle.

5.1 34-m Antenna Pointing Performance

5.1.1 Blind Pointing Model—There are two blind pointing models for DSS-25. The first is the standard X-band blind pointing model, which was used for the majority of the tracks. The second is a fourth-order Ka-band blind pointing model, which was used on an experimental basis. When the

fourth-order Ka-band model was used, it was observed that the Ka-band signal power was generally strong. When Conscan was turned on to bring the antenna on point, only 1 dB of increase in signal-to-noise ratio was observed with the fourth-order Ka-band model. With the X-band pointing model, this was not the case. At times, increases in the signal-to-noise ratio upwards of 5 dB were observed when Conscan was turned on and when the antenna was pointed using the X-band pointing model. This indicates that the Ka-band model is quite accurate and requires a minimum of active correction. For future Ka-band tracks it is recommended that the fourth-order Ka-band blind pointing model be used.

5.1.2 Conscan Mode—As mentioned above, we observed an increase of 1 to 5 dB in the signal-to-noise ratio when Conscan was turned on. This indicates that, given the current set of blind pointing models available, it is advisable to use an active pointing mechanism on the 34-m beam waveguide (BWG) antenna to take full advantage of the Ka-band performance. Furthermore, it was observed that when Conscan was turned on, fluctuations in the BVR symbol signal-to-noise ratio (SSNR) decreased from approximately ± 0.3 dB to ± 0.03 dB. Figure 10 is a plot of pointing residuals as a function of time for DOY 98-344. It is seen that pointing residuals of less than 4 mdeg were effectively maintained. This indicates that the antenna was on point because fluctuations in antenna pointing cause less degradation at peak. This is due to the fact that the roll-off in gain of the antenna at peak is not too sharp. Another thing to note is that several times Conscan was turned on for a few minutes, and pointing offsets were obtained; then Conscan was turned off, but the offsets were kept. This resulted in several hours of very good tracking. However, as the track proceeded, the pointing degraded and Conscan needed to be turned on to obtain new pointing offsets. This was especially true at high elevations, where the blind pointing model may not be as accurate when fast changes in azimuth occur.

5.2 Ka-band System Noise Temperature and Link Capacity Projection

During our initial tracks at DSS-25, there were problems with the reporting of the Ka-band system noise temperature (SNT). Once this was brought to the attention of DSN operations, there was improvement in the reporting of the SNT and reported SNT values were between 40 K and 50 K, depending on elevation (although some values were as high as 56 K at high elevations) (see Figure 11). Given that the contribution of noise sources other than atmosphere is about 27 K, this indicated a zenith atmospheric noise temperature of about 10 K to 12 K, which corresponds to 50% to 70% weather. It should be noted that these weather percentages are calculated not from the standard 810-5 numbers but from the latest set of water vapor radiometer data collected at Goldstone. These numbers reflect 46 months of

observations and are by far the most accurate source of atmospheric-noise-temperature data for the Ka-band at Goldstone.

The validity of the observed SNT values was verified in two ways. First, the theoretical SNT was calculated for a zenith atmospheric noise temperature of 10 K, based on DOY 1998-344 track elevation, and plotted against DOY 1998-344 SNT data (see Figure 11). Then, the Ka-band antenna-gain-to-system-noise-temperature-ratio (G/T) advantage over X-band was calculated, based on the received SNR values for DOY 1998-344 and DOY 1999-035, and plotted against the predicted Ka-band G/T advantage over X-band at 50% and 70% weather (10 K and 12.5 K zenith atmospheric noise temperature), respectively. The following method was used to calculate the G/T advantage. First, it was noted that if the spacecraft had the same amount of transmission power available for the X-band and the Ka-band over the same size antennas, with the exact same efficiency for Ka-band and X-band, the Ka-band Equivalent Isotropic Radiated Power (EIRP) would have been 11.6 dB higher than that for the X-band. However, for DS1, the EIRP for Ka-band is 3.4 dB *less* than that for the X-band. This means that in order to make a fair comparison between Ka-band and X-band performance we need to add 15 dB (11.6+3.4) to the measured Ka-band signal-to-noise ratio.

Therefore, we calculated the total signal-to-noise ratio (Pt/No) for each band from the estimates for carrier signal-to-noise ratio (Pc/No), symbol signal-to-noise ratio (Es/No), and, when applicable, ranging signal-to-noise ratio (Pr/No). Then, we added 15 dB to the Ka-band Pt/No and subtracted the X-band Pt/No from the total. The result is the Ka-band G/T advantage over X-band. These results are presented in Figure 12 and Figure 13.

As we can see from Figure 11, the observed SNT values for DOY 1998-344 closely match the predicted SNT values at the Ka-band. The mismatches that are observed occur at higher elevations, where the theoretical antenna models usually do not quite match the actual antenna performance. In Figure 12 and Figure 13, the G/T advantage for Ka-band over X-band is approximately 10 dB at the higher elevations. This is about 1 to 1.5 dB higher than those predicted for 50% weather for the antenna configuration (dual-frequency, diplexed configuration) that was employed during these tracks. There could be several reasons for the discrepancy between the theoretical and actual results. First of all, the theoretical model may not be accurate. This could lead to inaccurate estimates of antenna gain and system noise temperature for different elevations. This is the most likely source of error due to the large amount of error observed. Other factors, such as miscalibration of the SNT




Figure 10. DSS-25 Conscan Pointing Residuals, Showing that Pointing Error is Generally Less than 4 millirads




Figure 11. DOY 1998-344 Ka-band Measured SNT and 50% Weather-Predicted SNT




Figure 12. Ka-band G/T Advantage Over X-band, DOY 1998-344, DS1 Track at DSS-25




Figure 13. Ka-band G/T Advantage Over X-band, DS1 Track, DOY 1999-035, DSS-25

readouts, nonlinearities in the spacecraft modulator, and changes in the EIRP, could also effect the measurements, as could the pointing of the spacecraft antenna. The performance advantage observed during these tracks for the Ka-band over the X-band indicate that Ka-band could be, if it is not already, an evolutionary step to increase the capacity of the DSN by at least a factor of four.

There are several caveats to these observations. First of all, the tracks were performed under very good conditions. There was very little wind and humidity was low. Further tracks need to be performed, especially during summer, when the humidity at Goldstone is high, to observe the behavior of the Ka-band under adverse conditions. Secondly, DS1 carries a relatively low-gain Ka-band antenna. Due to this, the antenna pointing for DS1 Ka-band is not as stringent as it would be on a spacecraft that carries a higher gain antenna (say a 1-m dish). In that case, the secondary effect of antenna pointing error on the spacecraft, while negligible at the X-band, could drastically affect the performance of the Ka-band. Finally, it should be noted that the performance advantage that was calculated was only for the ground G/T performance. The actual end-to-end performance advantage of the Ka-band link depends also on spacecraft configuration. Lower efficiency of Ka-band amplifiers and lower efficiency of Ka-band antennas should also be factors in determining whether or not Ka-band should be used on a spacecraft.

5.3 Ka-band Performance Threshold

This test was designed to evaluate the quality of the Ka-band telemetry received from DS1. This was done by changing the received bit signal-to-noise ratio at the Ka-band by changing the telemetry mod index and then observing the lock status of the frame synchronizer subassembly (FSS) and the maximum likelihood convolutional decoder (MCD). Furthermore, telemetry gap reports were to be used to evaluate the decoding signal-to-noise ratio threshold for the (15,1/6) convolutional code, concatenated with the Reed-Solomon (255,223) interleaving depth 5 code.

Four days, DOY 1999-025, 027, 028, and 030, were scheduled for these tests. The spacecraft was sequenced so that the mod index would change every five minutes for cycles of 35 to 45 minutes, depending on the day. The mod index is the highest at the beginning of the cycle, producing the highest SNR, and the lowest at the end of the cycle, producing the lowest SNR. These mod indices were selected so that the test would produce SNR values both above and below the expected threshold of 0.65 dB bit SNR (corresponding to -7.05 dB symbol SNR). Shown in Figure 14 is a plot of the receiver SNR as a function of time. The steps in measured SNR result from modulation index changes.




Figure 14. Ka-band Threshold Measurement Data

The FSS and MCD lock statuses were obtained immediately after the pass from the monitor data and correlated with the measured receiver symbol SNR shown in Figure 14. The FSS starts losing lock when the symbol SNR is between -7.5 dB and -8 dB (bit SNR between 0.3 and -0.2 dB). This corresponds with the observations made on the X-band channel during tests for previous missions, where the FSS lost lock at about -7.5 dB symbol SNR. The MCD loses lock when the symbol SNR is between -8 dB and -8.5 dB (bit SNR between -0.2 dB and -0.7 dB). These values also match rather closely those observed for previous missions.

5.4 X/Ka-band Radio Science

X/Ka-band radio science is a potential objective for the solar conjunction. The frequency stability characterization performed during IC showed that a relative frequency stability of better than one part in 10^{13} can be achieved with an integration time of better than 10 seconds. The frequency stability appears to be limited by the relative shift of the X- and Ka-band phase center; the time scale is limited by the deadband of the spacecraft.

5.5 Ka-band Link Threshold at Low Bit Rate

As of this writing, this test has not been performed due to limitations imposed by the mission.

5.6 Ka-band Antenna Pointing and Gravity Compensation at 70 m

The DS1 Ka-band signal was used as part of a task to evaluate the performance of DSS-14 at the Ka-band. A complete report is being prepared by the task force on the experiments performed at DSS-14. Part of the report will address the use of the DS1 Ka-band signal to evaluate DSS-14 performance. As of this writing, the report is not complete. Following is a brief description of the systems that were used to evaluate and improve DSS-14 Ka-band performance, along with a summary of the conclusions presented by the task force.

5.6.1 Purpose of the Ka-band Tests at DSS-14—The purpose of these tests is twofold: (1) measure improvements in the antenna efficiency at Ka-band and (2) measure improvements in the pointing accuracy of the antenna for Ka-band tracking.

In order for DSN to use the 70-m subnet for tracking at Ka-band it must be shown that the antennas have sufficient gains and that they can be pointed accurately enough to justify their use at that frequency. It is, therefore, paramount to test various candidate technologies that improve the gain and pointing of the 70-m antennas to measure the potential performance of the 70-m subnet at Ka-band.

5.6.2 Candidate Technologies—There are three different technologies that are under consideration for use on the 70-m subnet for Ka-band pointing and gain compensation: (1) the monopulse pointing and compensation system, (2) the array feed pointing and compensation system, and (3) the deformable flat plate (DFP) compensation system.

Monopulse uses simple measurements in the antenna focal plane to estimate the peak of the antenna beam and to adjust the pointing of the antenna to help ensure that the beam is centered as the antenna tracks. Monopulse requires a coherent signal on which to peak; therefore, it can be used only on spacecraft signals. Monopulse improves only the pointing of the antenna and does nothing to improve its gain.

Array feed uses seven feed horns in the focal plane to estimate the peak of the antenna beam and to adjust the pointing of the antenna accordingly, so that the beam is centered. In addition, array feed has the potential to combine the output of the seven feed horns to compensate for decreases in the gain of the antenna due to gravity deformations. Array feed works with both coherent sources (i.e., spacecraft signals) and noncoherent sources (i.e., natural radio source such as stars, galaxies, and planets).

DFP is basically an RF mirror that changes its form according to the elevation of the antenna in order to compensate for the decrease in gain due to gravity. DFP compensates only for gravity deformations and does not affect the pointing of the antenna.

Since we need to improve both the gain and the pointing of the antenna, DFP and monopulse cannot be used by themselves. Therefore, there are three configurations that are considered during these tests:

- Array feed.
- Monopulse + DFP.
- Array feed + DFP.

5.6.3 Use of DSI—While natural radio sources could be used to measure the performance of the DFP and array feed systems, it is necessary to measure the performance of each configuration using real spacecraft signals, since the bottom line for the DSN is the quality of the returned data from the spacecraft. In addition, the monopulse system could not be tested with natural radio sources.

Currently, there are only four spacecraft that are operating at the deep space Ka-band (31.8–32.3 GHz): (1) Student Undergraduate Research Fellowship Satellite (SURFSAT), (2) Mars Global Surveyor (MGS), (3) Cassini, and (4) DS1. SURFSAT, which was supposed to act as a Ka-band beacon, orbits the Earth. However, due to SURFSAT's tumbling motion, its Ka-band signal power fluctuates and is,

therefore, unreliable for threshold-related experiments, for which a stable downlink is required. The MGS Ka-band signal was turned off during the time these experiments were being performed. Furthermore, implementation of the MGS Ka-band system, with large spurious signals at high modulation indices, resulted in uncertainties in total signal power and, thus, in unreliable threshold measurements. Cassini's Ka-band does not carry telemetry; Cassini is under strict configuration control.

DS1 is the only spacecraft that has a complete, independent, and stable Ka-band telemetry system. In addition, the spacecraft team has been more than helpful in meeting the needs of these tests at DSS-14. Therefore, we have naturally gravitated towards the use of DS1 during these tests.

The DS1 signal is used both to establish a baseline for each configuration and to measure the performance of each configuration when its constituent systems are activated.

5.6.4 Conclusions—The experiment was performed successfully. In addition, the DS1 Ka-band signal was used successfully to evaluate the performance of candidate configurations. It is the opinion of the task force that the combination of array feed and deformable flat plate provides the best option for receiving Ka-band at DSS-14. This is due to the fact that this combination provides the most gravity compensation for DSS-14 while providing accurate pointing. Another conclusion of the task force was that DSS-14 Ka-band performance is not adequately characterized at high elevations. Therefore, in the future, the DS1 Ka-band signal could be used in conjunction with DSS-25 to characterize DSS-14 performance at high elevations for the Ka-band.

6.0 SUMMARY AND CONCLUSION

The in-flight checkout activities and ongoing flight validation of the SDST provided confidence that the transponder functioned as intended. With the exception of nonlinear phase modulation and the temperature sensitivity of receiver best lock frequency, the SDST functioned exactly as intended.

One should also note that both the nonlinearity of the phase modulator and the variation in BLF/SPE have been corrected for the current generation of the SDST, scheduled to be flown on Mars 01 and SIRTf missions. Furthermore, unlike the DS1 SDST, which functioned only with a single-string C&DH, the Mars 01 SDST supports dual-string cross-strapping with the C&DH. These performance improvements and added capabilities, together with DS1's in-flight validation, make the use of the SDST truly low-risk for future flight projects.

7.0 LIST OF REFERENCES

- [1] G. Glass, “JPL Small Deep Space Transponder Functional Specification and Interface Control Document,” Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, JPL Internal Document ES-518193, May 2000.
- [2] A. Makovsky, “DS1 Telecommunications, FEM SDST/DSN Performance and Compatibility Motorola Test Report,” Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, JPL Internal Document D-19045, May 1998.

8.0 ACKNOWLEDGMENTS

So many people have contributed to the success of the SDST and its technology validation activities that it is impossible to acknowledge each of them individually. Instead, the authors would like to acknowledge the following teams.

The SDST development team: Development of the SDST was a large and complex team effort involving members from both JPL and Motorola. The authors would like to specifically acknowledge the following personnel for their leadership effort in the development:

Sam Zingales
Carl Nuckolls (Motorola)
Keith Siemsen (Motorola)
Dave Andersen (Motorola)

The DS1 telecom team: The successful integration and launch of the SDST would not have been possible without the direct support of the DS1 telecom team. Additionally, the DS1 telecom team participated in many of the technology validation activities, both pre-launch and post-launch. The authors would like to acknowledge the

following members of the DS1 telecom team for their support of the technology validation activities:

Marty Herman
Chien Chen
Sam Valas
William Hatch
Andrew Makovsky

The DS1 flight and mission support team: A large team of operational personnel have collaborated in the planning and execution of the technical validation activities. The authors would like to acknowledge the participation of the following individuals, without whose support the technology validation and characterization of the SDST would not have been successful:

Pam Chadbourne
Kathy Moyd
Rob Smith
Ben Toyoshima

Special thanks to Jim Taylor of the DS1 telecom team, without whose diligent planning and monitoring of SDST telemetry, the technology validation activities could not have taken place.

Finally, in addition to those cited above, the following people have contributed to the preparation of this report:

Shervin Shambayati
David Morabito
Miles Sue

The research described in this report was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Appendix A. List of Telemetry Channels and Names

Table A1 is a list of all of the telemetry channels that the SDST team collects and uses. Note the importance of "monitor" channels in this work. (Jim Taylor, 10/20/99.)

Table A1. Channels and Mnemonics

Channel	Mnemonic
T-0017	nar_band AGC
T-0018	carlock_acem
T-3252	sdst_evnt ct
T-3228	rvr spe
T-3116	aux_osc temp
T-3124	vco tmp
T-4002	XPA temp
P-2061	ess_bus v
T-3500	sdst_dc pwr
T-3501	xpa_dc pwr
T-3316	xpa_in pwr
T-3476	X Exc SPE
A-1637	bbc_CtrlErr0
A-1621	bbc_CtrlErr1
A-1625	bbc_CtrlErr2
T-3144	coherency
T-3240	cmd_datarate
T-0025	dnlink_rate
B-3090	DlinkClokRat
T-3156	x_tlm_mod
T-3188	ka_tlm_mod
T-3132	xtlm_coder
T-3136	katlm_coder

Channel	Mnemonic
T-3148	xsubcar_freq
T-3180	ksubcar_freq
T-3100	X_ranging
T-3101	ka_ranging
T-3224	ranging_gain
T-3104	X_Exciter
T-3105	ka_Exciter
P-3127	XPA_on_off
P-3160	DAM_on_off
T-3002	wts1_pos1
T-3004	wts2_pos1
M-0130	MCD1 SNR
M-0781	AB5 SS1 SNR
M-0773	AB5 PCN0
M-0777	AB5 PC
M-0775	AB5 SNT
M-0787	AB5 SPE
M-0618	RNG PRN0 X
M-0304	ANT A EL ANG
M-0305	ANT A AZ ANG
M-0308	A CNSCN
M-0309	A CNSCN LOOP

Appendix B. Date of Turn-on/off and Frequency of Data Capture

The SDST was turned ON as part of the launch script in fault protection. Per the ACE log, the downlink from the spacecraft was first detected at 98-297/14:35 UTC. The SDST has been on continuously since then, except for short

hiccoughs due to spacecraft safing. In fact, one could say that parts of it (e.g., receiver) have been on continuously. (Jim Taylor 10/29/99)