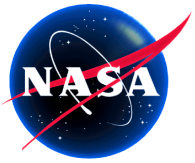

Autonomous Optical Navigation (AutoNav) DS1 Technology Validation Report

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

JPL

Table of Contents

Section Page

Figures

Figure Page

Tables

Table Page

Appendices

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 Borelly 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 *DS1 :* Technology Validation Fact Sheet

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

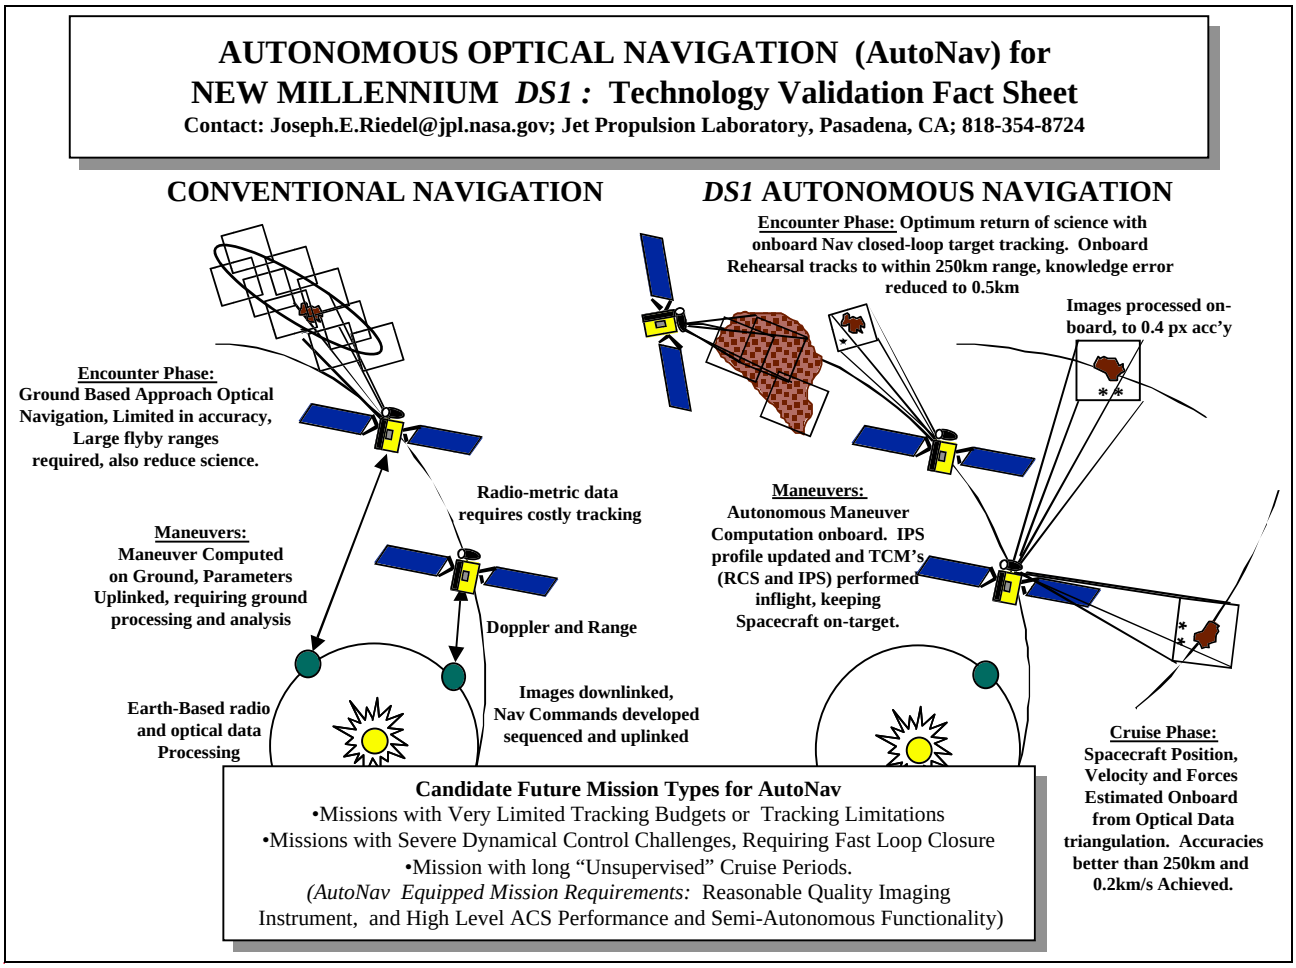

Diagrammatic and Comparative Descriptionof AutoNav Technology and Validation

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.

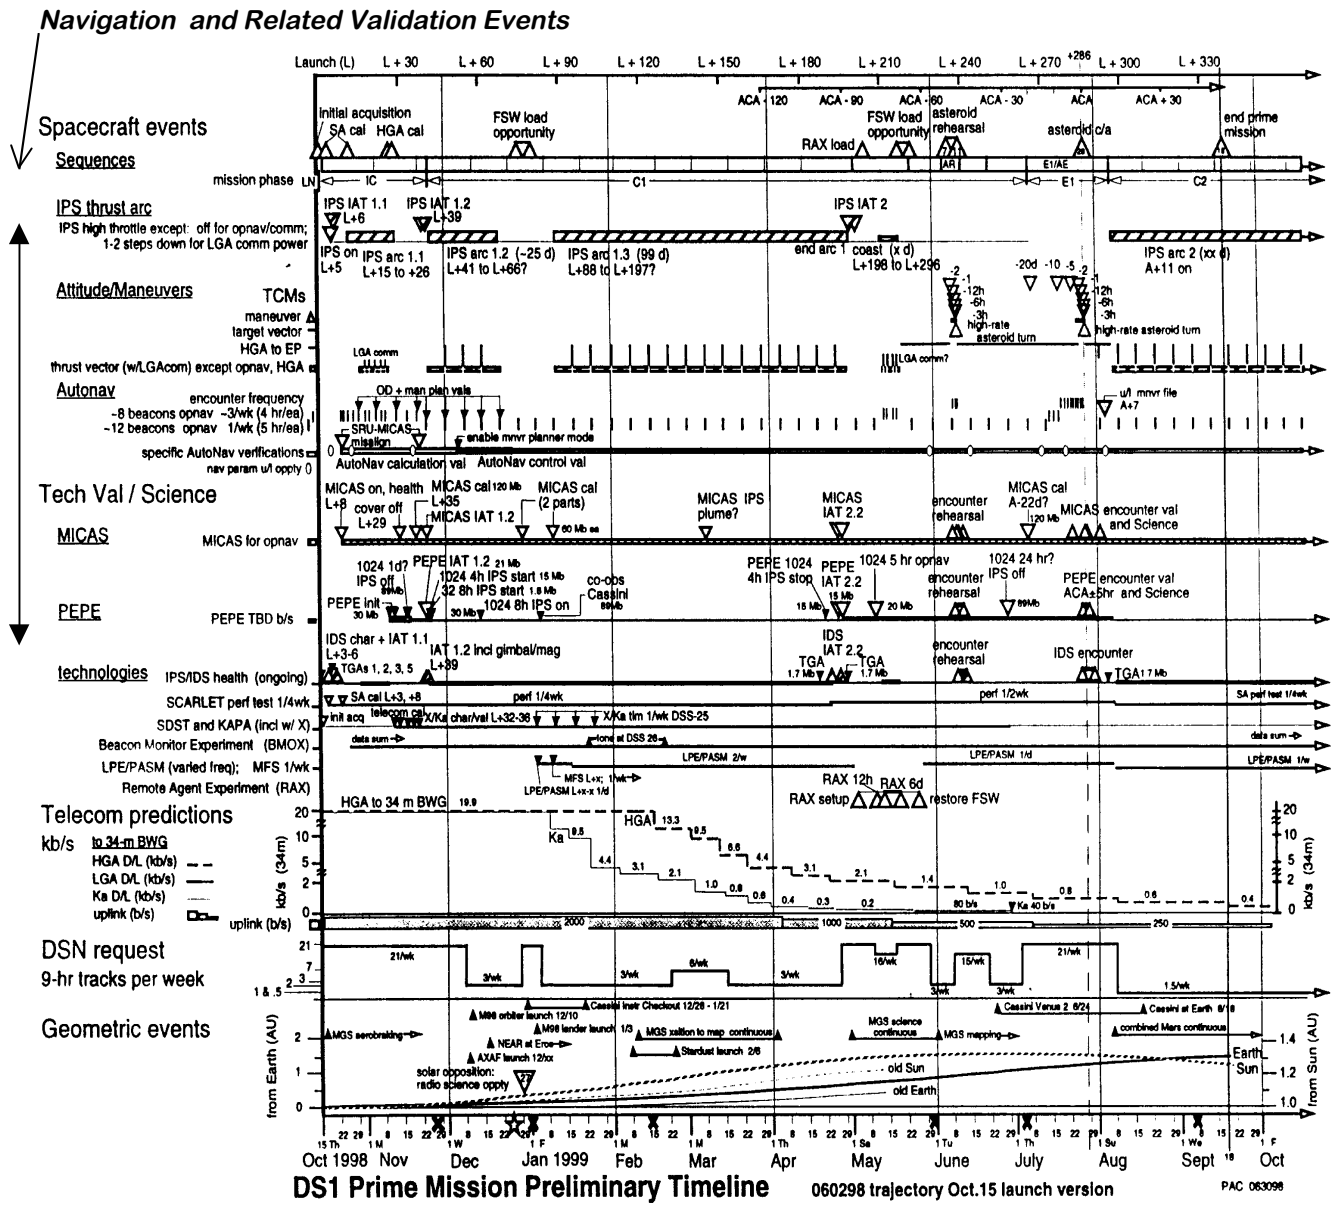

Figure 1. DS1 Mission and AutoNav Operations/Validation Plan and Schedule

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.

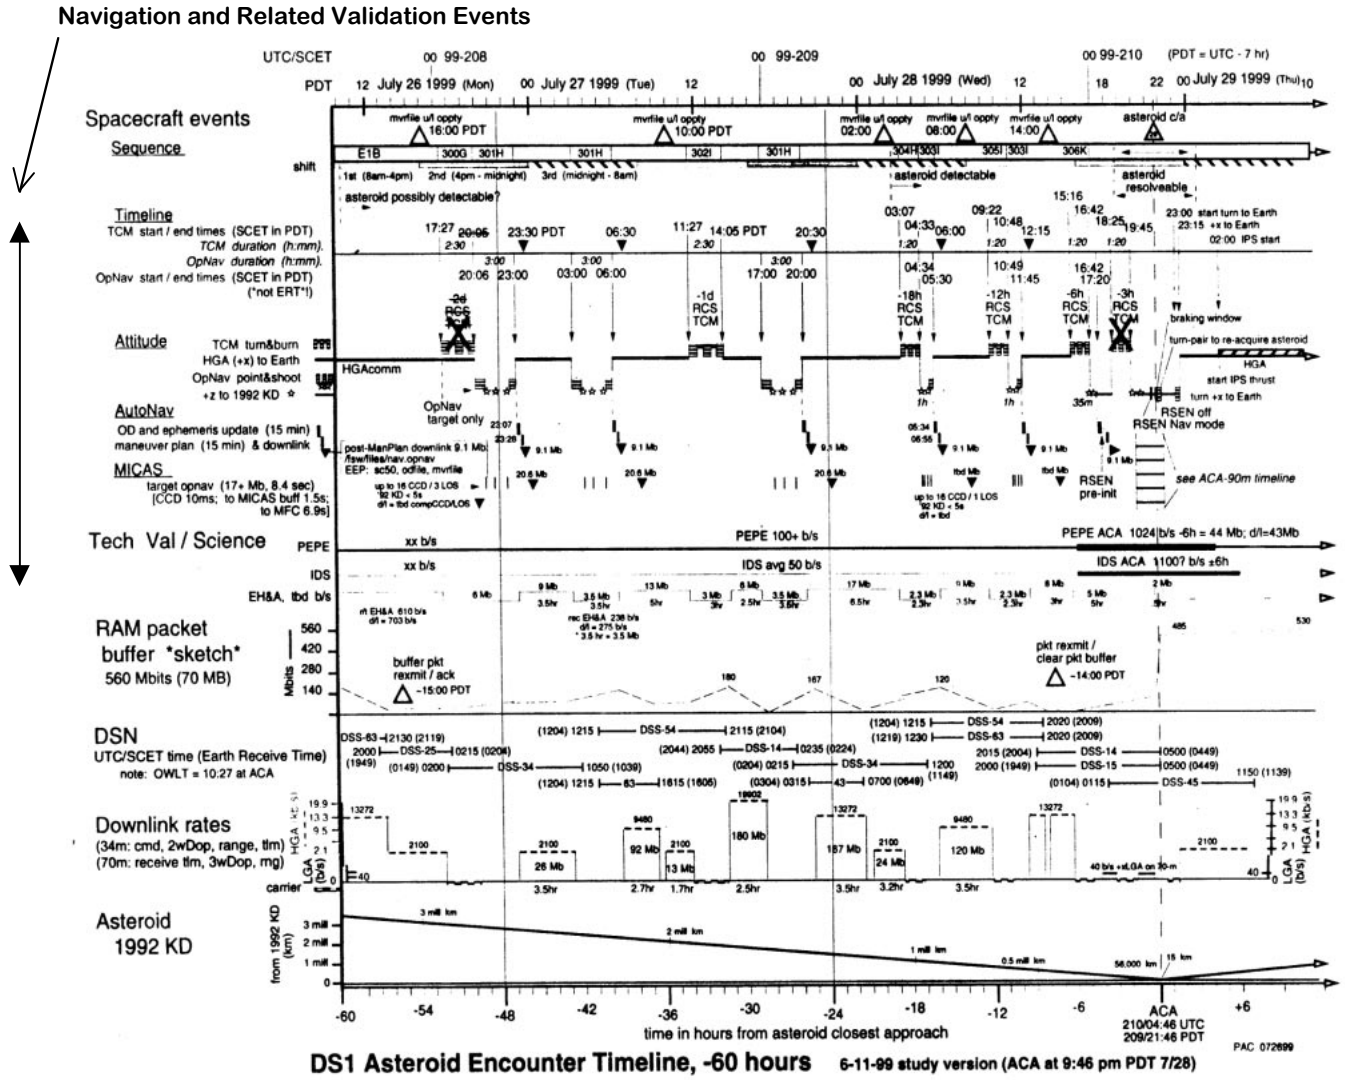

Figure 2. Mission and AutoNav Operations/Validation Plan and Schedule for Braille Close Approach

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 trans-mission 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.

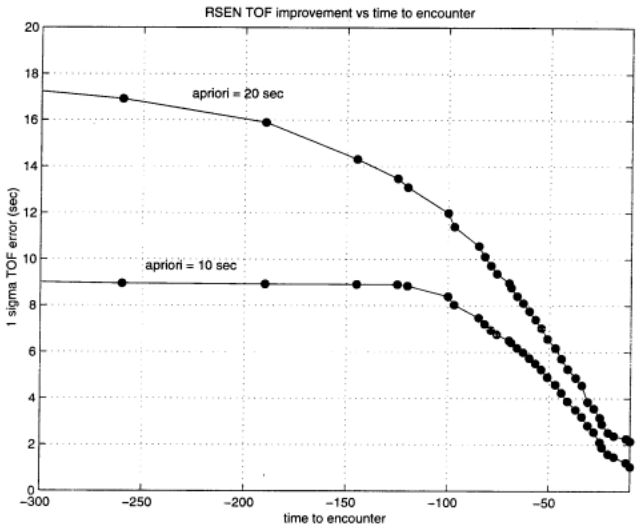

Figure 3. RSEN Time-of-Flight Performance

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 the 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.

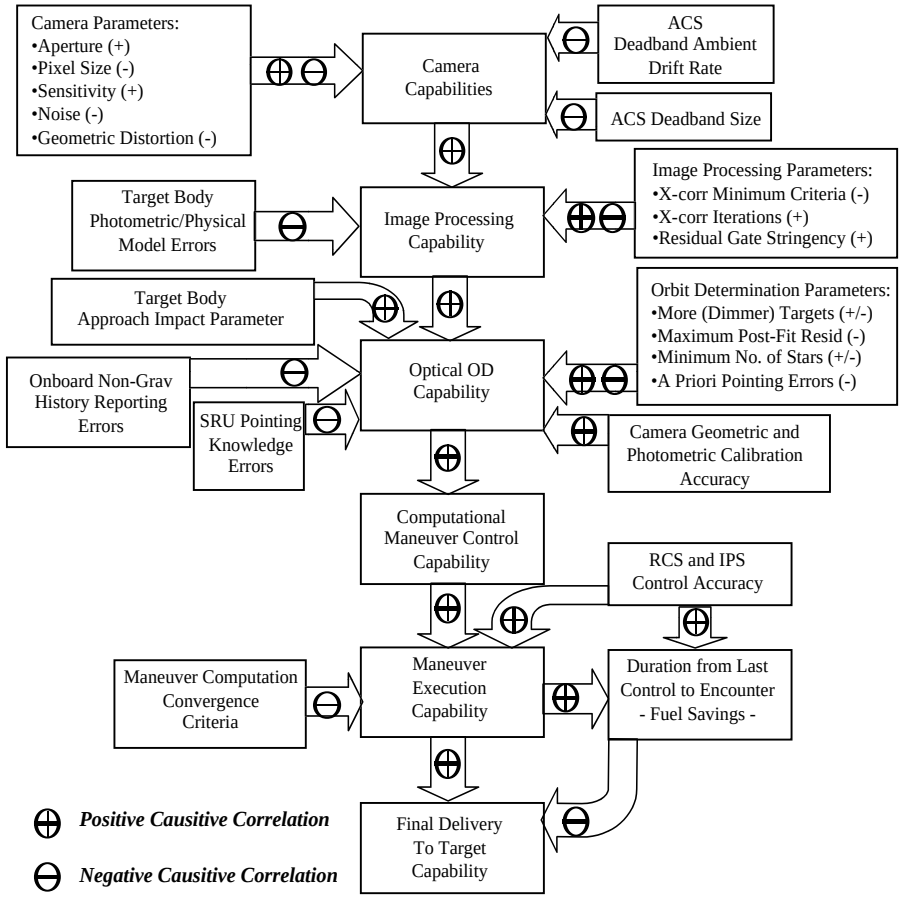

Figure 4. Subsystem Performance Influence on AutoNav

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 Astrographic-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

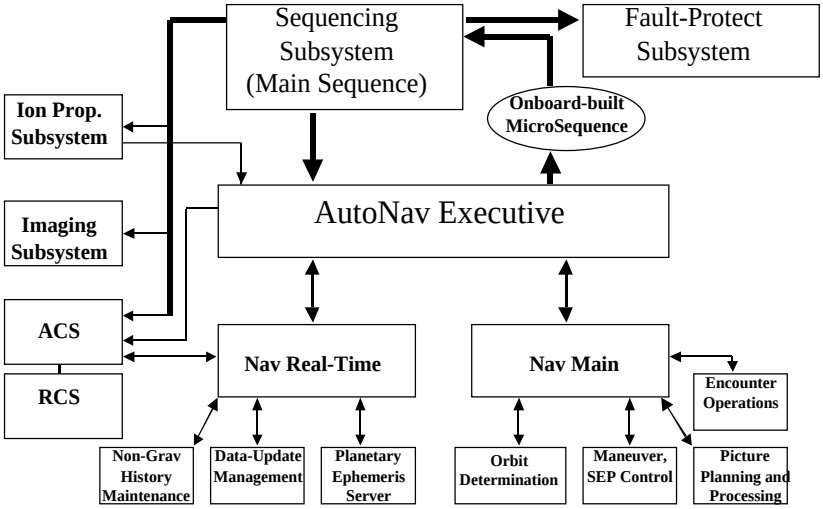

Figure 5. The AutoNav Software System and Interacting System Software

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_Encntr	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 Auto-Nav 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

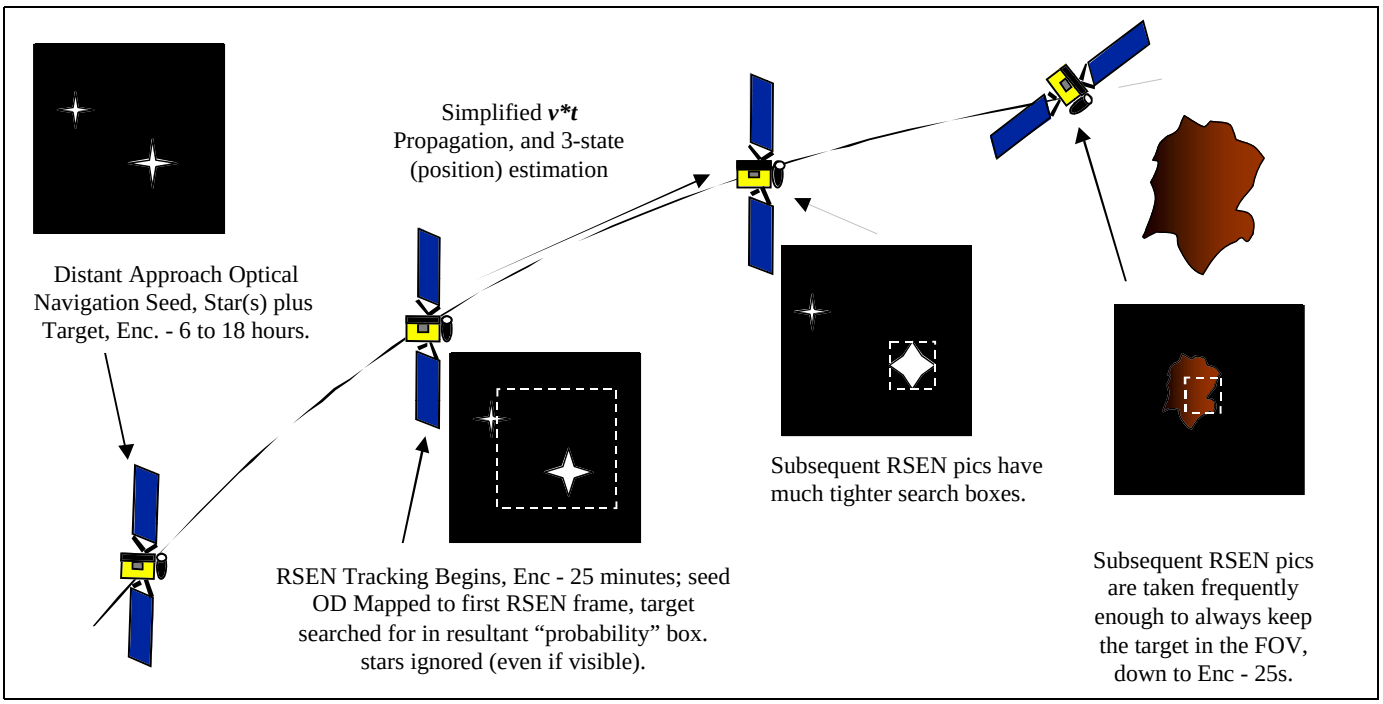

Figure 6. Reduced State Encounter Navigation Schematic Functional Overview

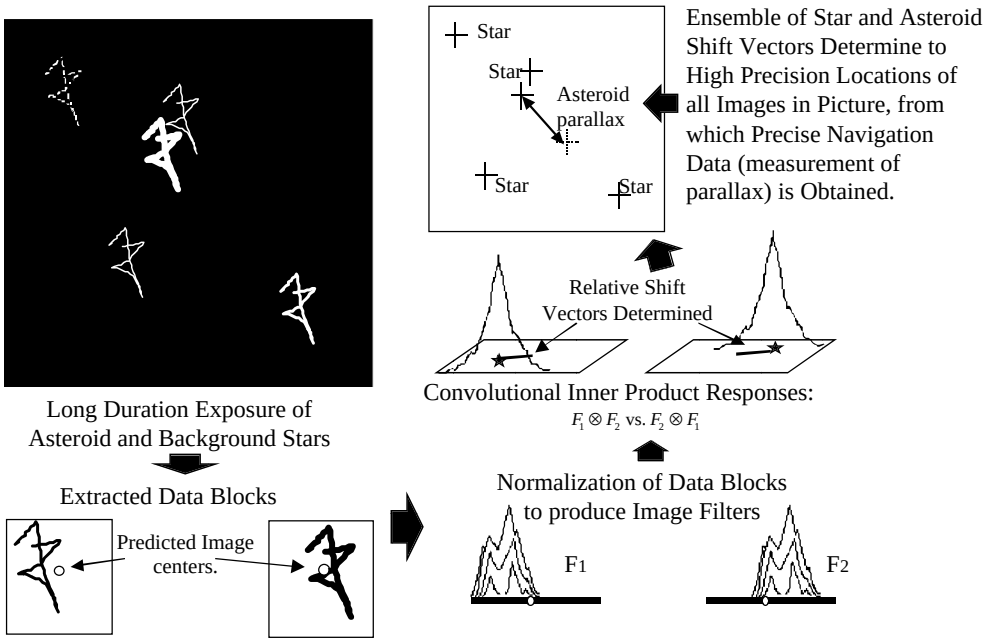

Figure 7. Multiple Cross Correlation of Asteroid and Stars

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* 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

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

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

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

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

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

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

(b)

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

$$\mathbf{H}(t) = \begin{bmatrix} \partial p/\partial \mathbf{r} & 0_{1x7} \\ \partial l/\partial \mathbf{r} & 0_{1x7} \end{bmatrix}$$

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

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

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

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

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

X_e = Target encounter conditions $(B \cdot R, B \cdot T, tof)$

ΔX_e = Deviations from desired encounter

α_i = Rt. ascention of thrust segment i

δ_i = Declination of thrust segment i

τ_κ = Burn duration of final segment κ

$\Delta X_e = K\Delta s$, where...

$$K^T = \begin{bmatrix} (\partial X_e / \partial \alpha_{k_1}) \\ (\partial X_e / \partial \delta_{k_1}) \\ (\partial X_e / \partial \alpha_{k_1+1}) \\ (\partial X_e / \partial \delta_{k_1+1}) \\ \vdots \\ (\partial X_e / \partial \alpha_{k_2}) \\ (\partial X_e / \partial \delta_{k_2}) \\ (\partial X_e / \partial \tau_\kappa) \end{bmatrix}, \text{ and... } \Delta s = \begin{bmatrix} \Delta\alpha_{k_1} \\ \Delta\delta_{k_1} \\ \Delta\alpha_{k_1+1} \\ \Delta\delta_{k_1+1} \\ \vdots \\ \Delta\alpha_{k_2} \\ \Delta\delta_{k_2} \\ \Delta\tau_\kappa \end{bmatrix}$$

Then, by the calculus of variations,

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

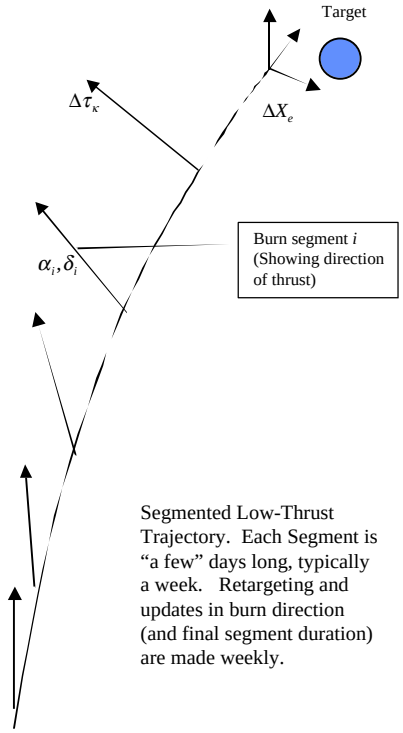

Figure 9. Adjusting a Low-Thrust Burn Arc

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	≥2 µrad	7 µrad	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 µrad 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 µrad. 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 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.

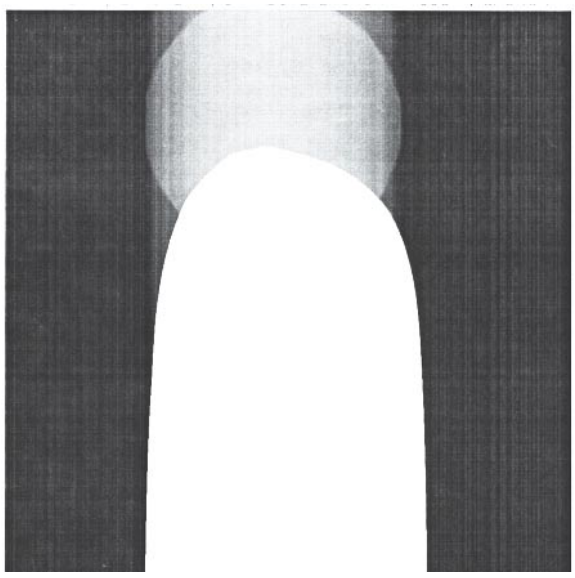

Figure 10. MICAS Extended Bright-Image Charge Bleed

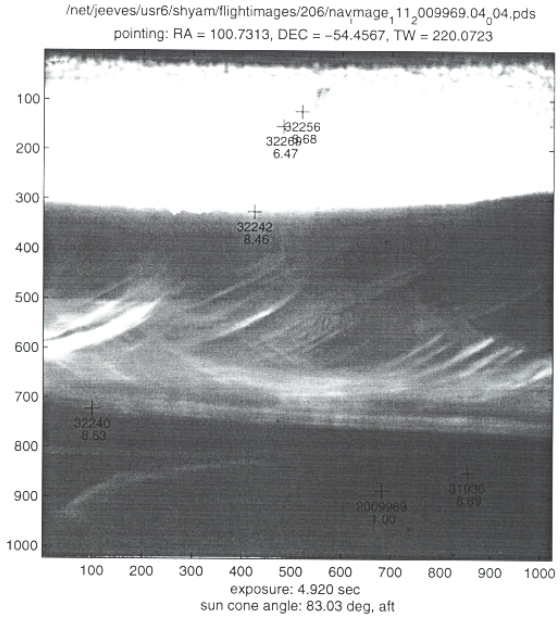

Figure 11. MICAS "Low Solar Cone Angle" Scattered-Light Picture

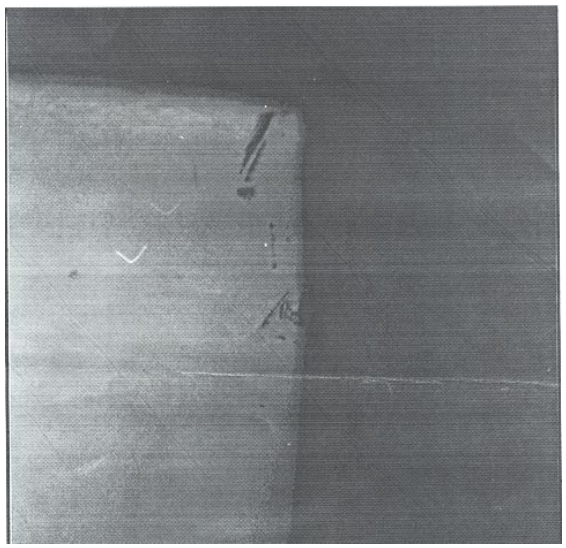

Figure 12. MICAS "High Solar Cone Angle" Scattered-Light Picture

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:

```
NJPL1I00PDS                 = 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
0000000000000000
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 descoped 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.

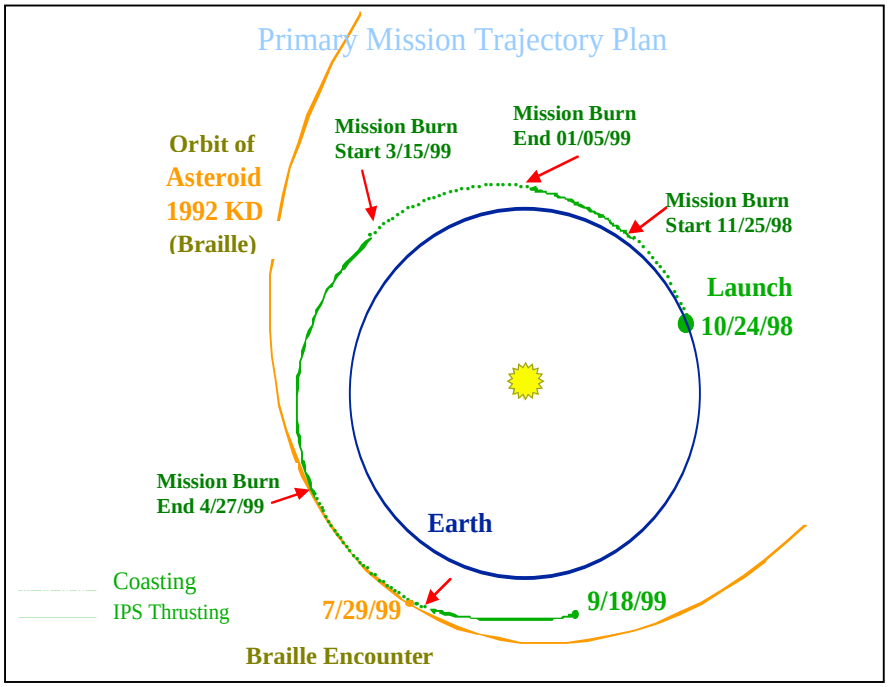

Figure 13. Primary Mission Trajectory Plan

- 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 activites began a period of 10 weeks of "normal" operations, which included weekly Photo-Op/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

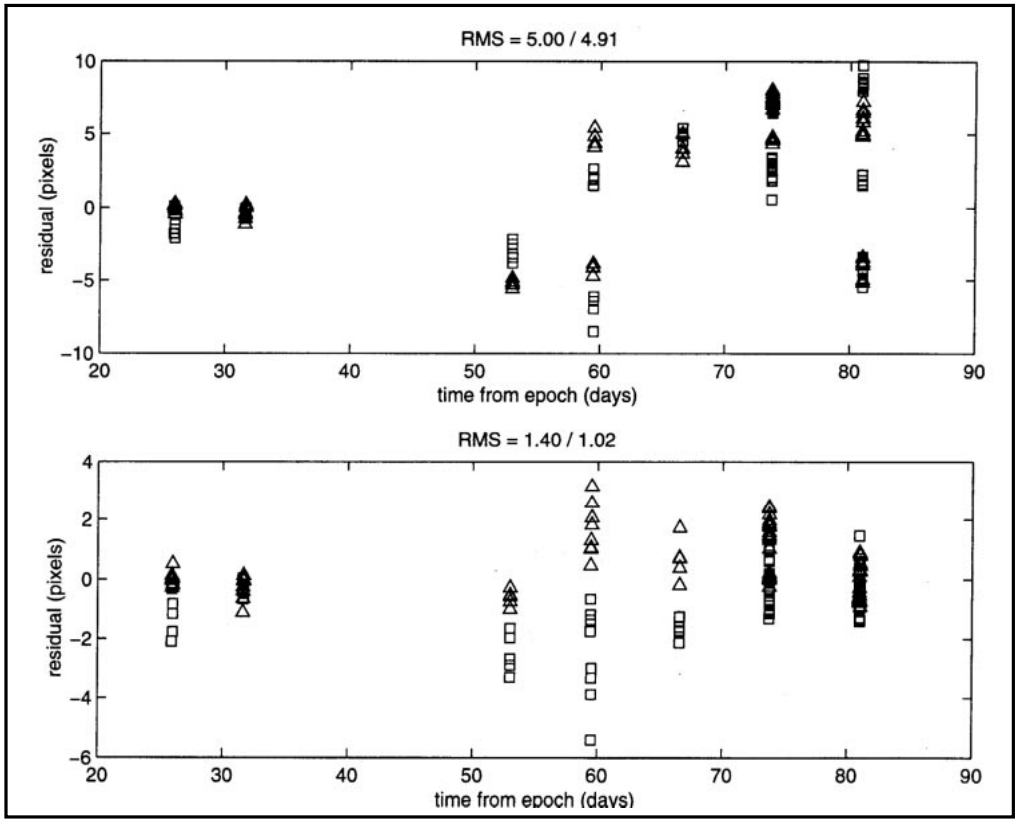

Figure 14. Pre-(upper) and Post-(lower) Fit Residuals from 3/22/99 Optical Solution

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 nonliearity 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.

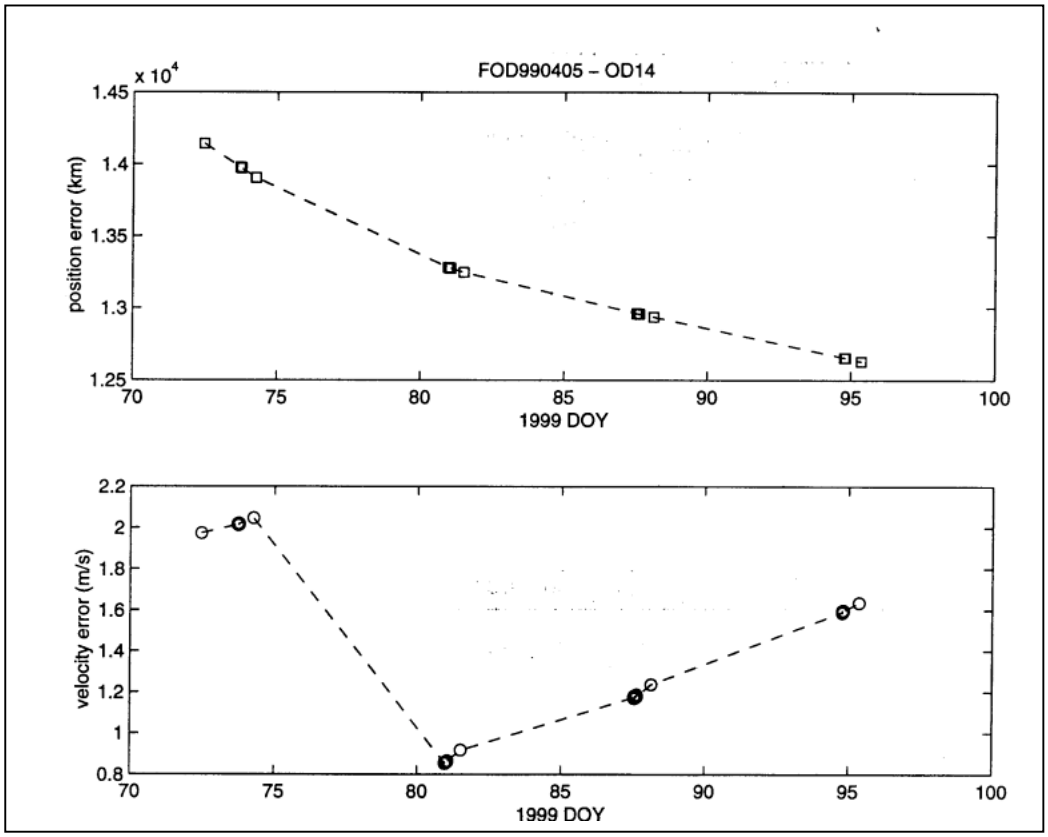

Figure 15. Flight vs. Ground-Orbit Determination April 5, 1999

- 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.

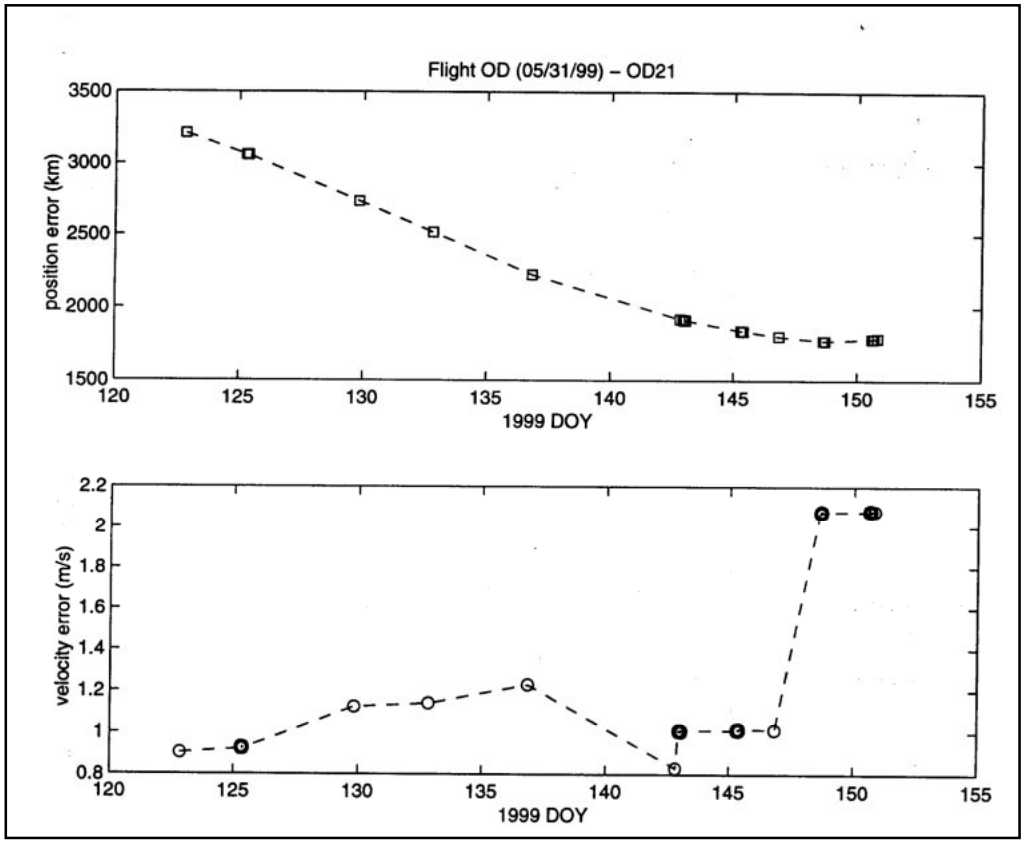

Figure 16. Flight vs. Ground-Orbit Determination 5/31/1999

- 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 per-formance 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 ΔVs 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	1st Close Approach Sequence	SEQ51
–2.5 min	1.5 min	2nd Close Approach Sequence	SEQ52
–90 s	65 s	3rd Close Approach Sequence	SEQ53
–25 s	25 s	4th Close Approach Sequence	SEQ54

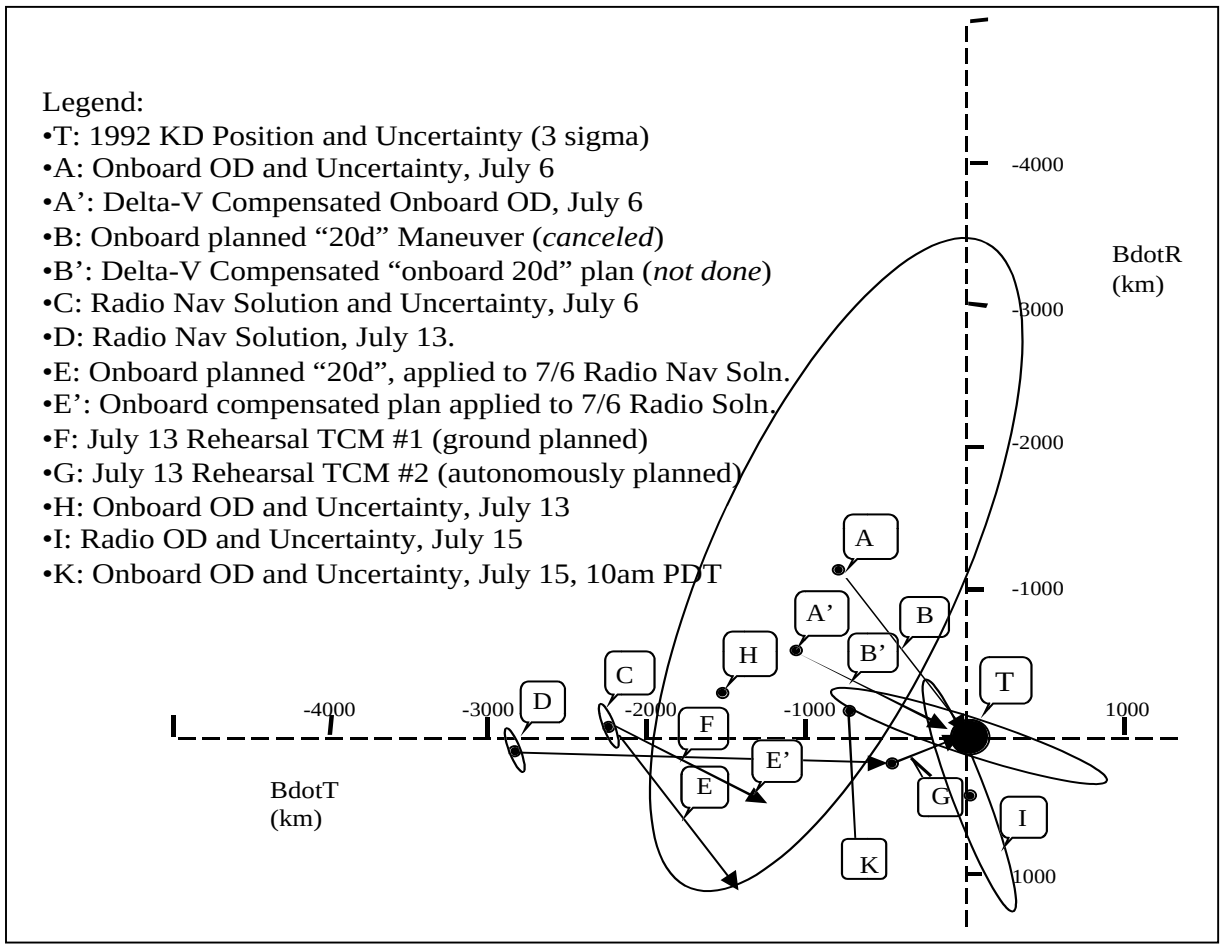

Figure 17. Current B-Plane Target Conditions at the – 20 – 10 Day TCMs: Decision Data from 7/15/99

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 serendipidously), 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.

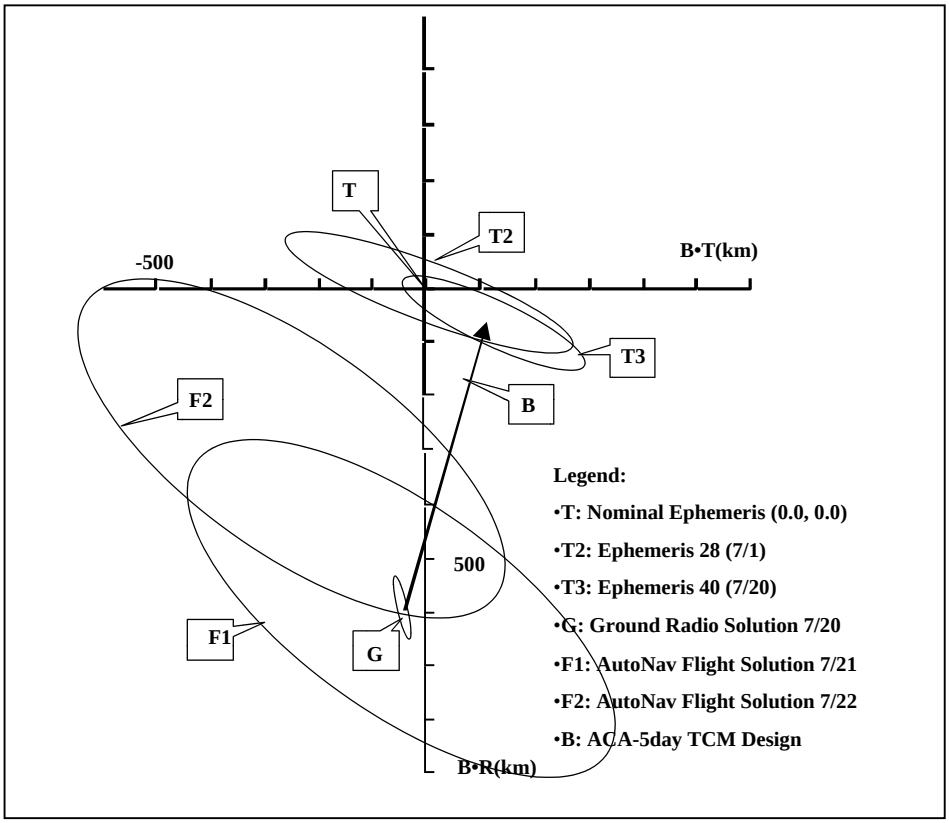

Figure 18. Minus 5 Day TCM Solutions

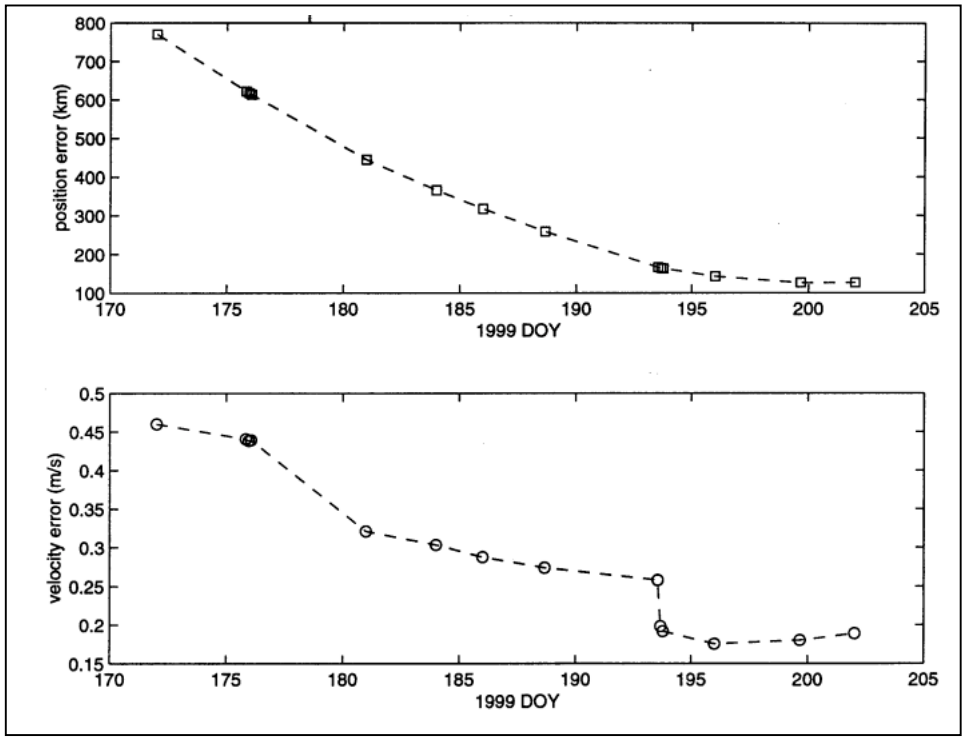

Figure 19. Flight OD vs. Ground OD#37, 7/21/99

- 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 where 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

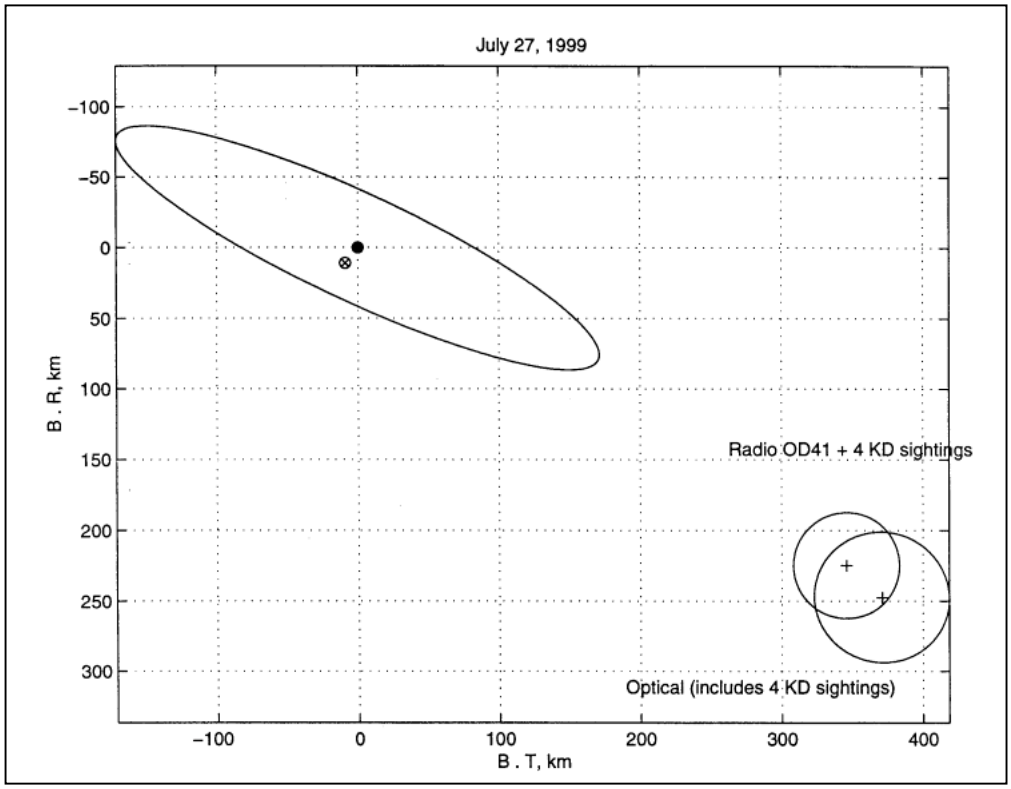

Figure 20. Pre-Minus –1 Day TCM, "Flight OD" Braille B-Plane

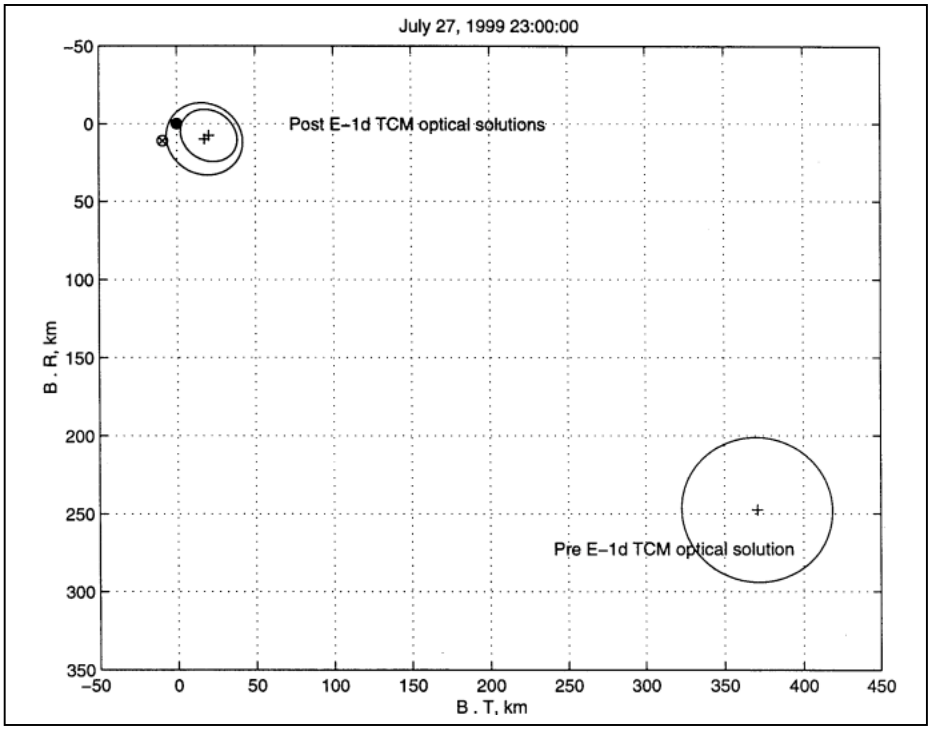

Figure 21. Pre-ACA – 6 hour TCM B-Plane July 27

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

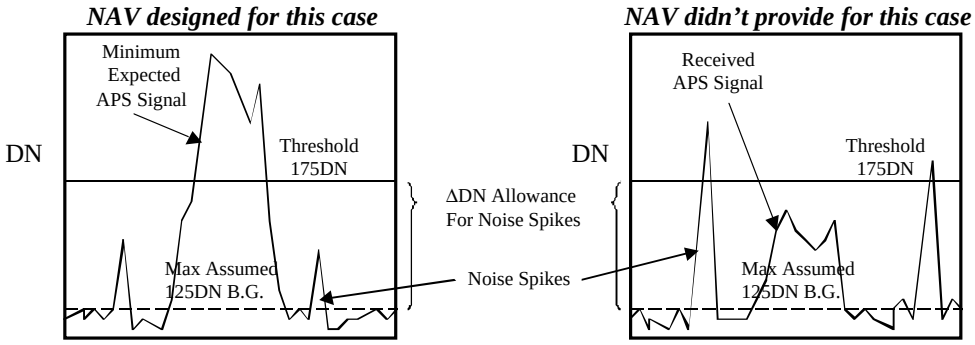

Figure 22. Diagrammatic View of Received RSEN Signal

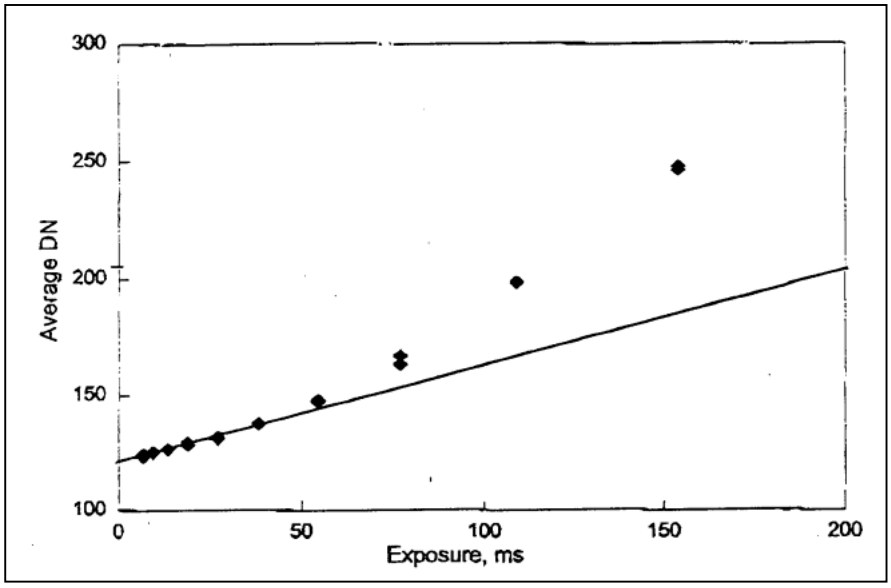

Figure 23. MICAS APS Channel Non-Linear Signal Response

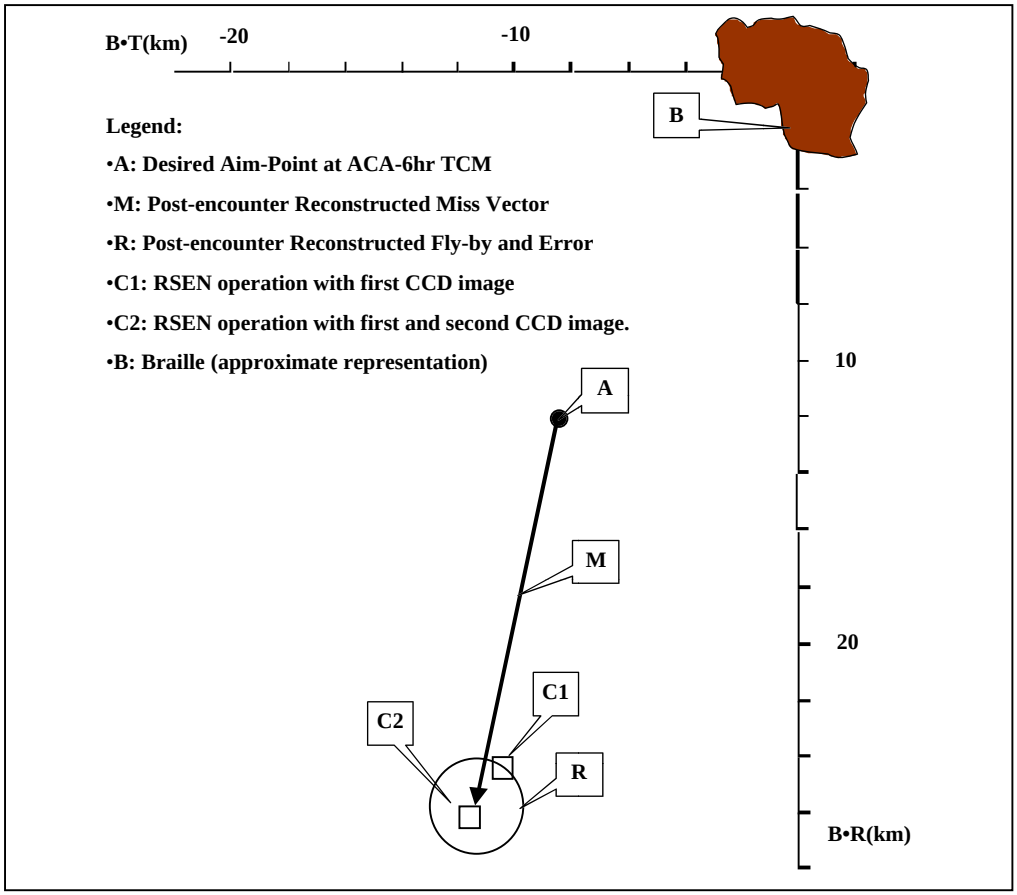

Figure 24. Encounter Results Using Post-Encounter CCD Braille Pictures

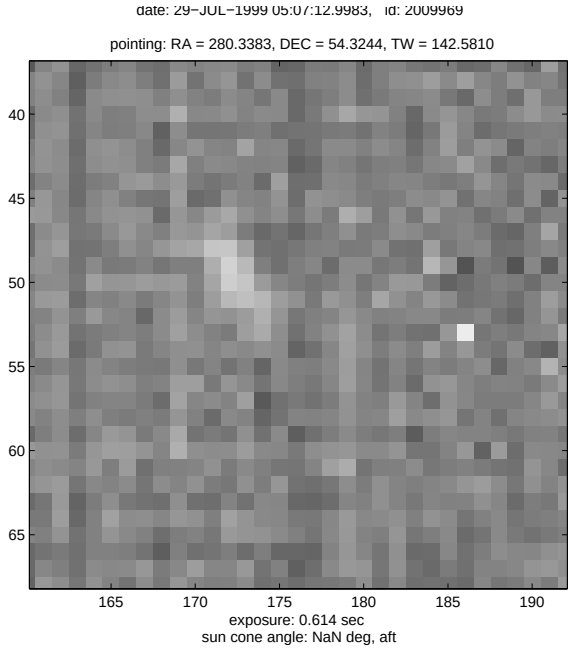

Figure 25. Post-Encounter APS Image of Braille

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

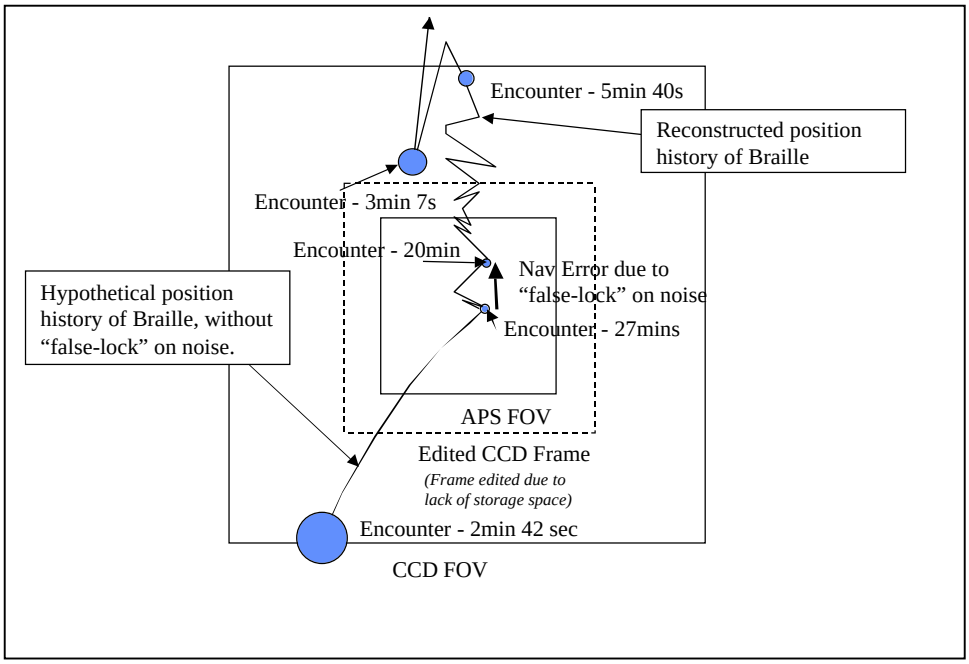

Figure 26. Reconstructed Nominal vs. Perturbed Braille Field-of-View Flight Path

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 evealed 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 begun 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.

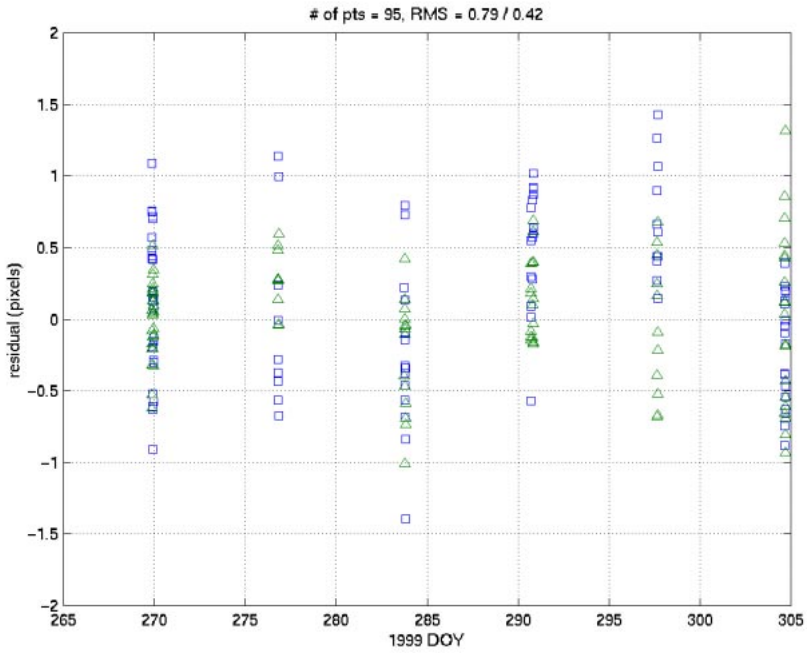

Figure 27. Post-Braille AutoNav DataArc and Residuals

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 unfavorble 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 AutoNav 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 Desireable 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 repoint 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 Astro-dynamics 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	InvldEpemReq
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 *DS1* 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 DS1* 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). *DS1* 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 Navi-gation and how it will be applied to *DS1* 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 *DS1*, the capability has been maintained to quickly intervene with, and modify the behavior of the system if mission emergencies require.

DS1 MISSION ATTRIBUTES

An overview of the New Millennium Program and *DS1* in particular is given in Ref. 2. The *DS1* 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 (WKI). 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.

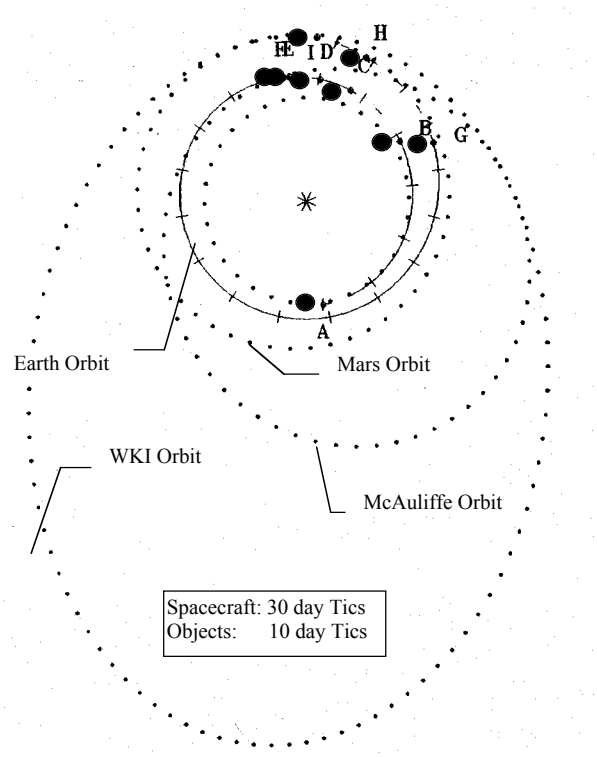

Figure 1. *DS1* Mission Design

For the McAuliffe flyby, the *DS1* 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	*DS1* 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	McAullife 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	WKI encounter

Table 1. Principal DS1 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 *DS1* 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

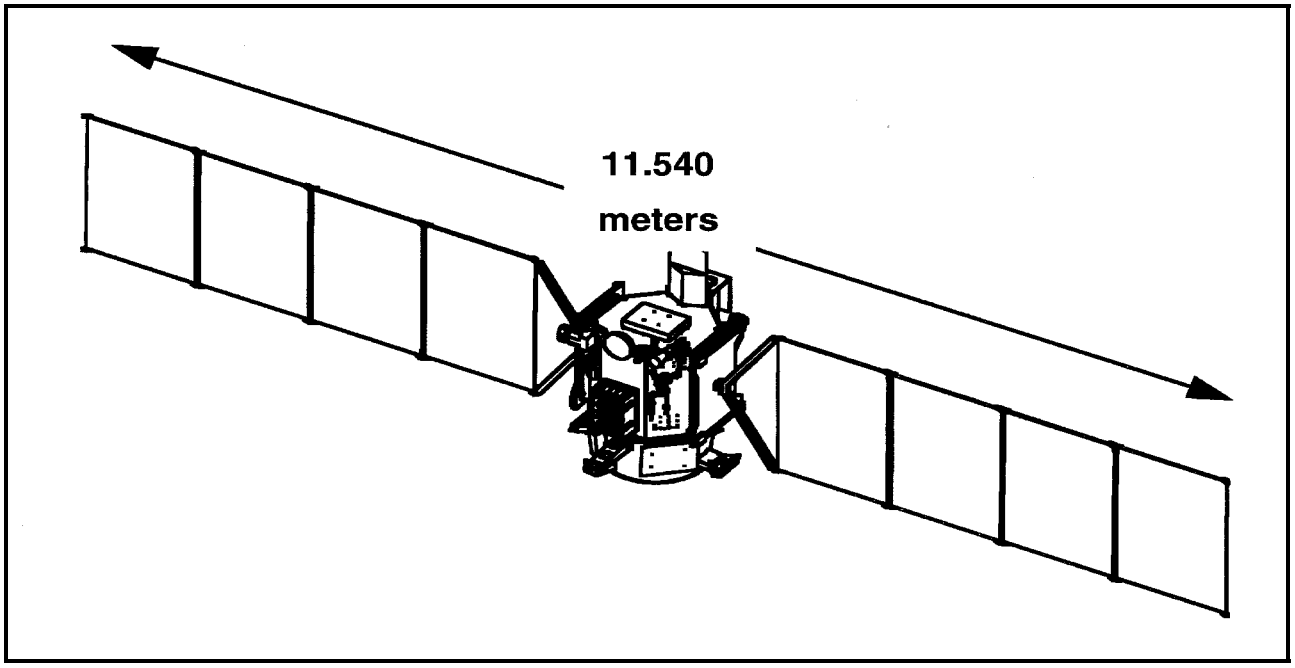

Figure 2. *New Millennium DS1* Spacecraft

technology, and one that has a fundamental influence on the nature of the *DS1* 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 *DS1* 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 *DS1*. 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.

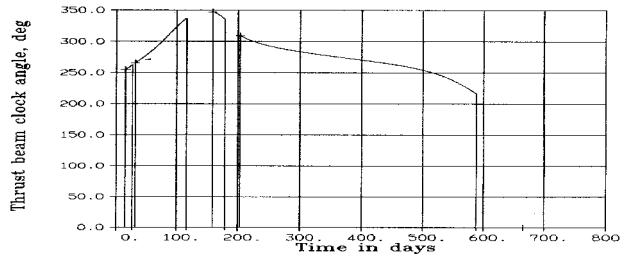

Figure 3. IPS Thrust Clock Beam Angle

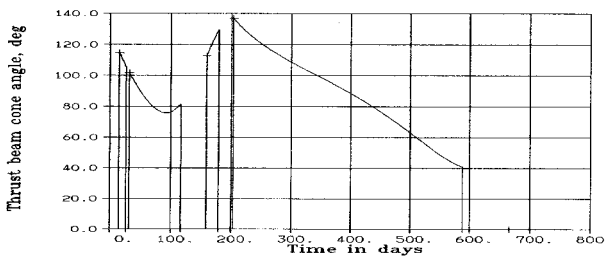

Figure 4. IPS Thrust Beam Cone Angle

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.

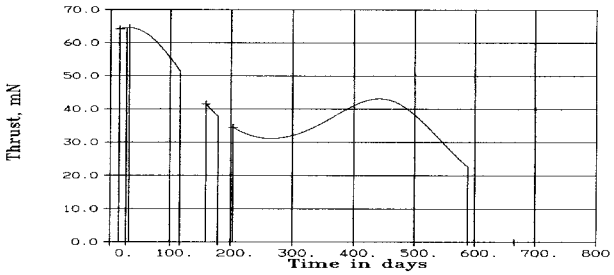

Figure 5. IPS Thrust Magnitude

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 *DS1* 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."

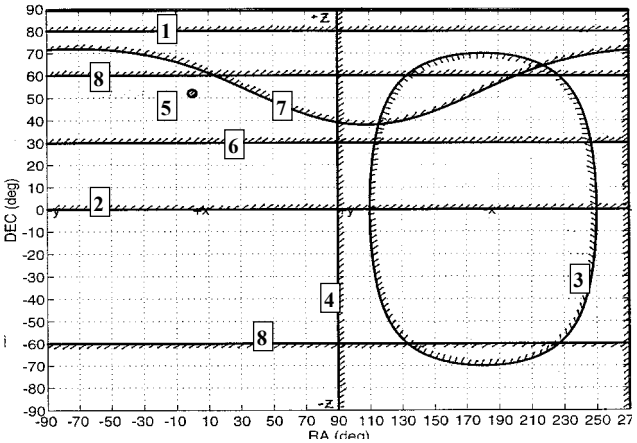

Figure 6. Illumination-Forbidden Regions of Spacecraft Body.

#	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: *DS1* 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 *DS1*) 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 flow by *DS1* 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 *DS1* 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 ephemeridies. 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 *DS1* 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 *DS1* 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 *DS1*. This compares favorably with the 7 to 10 individuals necessary to perform similar functions for the *Cassini, Galileo* and *Voyager* missions.

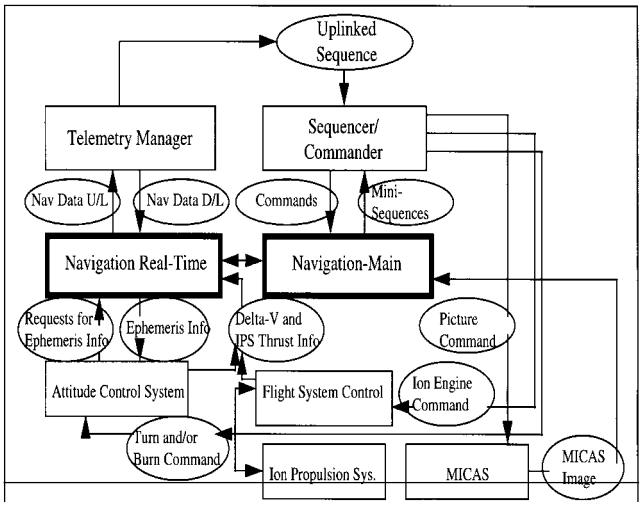

Figure 7: Navigation System Architecture

AUTONOMOUS NAVIGATION SYSTEM DESIGN:

Architecture
The *DS1* software system architecture, emphasizing the navigation system interactions, is shown in Fig. 7. The *DS1* 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-UDATE]	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 ephemeridies 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 *DS1*, 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 x n* in size denoted as ***F,*** and a sample area ***S****, M x N* in size, of which subset regions of *m x n* dimensions are extracted, then a function ***cij*** 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 $\mathbf{c}_{ij}$ represents the position of best match between ***F*** and the sample region

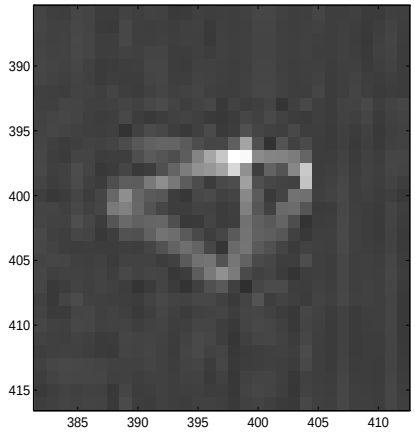

Figure 8. Simulated Cruise Asteroid Image

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 *DS1* 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 *DS1* 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 McAullife 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 signifi-cantly 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.

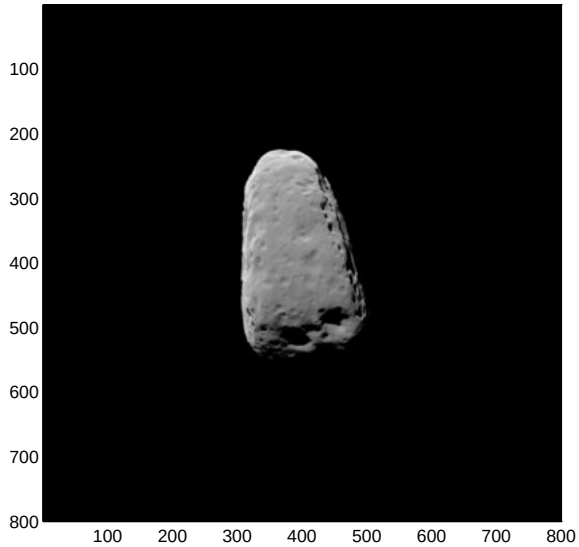

Figure 9: Simulated Asteroid Approach Frame

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/sec^2. 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/sec^2. The IPS engine thrust is believed to be predictable to about one percent, or about 1.2 mm/sec^2. 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.

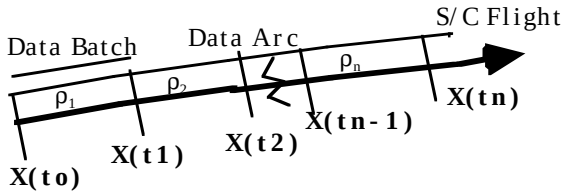

Figure 10. Schematic of Orbit Determination Data-arc Structure

Fig. 10 shows the subdivision of the data arc into batches over which an estimate parameter set is constant. $\mathbf{X}(t_O)$ 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 **A** can be formed:

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

$\mathbf{O}_n$ is the observation vector for observation n, and is a 2x1 vector, (pixel and line). The formulation of $\partial\mathbf{O}/\partial\mathbf{X}$ is documented elsewhere (Refs. 9,10). **q** is a vector of estimable parameters, and for batch 1, $\mathbf{q} = [\mathbf{X}(t_0),\rho_1,\mathbf{S}]$. **A** is combined into a covariance matrix referenced to t_o, Γ_{to}, 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 **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 $\mathbf{q}_n = [\mathbf{X}(t_0), \rho_1, \rho_2 \ldots \rho_m, \ldots \rho_n, \mathbf{S}]$:

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

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 **A** and Γ_{to}, an estimate of the parameters can be made:

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

where,

$$\Delta y_{1\times 2N} = O_{2\times N} - C_{2\times N}$$

where $\Delta\mathbf{y}$ is the residual vector formed as the difference between the observation vector **O** and the computed predicted value **C**. **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:

$$\Gamma_{t_{n/2}} = D\Phi_{t_0}^{t_{n/2}} \Gamma_{t_0} \Phi_{t_0}^{t_{n/2}\prime}$$

where $\Phi(\mathbf{t_0},\mathbf{t_{n/2}})$ is the state transition matrix from $\mathbf{t}_0$ to the midpoint of the data arc. 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 28days, 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 ephemeridies 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 ephemeridies are represented as Chebyshev function polynomials of varying order. This follows the standard representation of the planetary ephemeridies 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, **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 **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 **s,** is formed *via* the equations:

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

where,

$$K = \left[\frac{\partial X}{\partial \alpha_1} \quad \frac{\partial X}{\partial \delta_1} \quad \frac{\partial X}{\partial \alpha_2} \quad \frac{\partial X}{\partial \delta_2} \quad \cdots \quad \frac{\partial X}{\partial \alpha_m} \quad \frac{\partial X}{\partial \delta_m} \quad \frac{\partial X}{\partial \tau} \right],$$

$$\Delta s' = \left[\Delta\alpha_1 \quad \Delta\delta_1 \quad \Delta\alpha_2 \quad \Delta\delta_2 \quad \cdots \quad \Delta\alpha_m \quad \Delta\delta_m \quad \Delta\tau \right],$$

and,

$$\Delta X' = \left[\Delta B \cdot R \quad \Delta B \cdot T \quad \Delta ltof \right]$$

Δ[B•R,B•T,*ltof*] 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 McAullife (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 vectoraly 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.

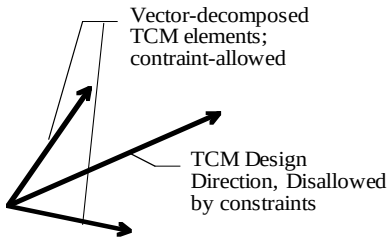

Figure 11. TCM Vector-decomposition

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 McAullife-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 infra-red 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 crosstrack 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 obs-ervations	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

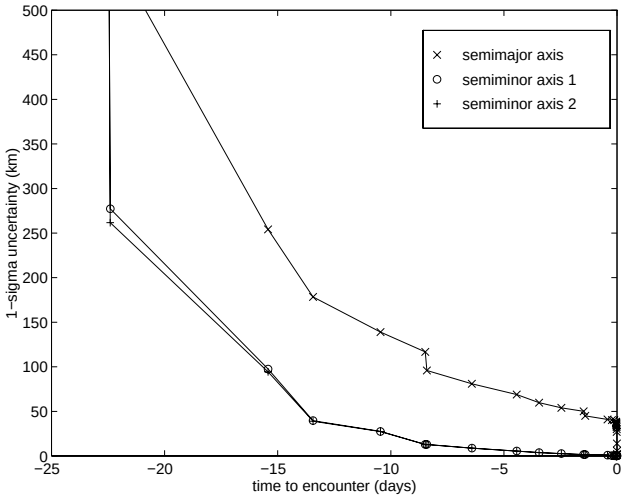

Figure 12: Autonomous Navigation System Orbit Determination Performance, Far Encounter

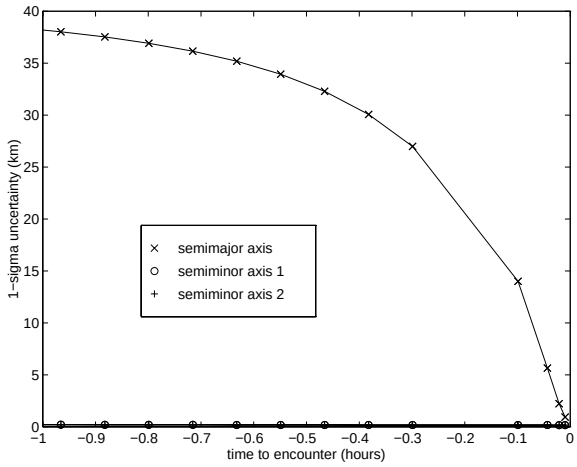

Figure 13: Autonomous Navigation System Orbit Determination Performance, Near Encounter

FIRST USE OF THE NAVIGATOR

Though in theory the *DS1* 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 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 *DS1* operations, does imply that the cost savings of navigational autonomy will not be seen on *DS1*. 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 *DS1* Mission Engineer, under whose auspices much of the *DS1* Navigation development has occurred, Abdullah Aljabri who lead the *DS1* 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, Hadeaegh, F., "Autonomous Guidance and Control for the *New Millinnium DS1* 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 Coordianate 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 sytem, 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 thrusted 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 μrad.

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:

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

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:

$$\dot{\mathbf{r}} = \mathbf{v} \tag{2}$$

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

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 ith perturbing planetary body

$\mathbf{r}_{ri}$ = the position of the spacecraft relative to the ith 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 ith 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 recieves 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 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 explictly 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 ΔVs 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 $\Delta\mathbf{V}$ 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 $\Delta\mathbf{V}$, and restart the integration. The accuracy of this method remains only if the RCS ΔVs are small (on the order of m/s) and therefore take only a short time for the thrusters to achieve; large RCS ΔVs 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

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

where

X^*, Y^*, Z^* = the cartesian position components,

$\dot{X}^*, \dot{Y}^*, \dot{Z}^*$ = the cartesian velocity components,

$k_1^* \ldots k_n^*$ = 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

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

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, $\mathbf{\Phi}$, as

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

where

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

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

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

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

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

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:

$$\tilde{\mathbf{H}} = \mathbf{H}\mathbf{\Phi}, \tag{10}$$

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:

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

and

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

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 $\mathbf{U}\,\mathbf{D}\,\mathbf{U}^{\rm T}$, 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

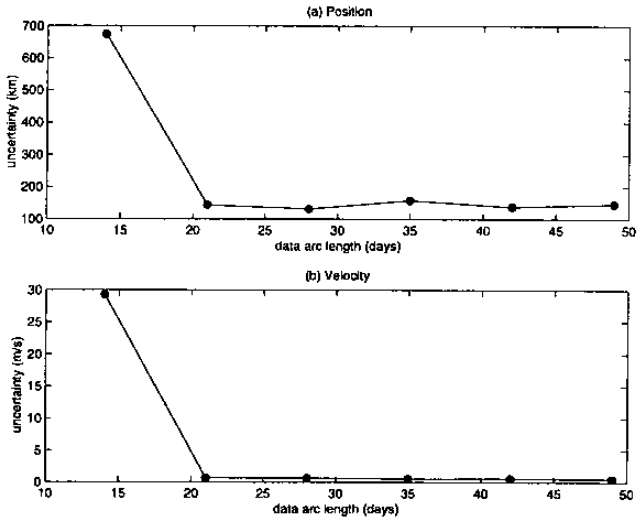

Figure 1: 1σ Uncertainties in Position and Velocity vs Data Arc Length

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 noticable 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

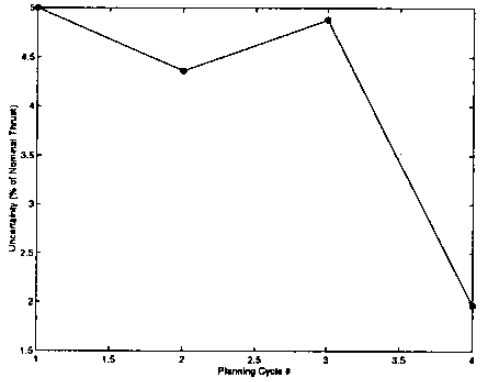

Figure 2: 1σ Uncertainties in Thrust Scale Factor Estimates

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

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

where

$$\begin{aligned} A_T &= \text{the amplitude of the additional thrust magnitude,} \\ t, t_0 &= \text{the current and epoch time, and} \\ \tau_T &= \text{the time constant of the magnitude variation.} \end{aligned}$$

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.*,

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

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.

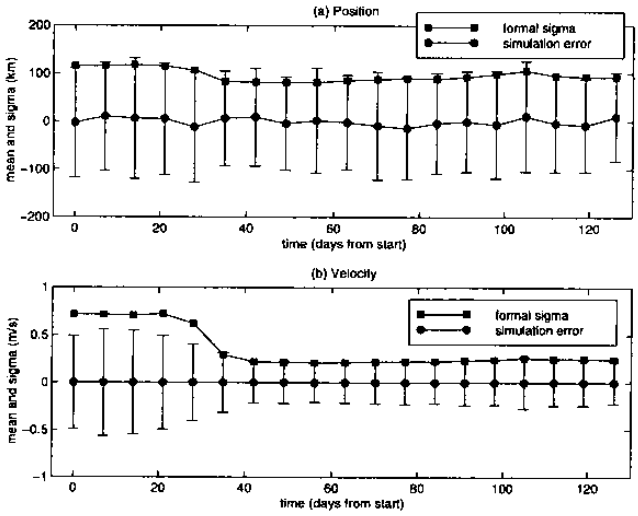

Figure 3: 1σ Formal Uncertainty from Filter and Actual Statistics from Simulations for Cruise

- 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 achievably 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.

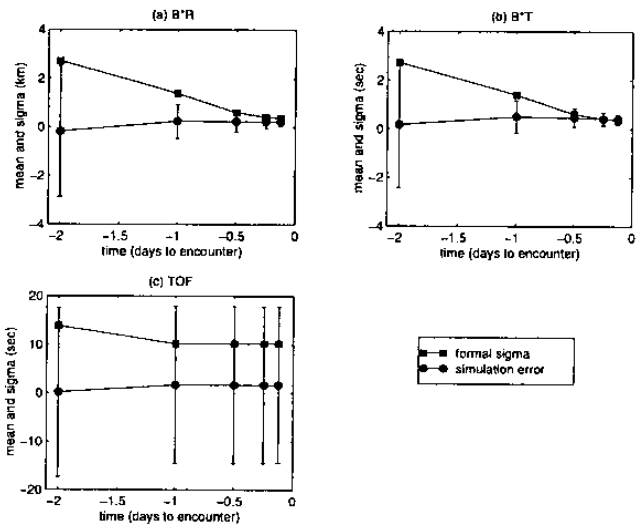

Figure 4: 1σ Formal Uncertainty from Filter and Actual Statistics from Simulations for Encounter

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.

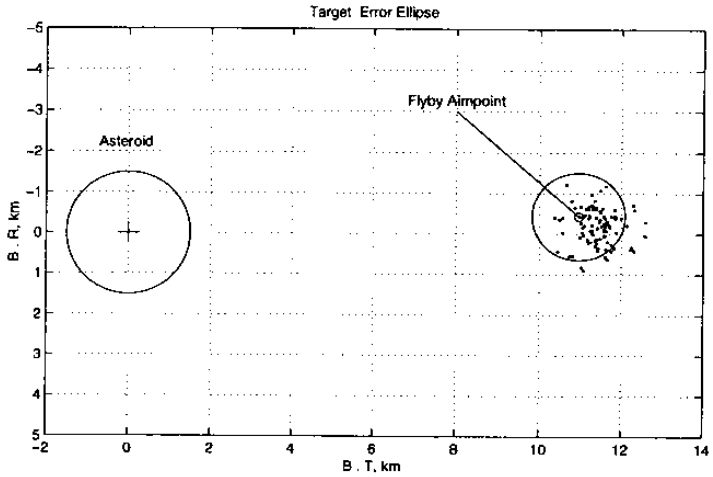

Figure 5: True B-plane Delivery Locations from Monte Carlo Simulations

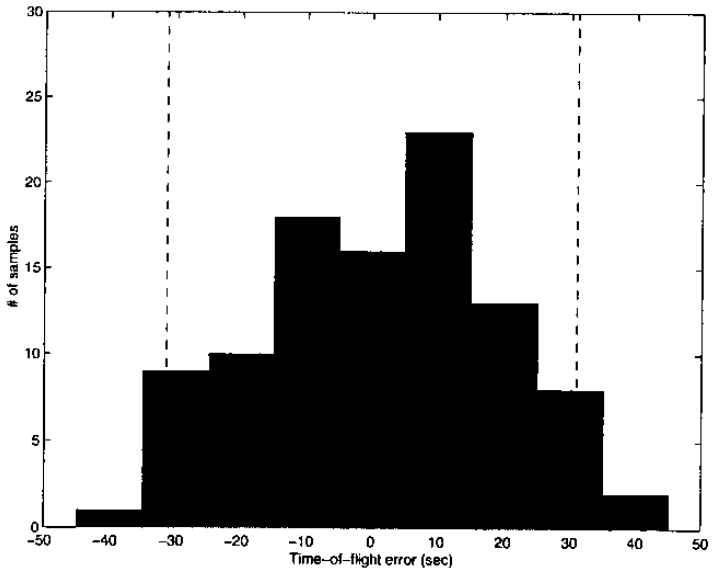

Figure 6: True Time-of-flight Errors from Monte Carlo Simulations

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

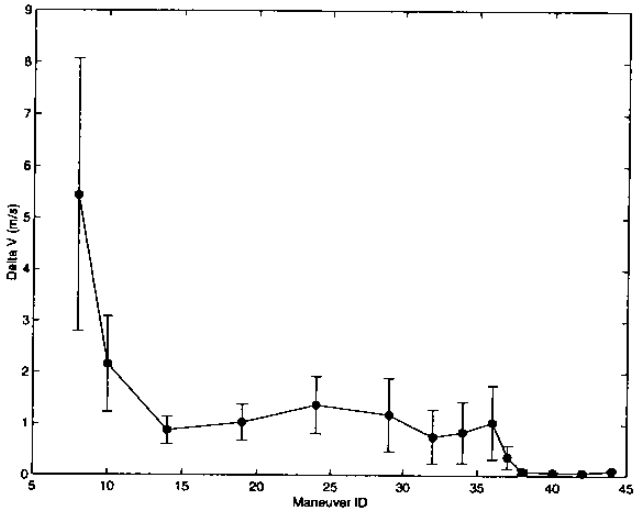

Figure 7: ΔV Statistics per Maneuver

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

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) mission[1], 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 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 system[5]. 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

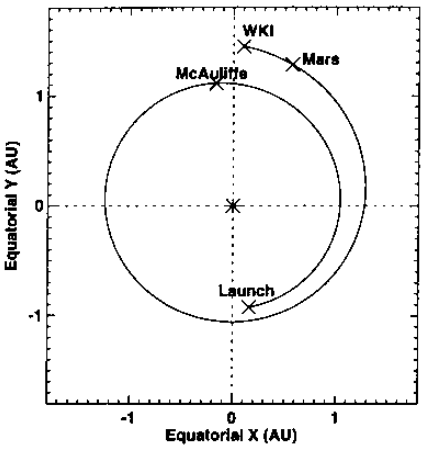

Figure 1: Sample DS-1 Trajectory to McAuliffe and West-Kahoutek-Ikemoura

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.

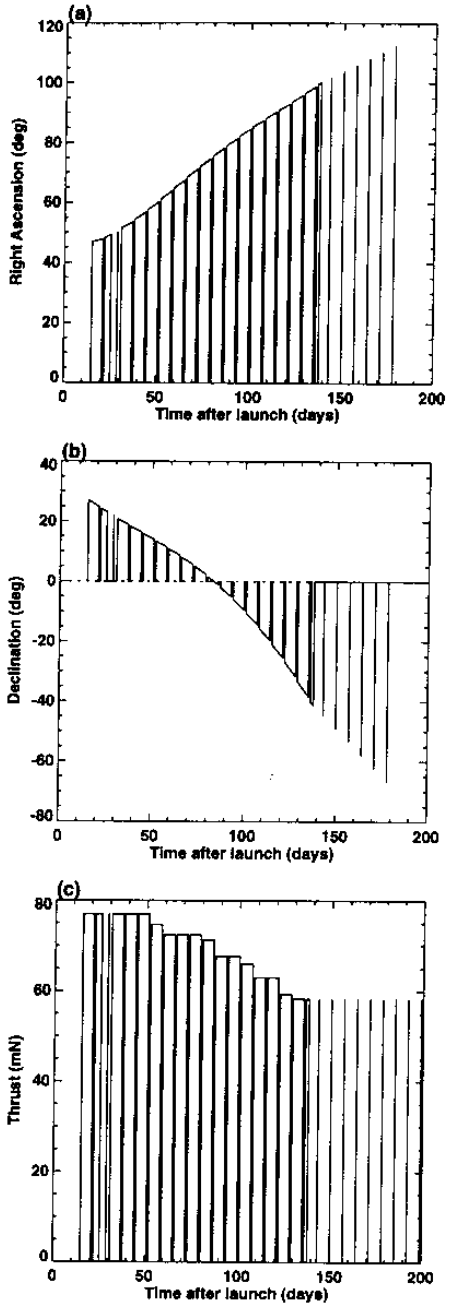

Figure 2: Right Ascension (a), Declination (b), and Magnitude (c) of SEP Thrust Pointing Vector for DS-1 Trajectory to McAuliffe

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

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

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

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.

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

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

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, $\dot{\alpha}_k$, δ_k, $\dot{\delta}_k$, 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_κ 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 $X_e(\alpha_k, \delta_k, \tau_\kappa)$ is a function of,

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

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 e, then adjustments to the parameters τ_κ, 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,

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

where the overbars ($\bar{}$) are used to indicate the adjusted SEP profile variables that are necessary to achieve the required target state. It is these variables, $(\bar{\alpha}_k, \bar{\delta}_k, \bar{\tau}_\kappa)$ 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.

$$k_1 \leq k_2 \leq \kappa \tag{9}$$

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.

$$\Delta X_e = K \Delta s \tag{10}$$

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.

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

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 $X_e(\alpha_k, \delta_k, \tau_\kappa)$.

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

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.

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

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

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

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.

$$\Delta X_e^T = \begin{bmatrix} \Delta B \cdot R & \Delta B \cdot T & \Delta TOF \end{bmatrix} \tag{15}$$

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 of 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.

$$|\, \Delta X_e \,| < e \tag{16}$$

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.

$$\Delta s = K^{-1}\Delta X_e \tag{17}$$

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.

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

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

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

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 e.

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:

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

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

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

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 $\delta J = 0$.

$$\Delta s^T - \lambda K = 0 \tag{22}$$

$$\Delta X_e - K\Delta s = 0 \tag{23}$$

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.

$$\lambda^T = (KK^T)^{-1}\Delta X_e \tag{24}$$

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

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.

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

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 $\hat{\bar{p}}$ 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 $\hat{p}'$.

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

$$\hat{\bar{p}} \cdot \hat{p} = \cos\left[\theta_k - \cos^{-1}(\hat{p}' \cdot \hat{p})\right] \tag{28}$$

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 e. 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.

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

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.

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

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.

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

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

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).

$$\Delta Y = K_A \Delta s \tag{33}$$

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

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: $(N + N_A) = M$

$$\Delta s = K_A^{-1} \Delta Y \tag{35}$$

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

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

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

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

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 redesigned 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

97-3819

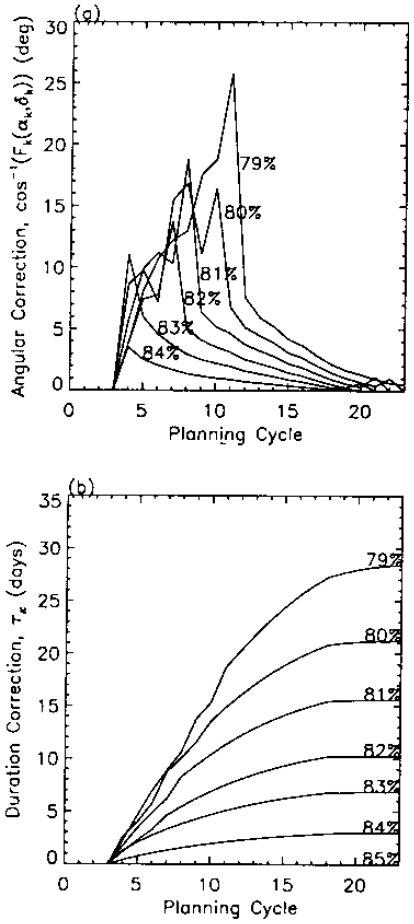

Figure 3: Angular corrections (a) and duration corrections (b) with respect to the designed SEP profile using a minimum of 1 SEP plan to correct prior errors in the SEP profile. No angular constraints are applied to the corrections.

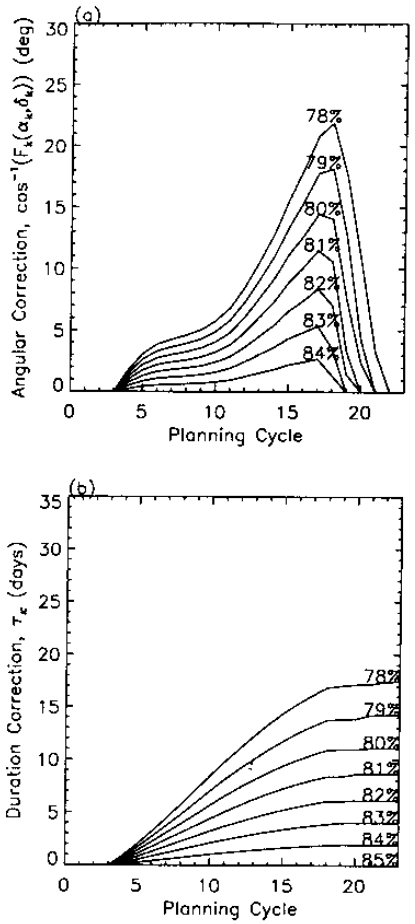

Figure 4: Same as Figure 3, but using all available SEP plans to correct prior errors in the SEP profile. No angular constraints are applied to the corrections.

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

97-3819

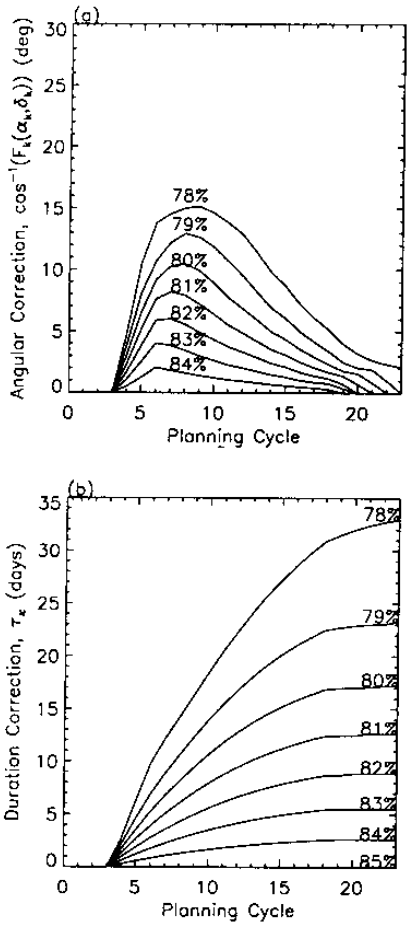

Figure 5: Same as Figure 3, but using a minimum of 3 SEP plans to correct prior errors in the SEP profile. No angular constraints are applied to the corrections.

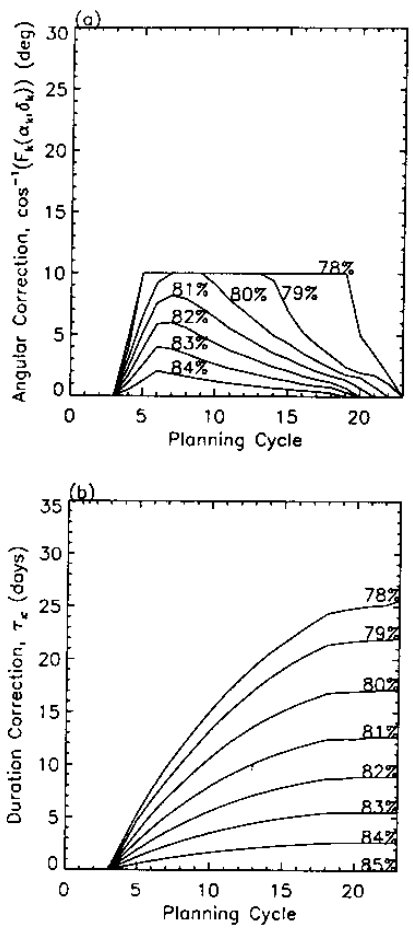

Figure 6: Same as Figure 5, but with an angular constraint of 10 degrees applied to the corrections.

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 energetically 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

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.

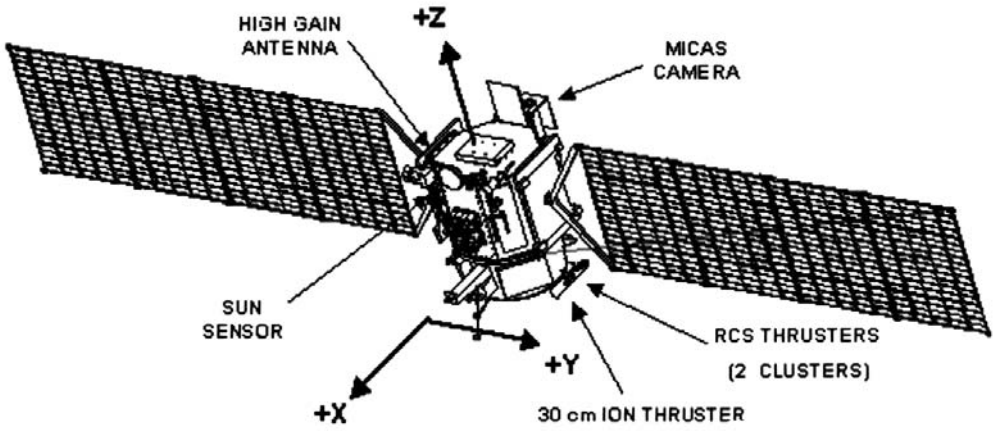

Figure 1. The DS1 Spacecraft

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 $0.5° \times 0.75°$[2], as compared to the $8° \times 8°$ 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.
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 thrustars 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 testing 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 thrustar 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 thrustar could not be obtained near the optimal thrust direction, then the thrust direction was vectorized to select several thrustars 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 thrustar 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 thrustars. 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 thrustars was based on the optimal thrust directions derived from SEPTOP. The locations of the thrustars 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 thrustars (usually one or two) at the beginning of the thrust profile or the thrust segments of interest, and hold the rest of the thrustars as fixed IPS TCMs. After the desired thrust directions were computed, THRUSTAR was used again to select new thrustars 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 plenam, and start

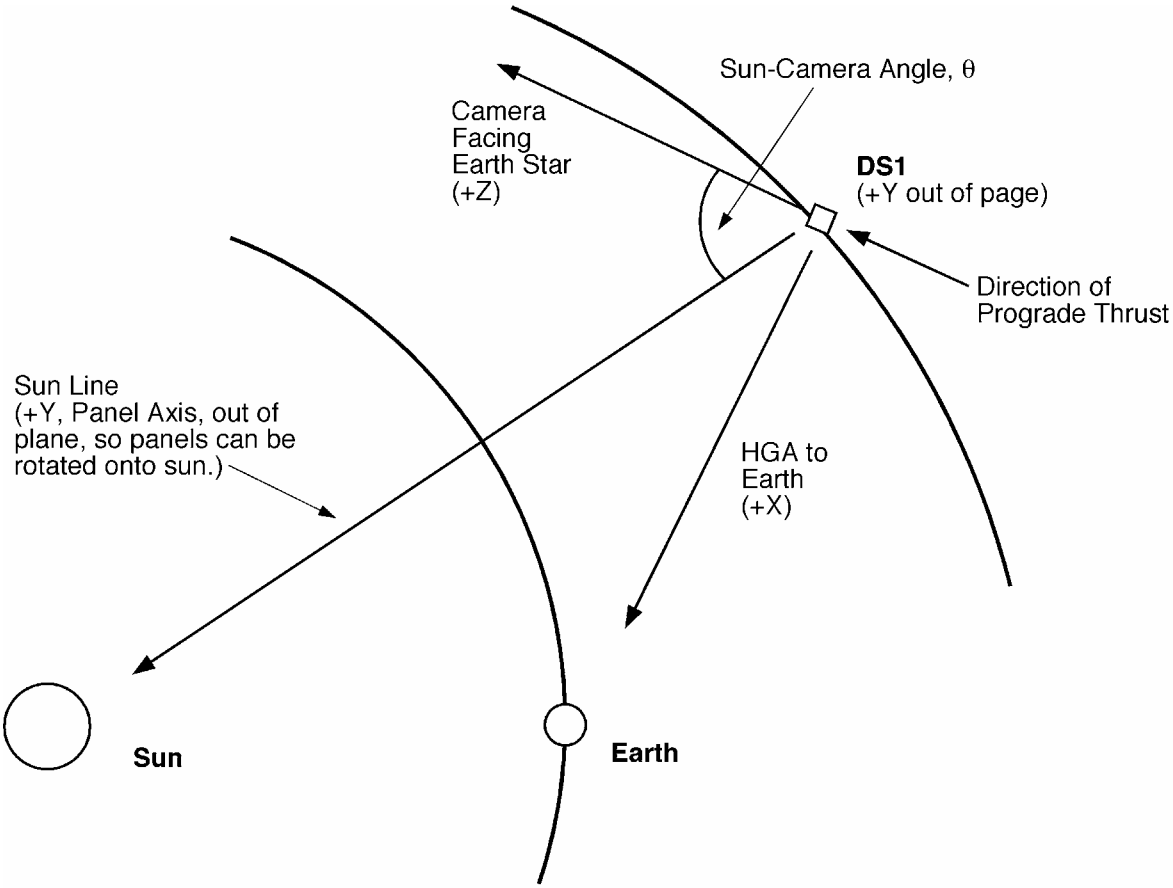

Figure 2. Earth Point Geometry

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 plenam 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 Borrely. 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 thrustar 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-ΔVs 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).

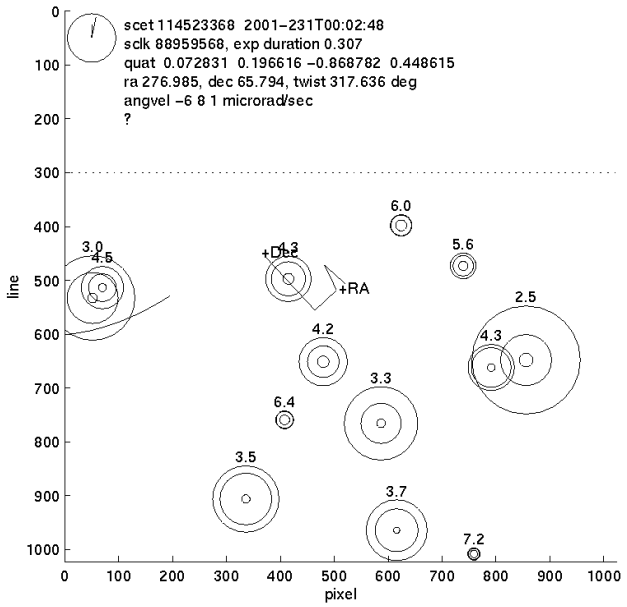

Figure 3. A recreated picture of one of the centroid data packets taken before recovery activities in LOL 5. It shows the 2.5 magnitude reference star that was locked onto. A 4.2 magnitude "companion" star is also visible, along with 11 false star signals caused by solar activity.

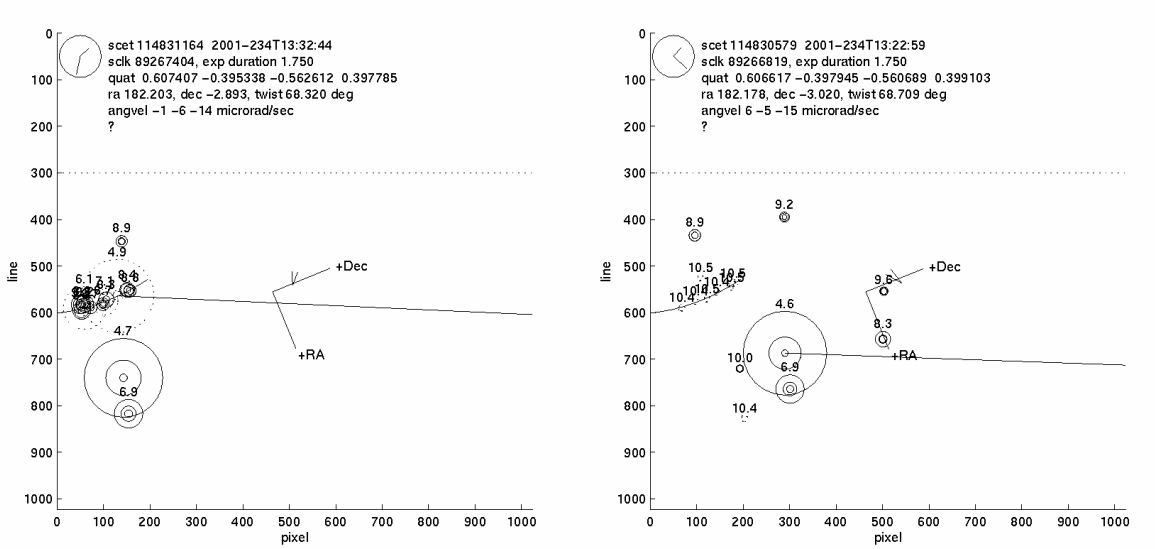

Figure 4. Centroid images taken 10 minutes apart. These images show three stars in the camera FOV, with magnitudes of approximately 4.5, 7, and 9. Other signals are stray light artifacts or cosmic rays.

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.

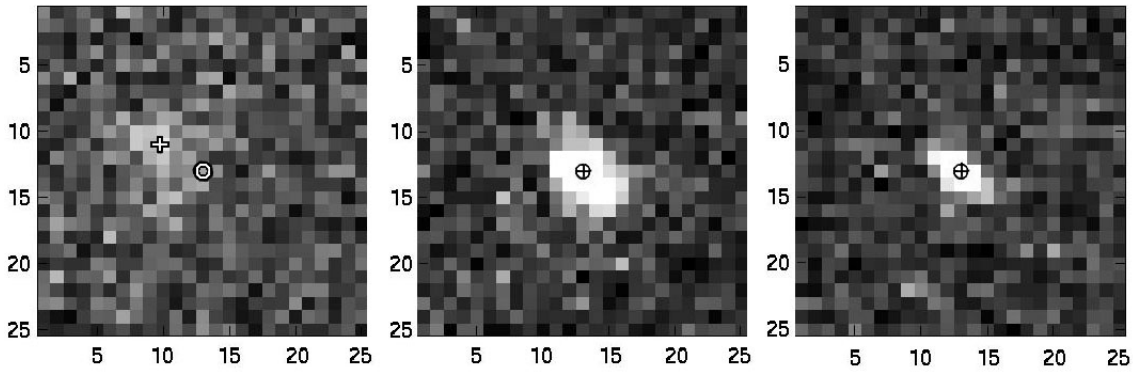

Figure 5. Left frame: Observed (+) vs. predicted (o) location of Borrelly using co-added images. Middle and right frames: registration performed on two stars seen in co-added images.

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-consumptive, and dangerous maneuver.

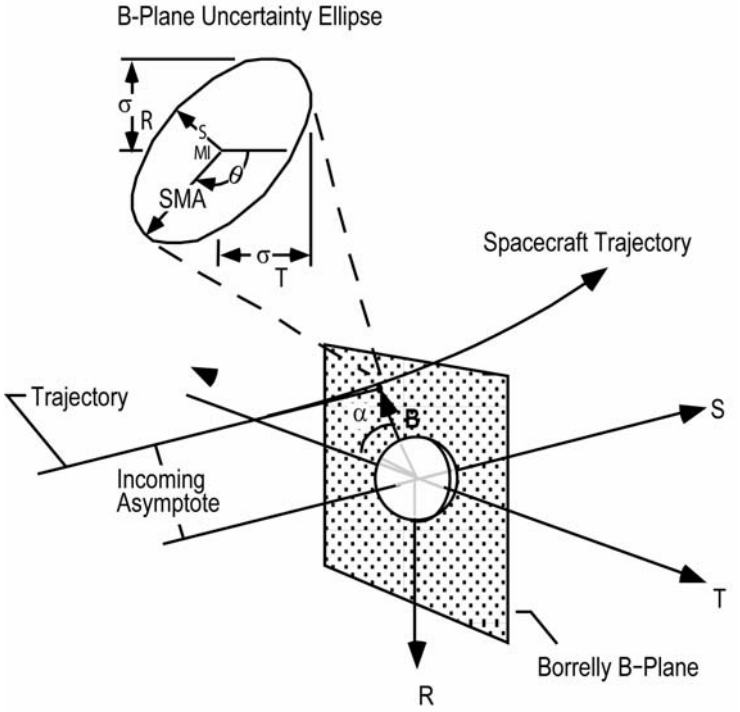

Figure 6. Targeting at JPL is performed in the so-called B-plane coordinate system.

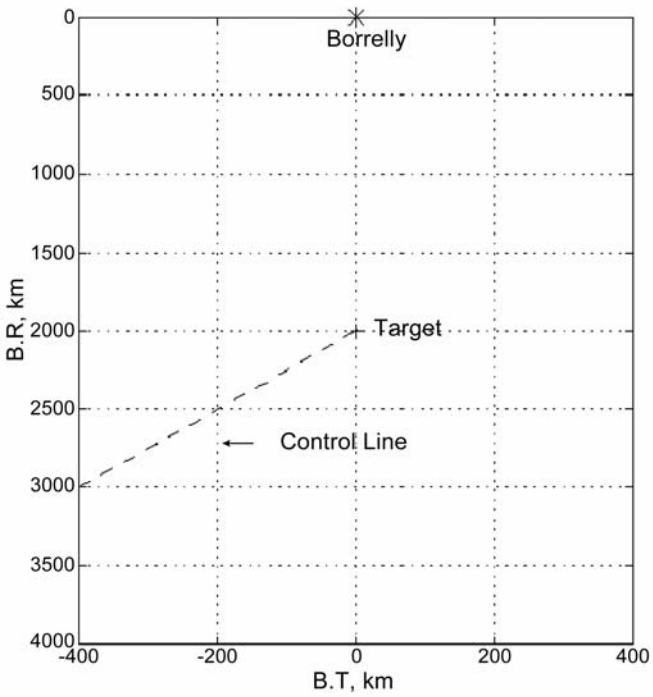

Figure 7. DS1 at Borrelly Encounter B-Plane

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.

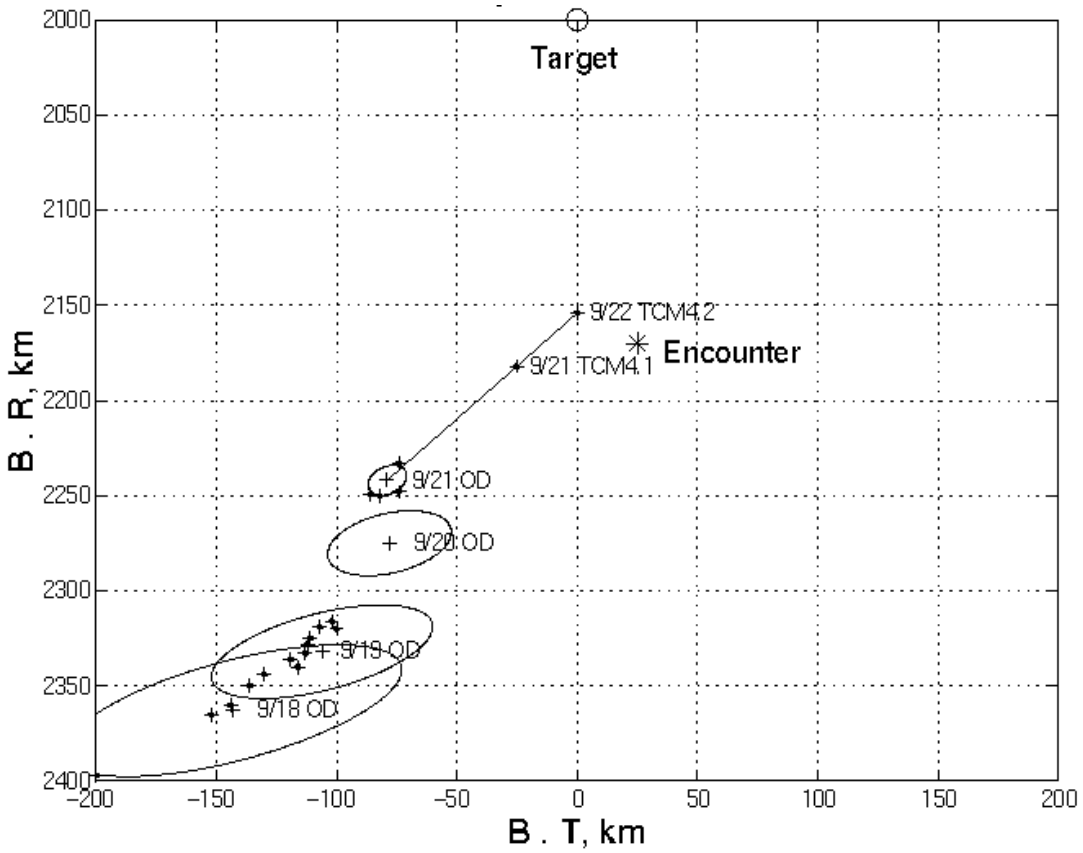

Figure 8. Final DS1 Borrelly Encounter B-Plane

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

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 ephemeredes 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 μrad 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•T and the vertical is B•R.

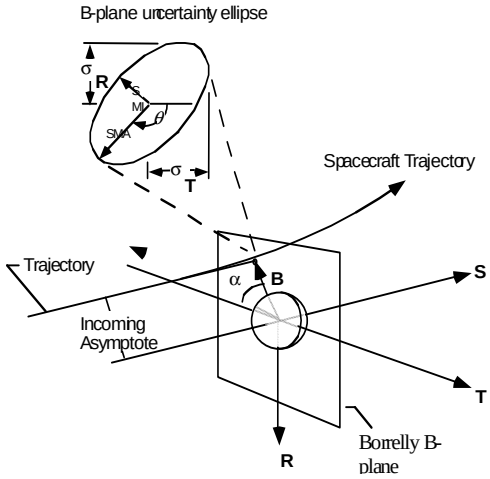

Figure 1: Borrelly B-plane

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 flyby[1]. 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 centroids[2]. 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 ephemeredes 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 nxn 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 maintainence 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

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

The maneuver strategy that was used during the approach phase turned out to be unusually complicated, partly due to the above two 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 overshot 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 SOB1were 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 DNs 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 coincidently 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 data[3].

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

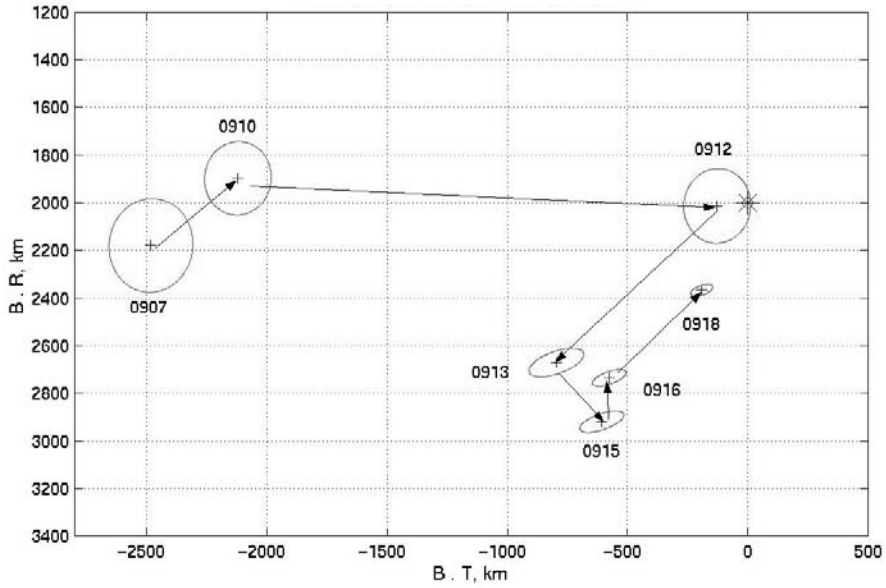

Figure 2: B-plane solutions prior to Sept. 18, 2001

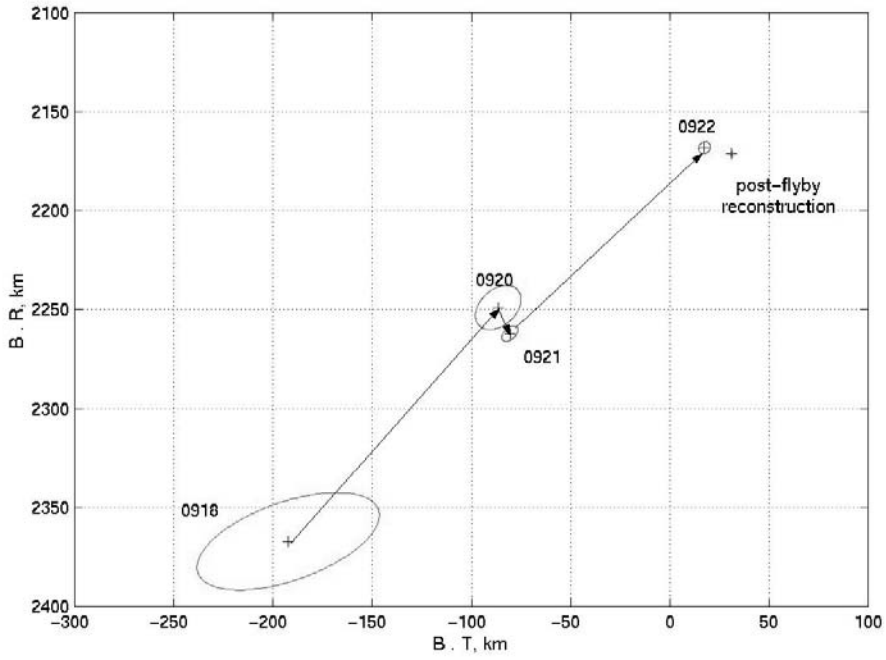

Figure 3: B-plane solutions after Sept. 18, 2001

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 elsewhere[4].

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

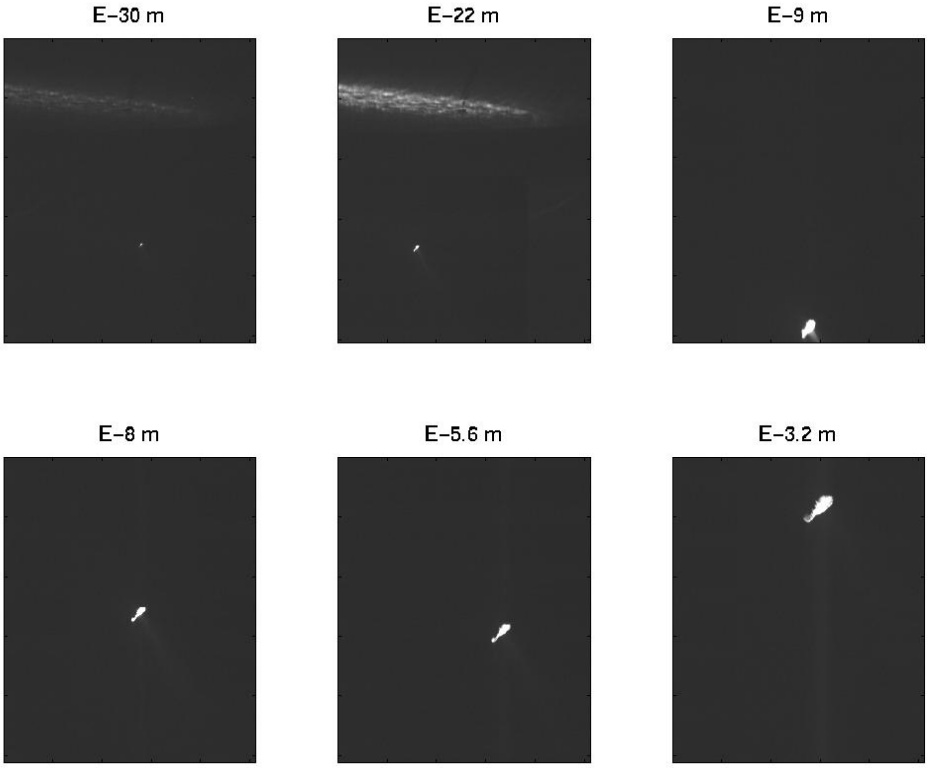

Figure 4: Sequence of RSEN images during final approach

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