AUTONOMOUS NAVIGATION FOR THE DEEP IMPACT MISSION ENCOUNTER WITH COMET TEMPEL 1

NIKOS MASTRODEMOS*, DANIEL G. KUBITSCHEK and STEPHEN P. SYNNOTT
Optical Navigation Group, Navigation and Mission Design Section, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, M/S 301-150, Pasadena, CA 91109, U.S.A.
*(*Author for correspondence; e-mail: Nickolaos.Mastrodemos@jpl.nasa.gov)*

(Received 19 August, 2004; Accepted in final form 8 December, 2004)

Abstract. The engineering goal of the Deep Impact mission is to impact comet Tempel 1 on July 4, 2005, with a 370 kg active Impactor spacecraft (s/c). The impact velocity will be just over 10 km/s and is expected to excavate a crater approximately 20 m deep and 100 m wide. The Impactor s/c will be delivered to the vicinity of Tempel 1 by the Flyby s/c, which is also the key observing platform for the event. Following Impactor release, the Flyby will change course to pass the nucleus at an altitude of 500 km and at the same time slow down in order to allow approximately 800 s of observation of the impact event, ejecta plume expansion, and crater formation. Deep Impact will use the autonomous optical navigation (AutoNav) software system to guide the Impactor s/c to intercept the nucleus of Tempel 1 at a location that is illuminated and viewable from the Flyby. The Flyby s/c uses identical software to determine its comet-relative trajectory and provide the attitude determination and control system (ADCS) with the relative position information necessary to point the High Resolution Imager (HRI) and Medium Resolution Imager (MRI) instruments at the impact site during the encounter. This paper describes the Impactor s/c autonomous targeting design and the Flyby s/c autonomous tracking design, including image processing and navigation (trajectory estimation and maneuver computation). We also discuss the analysis that led to the current design, the expected system performance as compared to the key mission requirements and the sensitivity to various s/c subsystems and Tempel 1 environmental factors.

Keywords: autonomous navigation, Tempel 1, simulations

Introduction

Mission Overview

Deep Impact is a dual spacecraft mission planned for launch in January 2005 with the engineering goals of impacting comet Tempel 1 on July 4, 2005, observing the impact event and ejecta plume expansion, and obtaining high-resolution images of the fully developed crater using the Medium Resolution Imager (MRI) and the High Resolution Imager (HRI) on the Flyby spacecraft (s/c) for the scientific purpose of exposing and understanding the interior composition of a comet nucleus.

After a 6-month cruise, the two spacecraft will separate 24 h prior to the expected time of impact (TOI). The encounter geometry will result in an illumination phase angle of approximately 64° for the Tempel 1 nucleus. The Flyby s/c will perform

Space Science Reviews (2005) 117: 95–121
DOI: 10.1007/s11214-005-3394-4

a slowing maneuver with a ΔV of approximately 102 m/s to provide 800 ± 20 s of post-impact event imaging and control the flyby minimum altitude to 500 ± 50 km. During the first 22 h following release, the Impactor s/c will acquire and telemeter science and navigation reconstruction images to the ground using the Flyby s/c as a data relay. The Flyby s/c will acquire and telemeter MRI and HRI visible and HRI infrared (IR) images of the nucleus and coma.

The fundamental requirements for the autonomous navigation (AutoNav) system are as follows:

1. Guide the Impactor to impact in an illuminated area on the nucleus surface that is viewable from the Flyby s/c for 800 s after TOI;
2. Provide comet-relative position information at TOI that will result in a Flyby HRI pointing accuracy of better than 100 μrad (3σ) to capture the impact event in the HRI 256×256 pixel subframe and observe the ejecta plume expansion dynamics;
3. Provide comet-relative position information and a crater-pointing offset that will result in a Flyby HRI pointing accuracy that is better than 1 mrad (3σ) at E+800 s to capture the fully developed crater.

The autonomous phase of the encounter will begin at 2 h before TOI. The type of autonomy used for Deep Impact can be classified as scripted autonomy (Zimpfer, 2003). A critical sequence running on-board the Impactor s/c will spawn science and navigation subsequences that issue Impactor Targeting Sensor (ITS) commands to produce navigation images at a 15-s interval. The Autonomous Navigation (AutoNav) software (Bhaskaran *et al.*, 1996, 1998; Riedel *et al.*, 2000), originally developed and demonstrated during the Deep Space 1 (DS1) mission, processes images to form observations for the purpose of trajectory determination (OD). OD updates are used to support computation of trajectory correction maneuvers (TCM) and provide relative position information for the purpose of pointing the navigation and science instruments. Three (3) Impactor targeting maneuvers (ITM) will be computed by AutoNav on the Impactor s/c and executed by the Attitude Determination and Control System (ADCS): ITM-1 at E−100 min (E−designates time before impact), ITM-2 at E−35 min, and ITM-3 at E−7.5 min. At E−4 min, the Impactor ADCS will point the ITS along the AutoNav estimated comet-relative velocity vector to capture and telemeter high-resolution (between 3 m and 20 cm) ITS images of the impact site prior to impact. Meanwhile, the AutoNav software on the Flyby s/c will process MRI images of the comet every 15 s and update the trajectory model of the Flyby s/c every minute to continuously point the MRI and HRI instruments at the nucleus. Figure 1 shows a schematic diagram of the encounter activities.

This paper describes the expected performance of the AutoNav terminal guidance on the Impactor s/c and the expected performance of the AutoNav tracking

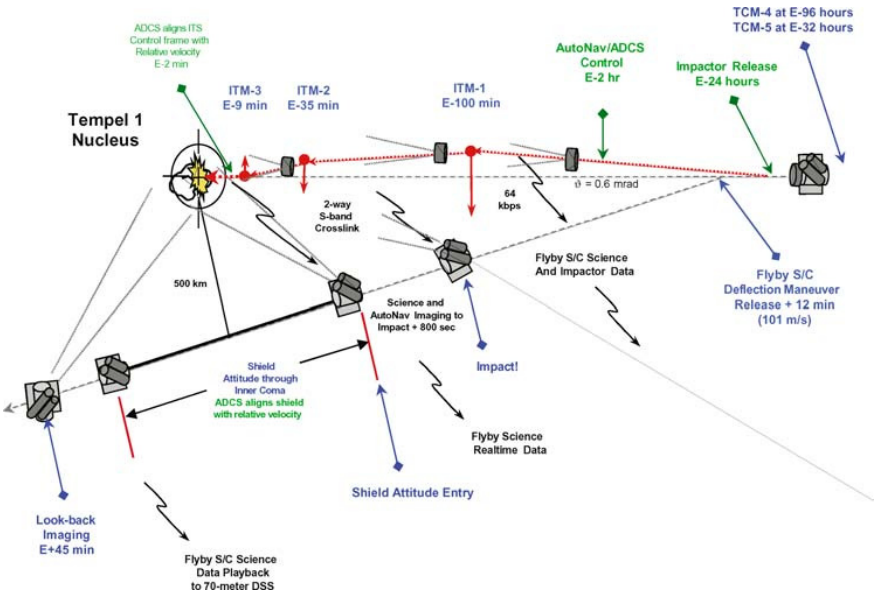

Figure 1. Tempel 1 encounter schematic for the Deep Impact mission.

process on the Flyby s/c at the two key science epochs: 1) TOI and 2) Time of highest-resolution imaging following the impact event.

FLYBY SPACECRAFT FLIGHT SYSTEM

The Flyby s/c, shown in Figure 2, was designed and built at Ball Aerospace Technologies Corporation. The Flyby s/c consists of a solar array to provide power to the combined Flyby and Impactor flight system, while the Impactor is mated to the Flyby s/c during cruise checkout and testing of the Impactor s/c; two redundant RAD 750 computers for processing; the MRI instrument that will be used for autonomous navigation during encounter; the HRI instrument with visible and IR detectors, which will be used for approach phase optical navigation (OpNav), Science imaging, and as the backup navigation camera during AutoNav operations; a high gain antenna (HGA) for real-time data return during encounter; a S-band antenna for communication with the Impactor s/c following release; a three-axis stabilized momentum wheel-based control system with a four divert/four RCS thruster hydrazine propulsion system for TCMs and momentum dumps; and an ADCS system that estimates the attitude in the ICRF inertial reference frame based on optical measurements from two StarTrackers and rates and linear accelerations from an Inertial Reference Unit (SSIRU). The mass of the Flyby s/c will be approximately 660 kg.

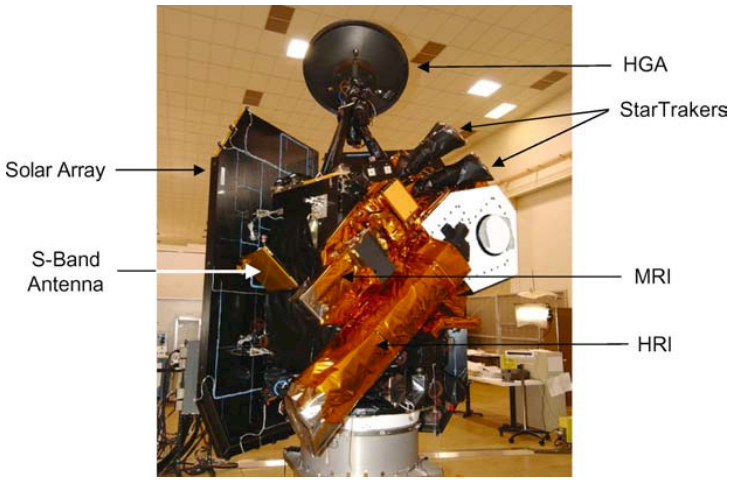

Figure 2. Flyby spacecraft in the clean room at Ball Aerospace prior to environmental testing.

The MRI camera has a 12 cm aperture (73.5 cm^2 collecting area with 35% obscuration), a focal length of 2.1 m and a 10 mrad field-of-view (FOV). The $1{,}024 \times 1{,}024$ pixel CCD is a split-frame transfer device with electronics that provide 14-bit digitization (16,384 DN full-well). The HRI instrument has a 30 cm aperture, a focal length of 10.5 m and a 2 mrad FOV with CCD electronics that are the same as the MRI.

IMPACTOR SPACECRAFT FLIGHT SYSTEM

The Impactor spacecraft, shown in Figure 3, was also designed and built at Ball Aerospace and consists of a battery for power during the 24 h free-flight, a single RAD 750 SCU for processing, an ITS consisting of a simple inverting telescope

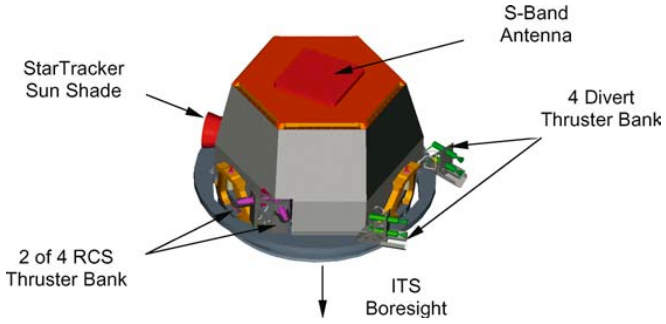

Figure 3. Impactor spacecraft flight system configuration.

and CCD detector, a S-band communications link to the Flyby s/c, a three-axis stabilized rate control system (RCS), a four divert/four RCS thruster hydrazine propulsion system with a ΔV capability of 25–30 m/s, and an ADCS subsystem that estimates the attitude in the ICRF reference frame based on measurements from a single StarTracker as well as rates and linear acceleration from an Inertial Reference Unit (SSIRU). The mass of the Impactor s/c will be approximately 370 kg with an all-copper fore-body cratering mass.

The ITS camera has a 12 cm, a focal length of 2.1 m and a 10 mrad FOV like that of the MRI. The 1,024 × 1,024 pixel CCD is also a split-frame transfer device with electronics that provides 14-bit digitization. The ITS serves a dual purpose: 1) provide navigation images and 2) provide pre-impact high-resolution (20 cm) science images.

Selection of the Impactor Targeting Strategy

General and specific guidance and control laws have been developed for several interceptor designs (Zarchan, 1997). There are two basic categories: 1) The use of proportional navigation or augmented proportional navigation techniques, and 2) the use of predictive guidance techniques.

The proportional navigation strategy is driven by measurement of the closing velocity and line of sight angular rates to control the acceleration via thrust vectoring. On the other hand, predictive guidance makes use of the dynamics (equations of motion) of both the target body and interceptor via state estimation using available measurements that can be related back to the state of the interceptor based on observations of the target. The former is what can be considered a non-dynamic or reduced-dynamic approach; the latter is a dynamic approach, which has the advantage of requiring fewer observations and being less susceptible to large, random errors in the measurements or observations.

If we consider the Impactor measurement system, which consists of ITS optical observations of the target body center of brightness; the Impactor targeting (maneuver) system characteristics; the well-known target body dynamics; the updating of the Impactor state based on the optical observations of comet Tempel 1; and the use of a few, discrete, lateral burns (ITMs) based on the predicted s/c and target body locations at the time of intercept, the Impactor targeting system for Deep Impact can be classified as an interceptor that uses predictive, pulsed guidance (few discrete burns) to achieve impact at the desired location.

It should be noted that the best quality optical observations are obtained during non-thrust periods, which suggests that a pulsed guidance system be used, as was selected for the Impactor s/c. Even though the *a-priori* position of the target body relative to the s/c can have a large uncertainty, which is reduced using optical navigation techniques, the dynamics are well known, except that (an important exception) the nucleus rotational dynamics and solar phase angle combine to induce

motion (acceleration) of the center of brightness (CB) with time, which can cause targeting errors and influence the impact location on the surface of the nucleus *via* over-estimation of the lateral velocity required to intercept at the desired location. These are mitigated, to some extent, in the batch filtering process by having some knowledge of the nucleus rotation period and by selecting the appropriate arc length over which to perform an orbit solution. This suggests that a predictive guidance strategy be used for Deep Impact and we have selected a 20 min OD arc length. In general, the OD arc length is based on a number of considerations such as rotation period, time from last OD to impact, redundancy against image loss, etc. Detailed simulations have shown that for a rotation period of 40 h or longer, impact statistics are statistically similar for OD arc lengths in the 10–35 min range but begin to increase for longer arcs.

In summary, the Deep Impact Impactor s/c uses a predictive guidance strategy and pulsed guidance system consisting of three (3) lateral, discrete magnitude burns (ITMs) based on the integrated equations of motion of the Impactor s/c and the evaluated position of the target (apriori comet Tempel 1 ephemeris) at the TOI to compute the "zero effort" miss distance, which is then used to compute the magnitude and direction of each ITM to achieve impact.

Autonomous Navigation

The autonomous navigation system for Flyby s/c tracking of the nucleus and terminal guidance of the Impactor s/c relies on both the performance and interaction of the AutoNav and ADCS flight software and the MRI and ITS navigation cameras. AutoNav consists of three (3) distinct modules: 1) Image processing; 2) Orbit determination; and 3) Maneuver computation (Impactor only). AutoNav was originally developed to operate in two different modes: Star-relative mode, which uses images that contain both the target body (beacon) and two or more stars for determining the orientation of the camera at the time of each image exposure; and Starless mode, which uses the ADCS estimated s/c attitude and camera alignment information to determine the orientation of the camera at the time of each image exposure. For Deep Impact, the Starless AutoNav mode is used based on the expected quality of the ADCS estimated attitudes. The combination of the CT-633 StarTracker(s) and SSIRU rate sensor provides an estimated attitude bias of no more that 150 μrad (3σ) on the Flyby s/c and 225 μrad (3σ) on the Impactor s/c, a bias stability of 50 μrad/h (3σ), and estimated attitude noise of 60 μrad (3σ) (Trochman, 2001).

The steps involved in the Flyby tracking and Impactor autonomous guidance process are as follows:

1. Acquire images of the comet nucleus, every 15 s, starting 2 h before the expected TOI.

2. Process the images to compute pixel/line location of the nucleus center of brightness (CB).
3. Use observed CB pixel/line locations to compute measurement residuals for comet-relative trajectory estimation.
4. Perform OD updates, every 1 min, starting 110 min before the expected TOI.
5. Perform three (3) primary ITMs at 100 min (ITM-1), 35 min (ITM-2), and 7.5 min (ITM-3) during the terminal guidance phase.
6. Acquire ITS images on the Impactor s/c for computing a Scene Analysis-based offset, relative to observed CB, just prior to ITM-3 maneuver computation and use the offset in the maneuver computation for ITM-3.
7. Acquire HRI images on the Flyby s/c at 23 min before the expected TOI for computing a Scene Analysis-based offset, relative to the observed CB, which is applied as a pointing correction to observe the impact site and track the crater during formation.
8. Align the ITS boresight with the AutoNav estimated comet-relative velocity vector starting 4 min prior to predicted TOI to capture and transmit high-resolution images of the nucleus surface to the Flyby s/c.
9. Track the impact site with the Flyby imaging instruments for 800 s following the actual TOI.

PROCESSING IMAGES

Image processing for the AutoNav system serves the purpose of providing observations of the s/c–comet relative trajectory over time. All images received by AutoNav are used to compute the target body's brightness centroid. A few selected images are also used in order to establish the most suitable impact site. AutoNav uses two different methods in order to extract the centroid of the image: 1) Brightness centroiding of all pixels above a brightness threshold and within a predetermined pixel subregion (Centroid Box); 2) Finding each separate "blob" in the image, i.e. the one or more contiguous regions of pixel brightness above a brightness threshold, computing the centroid of each blob and providing the centroid of the largest blob (Blobber). For either method, the brightness centroid of the nucleus in pixel/line coordinates is computed *via* a moment algorithm; or 3) Scene Analysis, which is used to compute the most suitable impact site and provide this information as a pixel/line offset, relative to the CB location, as determined from either 1) or 2).

The first step in image processing consists of removing the image background, most of which consists of a fixed bias value that is different for each quadrant, by subtracting a dark frame from each image. Based on the best estimate of the s/c trajectory relative to the nucleus and the navigation camera attitude at the time the image was taken, a predicted pixel/line location is computed. For processing with the Centroid Box algorithm, all pixels above the brightness threshold in a $N \times N$ pixel region surrounding this predicted location are used and the CB is

determined by computing the first moment values in the pixel and line directions. The brightness threshold can be a fixed value specified in the parameters file or dynamically computed based on a percentage of the "average peak" brightness. For Deep Impact we will use a 400 × 400 pixel region to cover uncertainties in the ADCS-estimated attitude and we will take 35% of the peak pixel brightness as the brightness cutoff. The brightness cutoff is used mainly to remove signal from the dust or coma surrounding the nucleus, so that the computed CB is entirely due to the nucleus.

The Blobber algorithm (Russ, 1999) has seen extensive in-flight use and differs from the Centroid Box in that it scans the full frame as it searches for each separate "blob" in the image. This removes the dependence on AutoNav's ability to predict where the nucleus will be located in the image array. The natural breakdown of the image into separate blobs, and subsequent selection of CB on the largest blob, offers a robust way to discriminate cosmic rays and other blemishes by means of minimum and maximum size criteria. This approach has many advantages over the Centroid Box approach, but has proven to be less stable for certain nucleus orientations where the illumination is such that the nucleus appears to have the shape of two disjointed large blobs of comparable size; in this case, a small change in the relative size of the two blobs can shift the CB abruptly and result in large targeting errors.

Scene Analysis images are processed using either Centroid Box or the Blobber algorithm to provide a reference CB location. In the first step, the image is scanned from the edge of the image towards the center to delineate the outer boundary of the nucleus by comparing each pixel with the brightness cutoff. Each on-nucleus pixel is further characterized as interior dark or bright by a similar comparison to the cutoff value. In the next step, make each on-nucleus pixel the center of a circle of radius equal to the 3σ error footprint, or control error. The control error itself is established *via* Monte Carlo simulations. For each candidate pixel, we compute a number of quantities: total number of on-nucleus bright pixels; total number of off-nucleus pixels and proximity to the direction of the Flyby s/c point of closest approach to the nucleus. In order for a pixel to qualify as a candidate for the impact site, it has to contain a certain portion of on-nucleus pixels within its control error. A typical such number is 95%. Qualifying pixels are then compared against each other with criteria considered in order of decreasing priority. The first criterion is maximizing the number of on-nucleus bright pixels within the control error (maximizing the lit area). When more than two sites meet the first criterion, the second criterion is minimizing the number of off-nucleus pixels. The third priority is proximity to the Flyby s/c's point of closest approach, which we call "biasing". Everything else being equal, biasing tends to move the impact site toward a location on the nucleus that is more likely to be visible from the Flyby s/c near its closest approach point where the view angle increases to 45°, thus enhancing the ability of the Flyby s/c to image the fully developed crater at the time of highest-resolution imaging. The difference between the CB reference location

and the selected site is computed and converted to an inertial correction vector that is stored and used by the AutoNav maneuver computation software. If no suitable site is found, then no correction vector is returned, which results in a default to targeting the CB. Note that the period of rotation does enter into the computations of the best impact site and there is no modeling of the pole or the orientation of the nucleus at encounter, since this information is not assumed to be available at any time prior to encounter.

TRAJECTORY DETERMINATION

For the Flyby and Impactor s/c, the trajectory is estimated and updated every minute during the last 2 h of encounter. The trajectory determination software supports both ADCS attitude control and AutoNav maneuver computations.

After images are processed, the important information needed to relate the observations back to the state of the spacecraft are stored in an on-board optical navigation (OpNav) file. This information consists of the time the image was exposed, the camera inertial orientation (right ascension, declination, and twist), the pixel and line location of the observed CB, and the data weight associated with a given observation and whether or not the observation was declared useful. When AutoNav receives the command to perform orbit determination, the best estimate of the s/c position and velocity is read from a second on-board file called the orbit determination file. The trajectory is integrated, making use of on-board accelerometer data that is stored in a non-gravitational history file, to the time of each observation and the predicted pixel and line location of the nucleus center of mass is computed. The difference between the computed pixel/line location and the observed pixel/line location corresponding to the observations contained in the OpNav file represents the residual, which is minimized in a least-square sense using AutoNav's batch-sequential processing algorithm to estimate position, velocity and two navigation instrument cross-line-of-sight attitude bias drift parameters. The on-board OD file is then updated with the best estimate of the s/c state vector. The OD arc length was selected to be 20 min for Deep Impact and with image processing every 15 s, each OD arc contains 80 observations.

Following completion of each orbit determination, the trajectory is updated and a representation of the estimated trajectory is generated in the form of Chebyshev polynomial coefficients. This polynomial represents a time series of the predicted Impactor position relative to the nucleus that is passed to the ADCS flight software. ADCS evaluates the polynomial, computes the relative position vector and aligns the ITS camera boresight on the Impactor and the HRI and MRI camera boresights on the Flyby s/c with the comet-relative position vector to center the nucleus in the instrument FOV. Figure 4 shows simulated MRI AutoNav images at E−2 h, E−0 min (Impact), and E+750 s, and the HRI AutoNav Scene Analysis image at E−23 min.

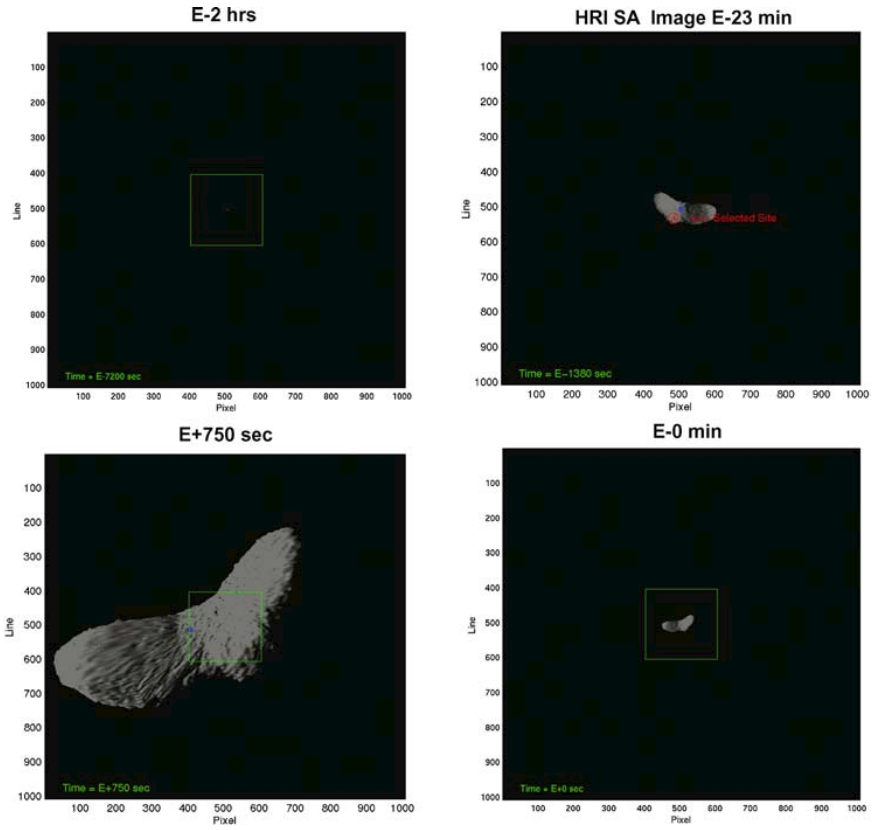

Figure 4. Flyby MRI AutoNav images with HRI FOV superimposed at E−2 h, E−0 min, and E+750 s. The upper right image is the Flyby HRI AutoNav Scene Analysis image at E−23 min.

IMPACTOR TARGETING MANEUVERS

Impactor targeting maneuvers are initiated *via* an AutoNav sequence command. These commands are issued from the engineering sequence that will be running during encounter. When AutoNav receives the command, an impulsive maneuver (magnitude and direction) relative to the ICRF frame is computed for the time contained in the command packet and passed to ADCS in the form of a command issued by AutoNav. ADCS receives the maneuver ΔV information in the command packet, computes the finite duration burn start time, and populates the necessary flight software current value table. When the current value table is populated, ADCS issues a command to spawn a trajectory correction maneuver (TCM) transition sequence consisting of the following commands:

1. Disable AutoNav image processing.
2. Command ADCS to transition (turn) to burn attitude.
3. Reset accelerometer accumulated ΔV.
4. Command ADCS to delta-V mode in preparation for ITM execution.

When the burn start time is reached, the burn is initiated. The ADCS flight software continually monitors the accumulated ΔV information based on incremental ΔV values measured by the SSIRU accelerometers and continuously adjusts the pointing and thruster duty cycle until the accumulated ΔV matches the desired maneuver ΔV, which results in burn termination. When the burn is complete, ADCS returns to instrument point mode and issues another command to spawn a second TCM transition sequence to re-enable AutoNav image processing. The need for TCM transition sequences and for disabling AutoNav image processing arises from the non-deterministic burn duration of each ITM. This allows instrument commands for AutoNav images to be issued every 15 s without regard to when a particular maneuver will occur or how long the burn will last. Since the quality of the images acquired during the burn may be degraded, this also allows AutoNav to simply ignore the navigation images received during these periods.

ITM-1 (ENCOUNTER−100 MIN)

The terminal guidance phase of the mission begins 120 min before impact. ITM-1 is the first targeting maneuver based on ITS observations and an updated trajectory. The primary purpose of ITM-1 is to remove the Flyby pre-release delivery errors. ITS observations have to be good enough to provide improvement over the pre-release orbit determination, but waiting too long for ITM-1 increases the amount of ΔV required to remove the delivery errors that are expected to be 6 km (3σ) due to the quality of the Flyby approach phase optical navigation solution. The delivery requirement is 30 km (3σ). Correcting for a 6 km pre-release delivery error requires $\sim$1 m/s ΔV for a maneuver placed at E−100 min. The maneuver execution error will be on the order of 4 cm/s for a 1 m/s burn. ITM-1 targets the nucleus CB (dot in Figure 5). Figure 5 shows a simulated image of the nucleus at a range of approximately 60,000 km. The optical signal from the nucleus spans approximately 10 pixels at the time of ITM-1.

ITM-2 (ENCOUNTER-35 MIN)

The second ITM, ITM-2, was placed at E-35 min to provide redundancy in the form of improved targeting over ITM-1 in the event ITM-3 fails to execute. ITM-2 will require approximately 11 cm/s of ΔV to remove the ITM-1 maneuver execution errors. The trajectory solution is based on ever improving ITS observations of the nucleus. As seen in Figure 6, the nucleus spans more than 25 pixels. ITM-2 also targets the nucleus CB (dot at tip of arrow).

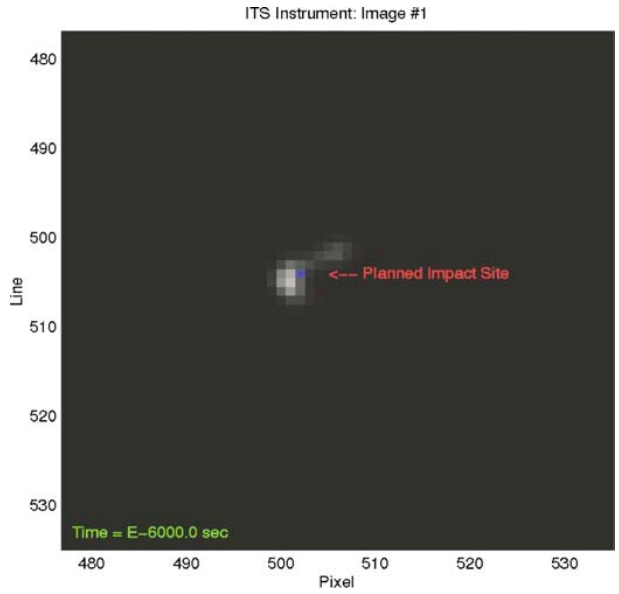

Figure 5. Simulated image of the comet nucleus with a 65° illumination phase angle at the time of ITM-1 (E−100 min).

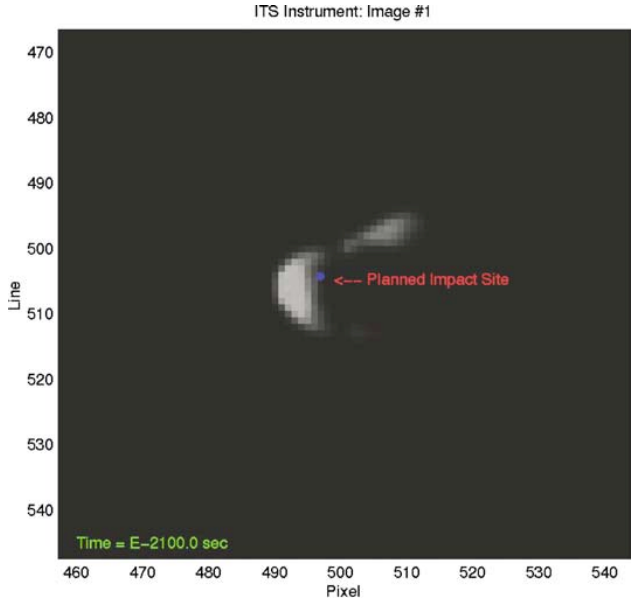

Figure 6. Simulated image of the comet nucleus with a 65° illumination phase angle at the time of ITM-2 (E−35 min).

ITM-3 (ENCOUNTER−7.5 MIN)

The third and final ITM, ITM-3, is the most important and provides the fine targeting for an illuminated impact. The trajectory determination for ITM-3 is based on CB observations of the nucleus. However, at E−11 min a Scene Analysis image is sequenced to provide a targeting offset from the CB. This offset, as previously mentioned, provides an increased probability of an illuminated impact and increases the probability that the crater will be seen from the Flyby s/c at the time of highest-resolution imaging. As seen in Figure 7, the nucleus spans nearly 100 pixels at E-11 min. The blue dot represents the computed CB. The red "+" symbol and surrounding circle represent the Scene Analysis selected impact site. It is important to note that for this nucleus orientation (particularly challenging), the CB lies in a shadowed region on the nucleus surface and would likely result in a dark impact had that been the target.

ITM-3 will nominally be centered at 7.5 min before impact and could require a B-plane correction of as much as 4 km to take the Impactor from intercept at the CB to intercept at the selected site. A 4 km correction at E-7.5 min requires approximately 7 m/s of ΔV. The associated maneuver execution error is expected to be no more than 10 cm/s, which results in a B-plane error of only 54 m. Placing

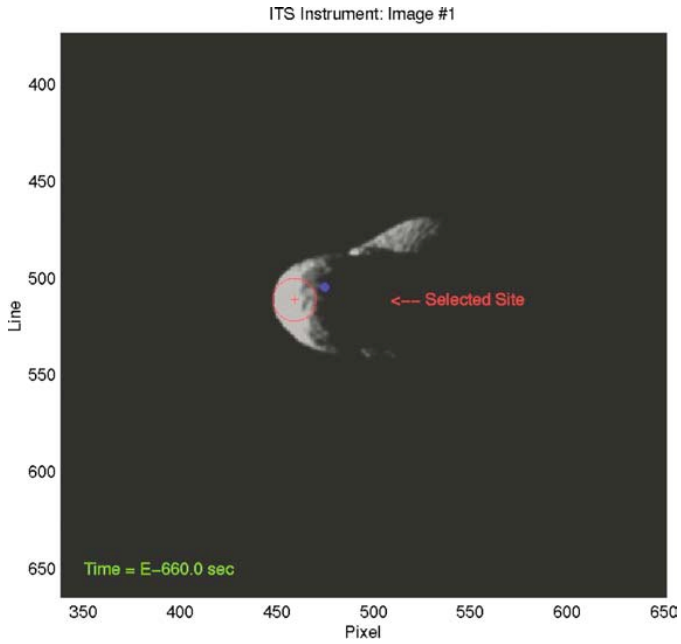

Figure 7. Simulated image of the comet nucleus with a 65° illumination phase angle at the time of scene analysis for ITM-3 (E−11 min).

ITM-3 as late as possible is ideal for minimizing maneuver computation errors and maneuver execution errors, however, the Impactor s/c is constrained by the amount of propellant (25 m/s ΔV capability) available for ITMs and the worst-case coma environment could result in ADCS control upsets due to particle impacts. To ensure a high-probability of successful maneuver completion it was determined that ITM-3 be completed prior to E−6 min, which represents a range of approximately 3,600 km from the nucleus.

Expected Autonav Performance

The performance of AutoNav is examined by simulations and testing at increasing levels of fidelity:

a) Monte Carlo simulations, which replicate the basic image processing, orbit determination and maneuver computation functions including error sources such as those from ADCS attitude estimation and maneuver execution. These simulations form the bulk of the analysis and are used not only for insight into the sensitivity of the various parameters, but also to establish important encounter parameters.
b) Monte Carlo simulations, which trigger the actual flight code and are used to confirm that the flight software will meet its expected performance.
c) Software test bench runs that utilize the closed-loop interaction between AutoNav and ADCS in the flight system environment which are used to verify the encounter sequence design and the AutoNav/ADCS interface. All three levels of testing are essential for performance validation at encounter.

This paper discusses the results from the Monte Carlo simulations from category a) described earlier.

Nucleus Test-Model Development

One of the unique technical challenges of the Deep Impact mission is the lack of detailed knowledge of the Tempel 1 nucleus prior to the actual encounter and the need for robust performance within a broad range of cometary characteristics such as different shape models, topographic relief maps, average size, rotation period, pole direction and dust environment. The most recent data from ground- and space-based observations, presented in detail elsewhere in this issue, indicate that the Tempel 1 nucleus has an average radius just over 3 km, a mean albedo in the range 0.03–0.04, an axial ratio of ∼3:1, a rotation period in the 38–41 h and a shape that may be irregular and not ellipsoidal.

Over the past few years, the Optical Navigation Group at JPL, in collaboration with the Deep Impact Science Team, has constructed a number of shape models driven by past observations of Solar System small bodies or theoretical predictions

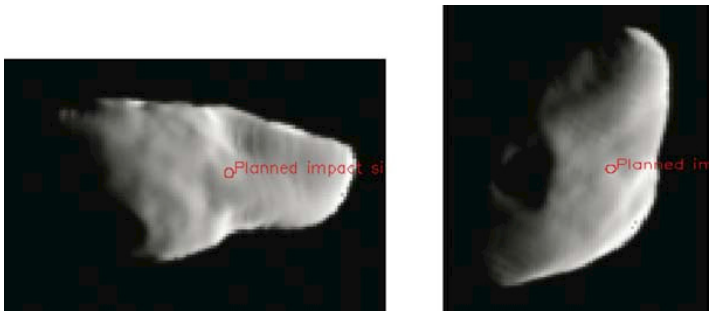

Figure 8. Examples of a Halley-like nucleus with an ellipsoidal shape and few topographic variations.

that span a range of possibilities that bound what is currently predicted for comet Tempel 1. Three such categories of models are highlighted as follows:

1. Halley-like models (Figure 8) derived from the unpublished shape models developed by P. Stooke and A. Abergel (Halley Nucleus Shape Model, personal communication via M. Belton, 2000). These models have a general ellipsoidal shape with an axial ratio of 2:1 and a relatively smooth surface;
2. Borrelly-like models (Figure 9) derived from the DS1 observations of comet Borrelly (Kirk *et al.*, 2004). They have axial ratios in the range of 2.5:1 (for the baseline model) to 3.5:1 (for the worst-case model) and a double-lobed structure that results in significant shadowing. These models cause the largest targeting errors and their center of brightness can be very near the limb; and
3. Accretion models (Figure 10) that result from theoretical considerations (Weidenschilling, 1997). These models have extreme topographic variations on every length-scale with heights up to half the size of the horizontal length-scale of the surface fluctuations. They tend to result in a higher frequency of dark impacts.

KEY AUTONAV ASSUMPTIONS AND ERROR SOURCES

The Monte Carlo analysis described in this paper relies on a number of key assumptions and sources of error that directly affect the performance of the AutoNav targeting and tracking algorithms: 1) Cometary characteristics; 2) ADCS attitude estimation errors; and 3) Maneuver execution errors.

Cometary Characteristics

In addition to the unknown shape, axial ratio and topography of Tempel 1, which is addressed by the variety of nucleus models discussed earlier (Figures 8–10), other comet model parameters that affect navigation performance include: the coma brightness, average size, rotation period, pole direction, nucleus orientation at the time of the encounter and the time-of-flight, or downtrack position of the s/c relative to the nucleus.

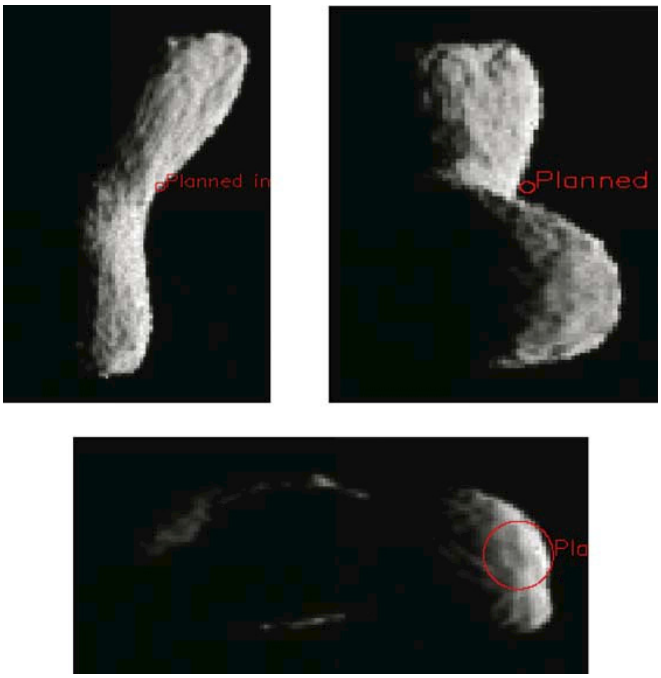

Figure 9. Typical examples of Borrelly-like nucleus models with a 3.5:1 axial ratio. The small circle indicates the center of brightness in the upper left and upper right images and the large circle is the site selected using scene analysis in the bottom image.

The coma brightness is one of the key considerations for a successful encounter. The need to discriminate off-nucleus coma, so as to guide the Impactor towards the nucleus, and on-nucleus foreground coma, so as to select an intrinsically illuminated impact site is paramount. The coma brightness is specified in terms of the brightness ratio between the peak brightness of a jet and the peak brightness of the nucleus when both are fully resolved and is bounded between the most-likely case, which is the current best expectation and the worst-likely case, which represents scenarios whereby all observed dust is created by a single narrow jet. The brightness ratio for the most-likely case coma model is 1/32 and for the worst-likely case coma model it is 1/6. AutoNav computes an autonomous coma cutoff, relative to an "average peak" brightness value per image, which for conservatism is currently set to 35% of the peak or $\sim$2 times larger than the worst-case coma. As the coma cutoff value increases, a larger portion of the nucleus image is discarded as being part of the coma. This has the undesirable side effect of moving the CB closer to the bright limb as well as causing a less-stable CB from image to image, which maps directly into an OD error. Figure 11 illustrates this point for two different values of the coma

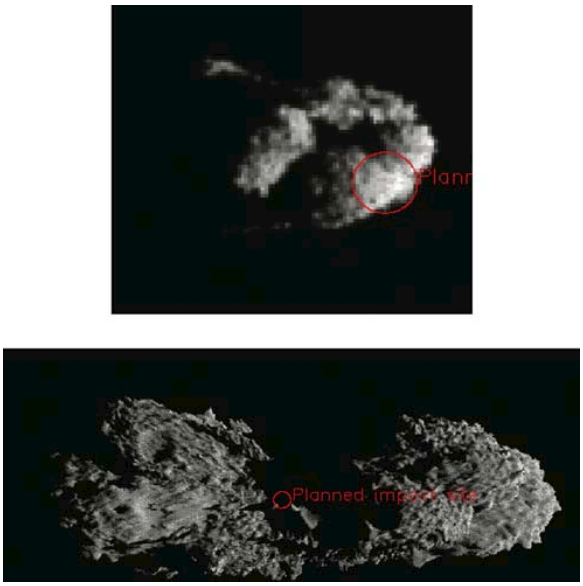

Figure 10. Examples of accretion models at different levels of spatial resolution. The top image is taken at E−11 min and shows the site selected using scene analysis; the bottom image was taken at E−3 min and shows the CB, which is not illuminated.

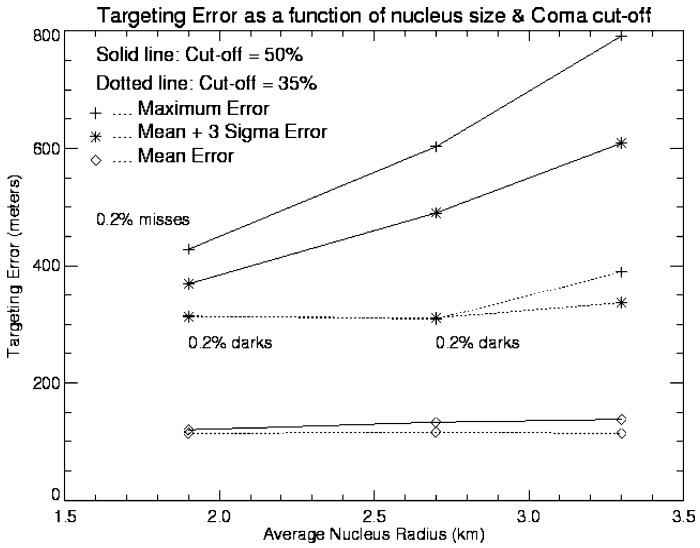

Figure 11. Targeting error as a function of mean radius and coma brightness cutoff for a Halley-like nucleus. While the mean targeting errors of 100–120 m are not sensitive to the size of the body, the maximum and 3σ targeting errors are significantly reduced with a lower coma cutoff.

cutoff. The larger targeting error for a 50% brightness cutoff results in a larger miss distance and a higher probability of a miss, 0.2% in this example, compared with the 35% brightness cutoff, where the probability of a miss is traded with a probability for a dark impact (0.2%). The final trade between completely eliminating all influence of coma without discarding too much of the nucleus, will be made in-flight after the dust environment of Tempel 1 has been analyzed using HRI image data during the approach phase of the mission (Encounter-60 to Encounter-10 days).

The rotation of the nucleus is very important and affects the magnitude of the targeting error by moving the CB about the center of mass. In the presence of significant shadowing, the apparent motion of CB can be quite large, giving rise to apparent cross-track velocities up to a few meters per second, which result in an incorrect orbit solution relative to the true center of mass of the nucleus. This effect becomes very pronounced for very short rotation periods, with an attendant increase in the mean targeting error.

The OD errors described earlier are also proportional to the nucleus size *via* a lever-arm effect: any apparent CB motion in the plane-of-sky is amplified by a larger nucleus, which in-turn proportionally increases the OD error. In this regard, we have the counter-intuitive result that a larger nucleus, although offering a larger cross-sectional area for targeting, causes a targeting error proportional to its size. The end result is that a larger nucleus does not necessarily result in better targeting with the Impactor and degrades pointing for the Flyby s/c. In Figure 12 we show

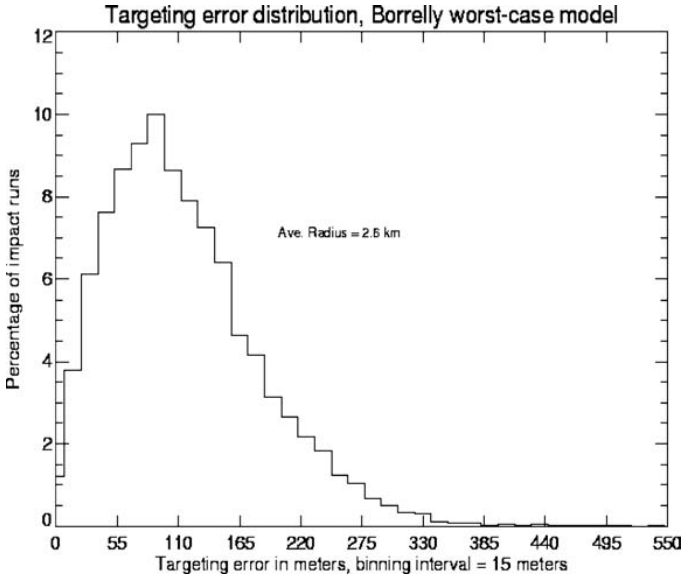

Figure 12. Distribution of targeting errors, which represent a typical example of the skewed distribution seen for all nucleus models examined. Impactor misses occur at the tail-end of the distribution.

a histogram distribution of the absolute value of the targeting errors for a 2.6 km mean nucleus radius. In our parametric studies, we have assumed a range of mean radii between 1.8 and 4 km, which spans the currently predicted mean radius of Tempel 1.

The inertial direction of the rotation pole and the orientation of the nucleus at encounter are treated as unknown in the AutoNav targeting and tracking process. In order to have a realistic picture of how different body orientations affect targeting and tracking, we perform large-scale simulations where the pole spans the celestial sphere at 124 different directions. For each pole direction the nucleus is placed at 72 different orientations, relative to the s/c, each 5° apart, before it begins to rotate for a given simulation run.

Time-of-flight errors are particularly large for cometary encounters compared to other small bodies. Owing to non-gravitational accelerations from outgassing that vary from orbit to orbit, and systematic errors in astrometric observations due to the displaced center of light from the center of mass, long-term predictions of cometary orbits are difficult. As a result, the actual position of the comet along the line-of-sight of the navigation instruments will not be reliably known until the Flyby s/c determines its TOF *via* direct geometric observations during the last 1 h of the encounter. That information will not be available to the Impactor for navigation, since observations from the ITS camera have no strength in establishing TOF due to having zero parallax. TOF errors map into targeting errors, which are to first-order proportional to the magnitude of a given maneuver. The nominal, 1σ, TOF error at E-24 h is expected to be 9 s. However, past experience has shown that TOF errors up to 30 s can be experienced, which for an ITM-3 of $\Delta V = 13.5$ m/s (the 99% figure) will contribute to ~400 m of targeting error.

Finally, direct parallax measurements of the Flyby s/c during nucleus tracking will allow AutoNav to determine its position relative to the nucleus center of mass only to within a few kilometer, because the actual location of the impact site relative to the center of mass of the nucleus is unknown and cannot be established by the observations. This effect introduces pointing errors when viewing the crater at large angles of 30° or more, relative to the approach asymptote. This cannot be compensated for, unless the shape, size and orientation of the body are known prior to the encounter. This type of pointing error increases with the size of the body.

The error sources given earlier, associated with the observations of an unknown nucleus that are used to establish a stable orbit solution, constitute by far the largest contributors to the targeting error for the Impactor and pointing error for the Flyby s/c.

ADCS Attitude Knowledge Errors

The ADCS subsystem, which estimates and controls the spacecraft's attitude, provides AutoNav with the attitude information necessary to determine the inertial pointing of the navigation camera's boresight. Errors in estimation of the

spacecraft's attitude map directly into errors in the estimated position and velocity when the s/c state is determined from optical observations from the navigation camera.

The attitude estimation errors are broken down into high-frequency noise, which amounts to 60 μrad (3σ) about any axis, a fixed attitude bias of 150 μrad (3σ) about any axis for the Flyby s/c, a fixed attitude bias of 220 μrad (3σ) about any axis for the Impactor s/c, and bias stability of 50 μrad/h (3σ) on both s/c. These errors scale with range, frequency of observations and span of time over which the observations are used to determine the orbit solution (OD arc length). With the current baseline of a 20 min arc length, one image every 15 s and the last picture prior to ITM-3 at E-11 min, the 3σ ADCS-induced targeting error for the Impactor is $\sim$240 m. The magnitude of this error is quite modest and analysis has shown that AutoNav is tolerant to ADCS attitude noise that is six times larger and a bias drift that is $\sim$3 times larger than given earlier before noticeable performance degradation begins to occur. The estimation of two additional cross-line-of-sight attitude bias drift parameters in the initial state allows AutoNav to tolerate much larger bias drift values up to 1 mrad/h without a performance penalty.

Impactor Targeting Maneuver Execution Errors

Maneuver errors are the largest errors expected to influence targeting on the Impactor s/c. Their behavior is such that the largest component is proportional to the magnitude of the requested ΔV. For conservatism, our assumptions on maneuver execution errors exceed the expected performance figures by 90–35% for maneuvers in the range of 10–25 m/s, respectively. In Figure 13, the targeting error is shown as a function of the maneuver magnitude, ΔV. For a ΔV of 13.5 m/s, (the 99% tile), the 1σ maneuver error contributes $\sim$200 m to the overall targeting error.

IMPACTOR SPACECRAFT TARGETING RESULTS

In Table I we summarize the overall expected Impactor performance based on the assumed ADCS and maneuver error models described in the previous section, and for the range of different nucleus shape models, nucleus sizes, ITS point-spread function (PSF) ranging between 0.8 and 1.7 pixels, OD arc lengths in the 10–40 min range and data weights in the 2–15 pixel range. For comparison, in Table II, we also show one representative set of results for targeting the CB of the nucleus as opposed to targeting the best site based on Scene Analysis.

The range in statistics for the Borrelly–worst-case model is the overall range of results found under the previously described parameter variations. The few misses observed are the result of unfavorable nucleus orientations at encounter that give rise to large OD errors. It is also evident that the simplest approach, CB targeting, is not adequate in meeting the mission's goals, but only for a small range of possible shape and topography models such as those derived from Halley.

TABLE I
Scene analysis targeting results.

Nucleus model	Probability of illuminated impact (%)	Dark impact probability (%)	Probability of a miss (%)	Approximate number of simulations
Halley	99.98	0.02	0	9,000
Borrelly–worst case	99.98–99.88	0.01–0.1	0–0.05	100,000
Accretion	97.28	2.7	0.02	9,000

TABLE II
Impactor CB targeting results.

Nucleus model	Probability of illuminated impact (%)	Dark impact probability (%)	Probability of a miss (%)	Approximate number of simulations
Halley	99.8	0.2	0	1,500
Borrelly–worst case	90.75	6.75	2.75	9,000
Accretion	65	35	0.02	700

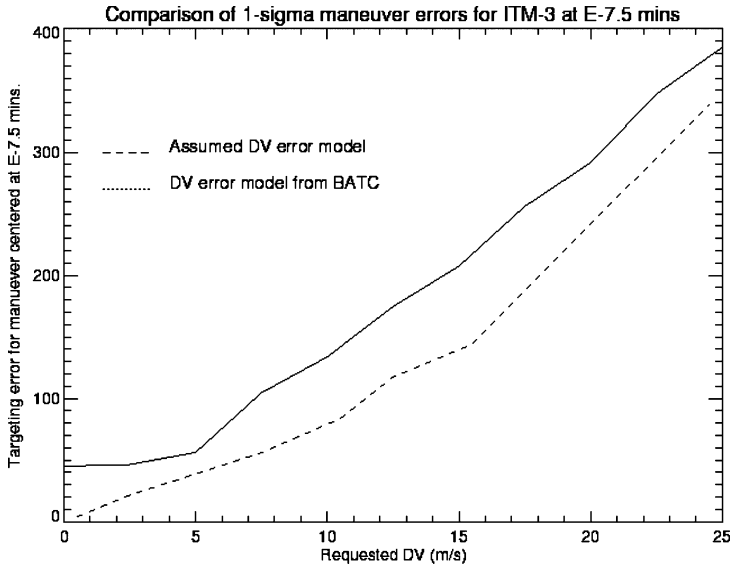

Figure 13. The 1σ Impactor targeting error that results from maneuver execution errors over the range of expected ΔV values.

KEY FLYBY RESULTS

One could break down the Flyby spacecraft's encounter goals into two main categories: 1) establish an accurate trajectory relative to the center of the nucleus; and 2) point the instruments at the impact site and continue tracking the crater during formation and for as long as possible to the end of the post-impact Science observations.

Compared to the Impactor, the Flyby s/c is able to determine its position relative to the nucleus in all three dimensions, including the downtrack direction or TOF, by parallax observations (i.e. measuring the nucleus CB at different view-angles over time). The strength of those observations is proportional to the ratio of the Flyby s/c miss distance to the downtrack distance or range to the point of closest approach. Significant TOF corrections occur during the last 40 min of observations. Determining the TOF is important for the temporal optimization of Science imaging to capture the impact event with the Flyby s/c instrumentation and for optimizing the pre-impact imaging on the Impactor s/c. The TOF is transmitted from the Flyby to the Impactor *via* the S-band cross link at ∼E−5 min and E−3 min. As seen in Figure 14, the initial TOF error converges rapidly towards zero as the s/c closes the range to the comet. Because of the direct observations, the reduction in the TOF error is practically insensitive to the initial conditions as long as the apriori covariance of the comet is sufficiently large to accommodate the magnitude of the necessary adjustments. The TOF requirement allocated to AutoNav is 2.7 s between E-5 min and E-3 min. This requirement is met at the 98% confidence level when the

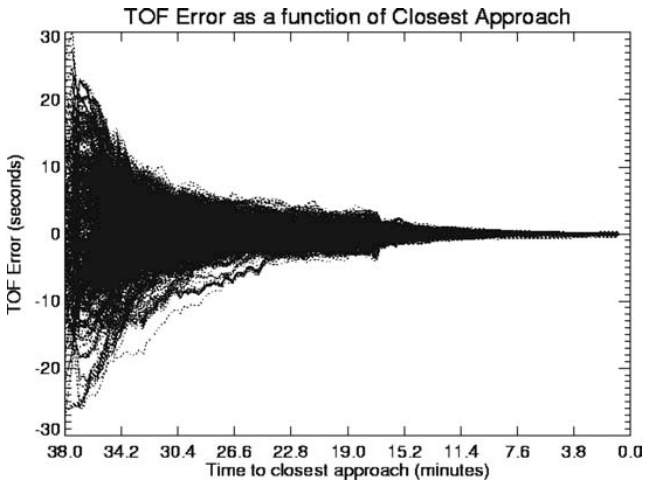

Figure 14. The decrease in the TOF error as a function of time to closest approach of the Flyby s/c from ∼4,500 Monte Carlo simulation runs. The initial conditions simulated a 10 s (1σ) TOF error.

initial TOF error at E-2 h is assumed to be 10 s (1σ) and to the 97% confidence level when the initial TOF error is 30 s (1σ). At the TOI the largest TOF error is 1.6 s. At E+12 min and later the mean TOF error is <0.15 s, or equivalently, the error in the location of the center of mass of the nucleus is <1.5 km in the downtrack direction.

Pointing the instruments at the expected impact site prior to the impact event and to capture the actual crater location following the impact event, requires both knowledge of the nucleus CB and of the impact site relative to the CB. Currently, the impact site chosen *via* Scene Analysis on the Impactor s/c is not communicated to the Flyby. Instead, the Flyby estimates the best impact site independently and by the same process used on the Impactor. This is done on the Flyby s/c using the HRI at a time necessary to match the resolution element of the ITS on the Impactor s/c when the Scene Analysis image is processed, where the figure of merit for the resolution is the full width half maximum (FWHM) of each instrument. Currently, with a 1.5 pixels FWHM for the ITS and 2.25 pixels FWHM for the HRI, the Scene Analysis images on the Flyby will be acquired at ~E-23 min. This independent site selection process is affected by: 1) difference in the scene due to nucleus rotation between the time when the HRI Scene Analysis image is taken and the time when the ITS image is taken; 2) the $\sim 3^\circ$ nucleus view angle difference due to the Flyby and Impactor being on different trajectories relative to the nucleus; and 3) different noise characteristics between the ITS and HRI detectors.

Consequently, the two s/c do not always select the same impact site. In addition, the actual impact site may differ from the one selected by the Flyby due to targeting errors, such as maneuver execution errors, on the Impactor, which result in an impact site that differs from the one targeted by the Impactor. Figure 15 shows the error distribution in the impact site selection, defined as the difference between the impact site computed on the Flyby and the actual impact site achieved by the Impactor. The figure of merit used to describe the Flyby tracking performance is a site error of 1,750 m, which is equivalent to half the HRI FOV at the time when a 7 m resolution image becomes achievable. The probability of such an error is presently 3.6%, although the mean error of ~400 m is significantly smaller.

The error in the identification of the impact site is the largest contribution to Flyby instrument pointing errors at TOI. The current requirement for a total pointing error at TOI of no more than 100 μrad (3σ), for a 128×128 pixel HRI subframe on the impact site, can be met with a 91% confidence. Using a larger HRI subframe, 512×512 pixels for example, can be met with 99% confidence. Figure 16 shows the pointing error distribution at TOI, in microradians, with a mean value of 61 μ rad.

Following imaging of the impact event and the ejecta plume expansion, the main engineering goal is to track the crater through its expansion and to capture high-resolution images of the crater. The main requirement is to capture a 7 m resolution image of the crater, with the goal of capturing a even higher, 3.4 m resolution image of the fully developed crater.

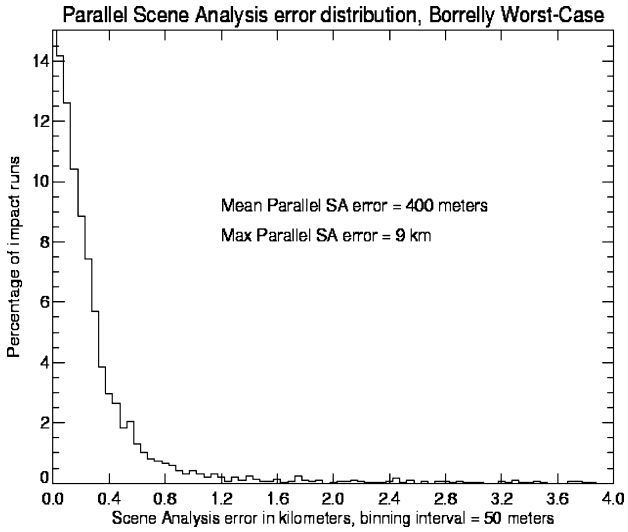

Figure 15. Error distribution in kilometer, between the actual impact site and the expected impact site as computed on the Flyby s/c at E−23 min.

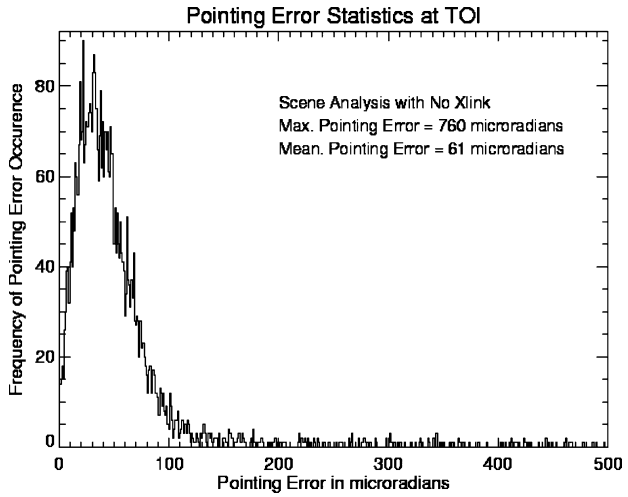

Figure 16. Pointing error at TOI. The large values at the long "tail" of the distribution are due to misidentification of the actual impact site.

As the Flyby s/c reaches its closest approach point (CA), the pointing errors present at time of impact propagate to later times. In addition, OD errors, as discussed earlier, associated with projection effects of the downtrack direction and by the rapid change of the nucleus brightness profile with the view angle, begin to

dominate. As a result, the probability of keeping the crater in the HRI FOV drops sharply after CA-120 sec, i.e. during the last 70 s of imaging. On the other hand, the pixel scale and resolution of the HRI increases linearly with time to closest approach and the competition between these two factors determine the resolution attained at any given time. For the current analysis, the resolution element is the FWHM PSF of the instrument with the addition of the smear induced by the slew of the instrument as it tracks the comet. Smear is very small up to ∼CA-120 s and increases rapidly to a mean value of 0.7 HRI pixels for a 50 ms exposure at CA-50 s. The actual crater resolution ultimately depends on crater tracking, the instrument FWHM, and the smear rate and desired exposure, with the exposure duration affecting the amount of smear.

Figure 17 shows a typical example of how the probability for attaining certain crater resolution varies relative to CA. The sharp drop during the last 1 min of imaging is typical of all scenarios examined, although, the actual probability of the crater in the HRI depends on the nucleus model with a more spherical shape giving a higher probability, and the nucleus size with a larger size giving a smaller probability. In this particular example, there is a small difference in the probability curve between the baseline and the worst-case nucleus models. The vertical dashed lines mark the first time that a 3.4 m resolution becomes possible for a certain PSF value of the HRI and for a 20 ms exposure. The two vertical dashed-dotted

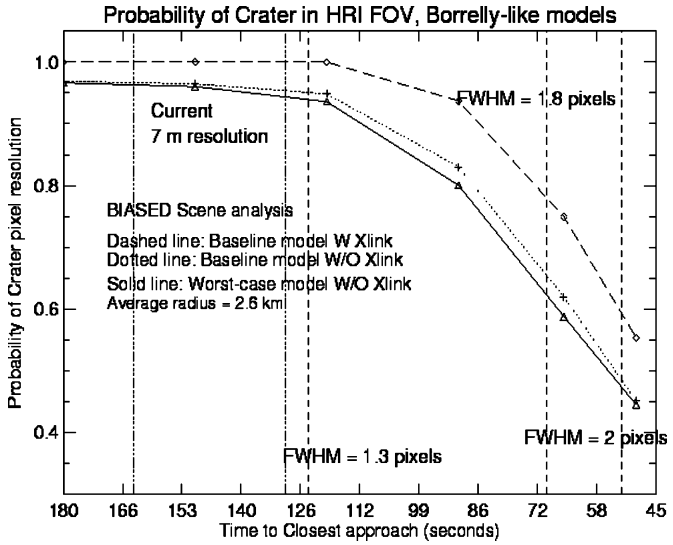

Figure 17. The probability curves for the crater in the HRI FOV as a function of time under a number of different scenarios.

lines mark the first time where a 7 m resolution image becomes possible for PSF values in the range of 2–2.5 pixels FWHM. Currently, the probability of capturing a 7 m resolution image is in the range of 95–97%; the result of propagating errors from performing Scene Analysis independently to the time of ∼CA-150 s. The corresponding probability curve, if Impactor were to send its predicted impact site to the Flyby *via* the cross-link (dashed line), shows 100% probability for a 7 m resolution image, because in this case the crater identification is correct to within 120 m (1σ). Currently, the probability for a 3.4 m resolution is quite small: 45% for a 2 pixel FWHM and for larger PSF values a 3.4 m resolution becomes untenable regardless of whether or not AutoNav can track the crater throughout the encounter.

Summary

In this article, we have described an overview of the autonomous navigation (AutoNav) as applied to Deep Impact and we have discussed the expected system performance. Clearly, this is complicated and challenging mission, which relies on the successful interaction of the three instruments (ITS, MRI and HRI), AutoNav and ADCS in a closed-loop fashion. From a navigation standpoint, robustness in the face of unexpected cometary properties, has been one of the focal points of the encounter design, operations of AutoNav and the various trades that are necessary on a cost-constrained mission. Any future work, prior to encounter, on design issues will likely be directed towards alternative scenarios for capturing high-resolution images of the crater. Such scenarios include a closer flyby altitude to increase instrument resolution and a simple one-dimensional mosaic along the slew direction, which is also the direction of the dominant pointing errors. The remainder of the work will be focused on AutoNav testing, both at the test-bench level and in-flight during Earth departure as well as on developing a limited number of encounter contingency scenarios. These are topics, which out of space limitations, are not discussed in this paper.

Acknowledgments

We are grateful to Bob Gaskell of the Optical Navigation Group at JPL for constructing all of the nuclei shape models used for navigation analyses and to George Null of the same group (retired) for establishing many of the essential guidelines and procedures upon which the current AutoNav strategies are based. We would also like to thank the anonymous referee for many helpful comments. The research 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 for the Deep Impact Project.

References

Bhaskaran, S., Riedel, J. E., and Synnott, S. P.: 1996, Autonomous Optical Navigation for Interplanetary Missions, *Science Spacecraft Control and Tracking in the New Millennium, Proc. SPIE*, pp. 32.

Bhaskaran, S., *et al.*: 1998, Orbit Determination Performance Evaluation of the Deep Space 1 Autonomous Navigation Software, *AAS/AIAA Space Flight Mechanics Meeting*, Monterey, CA.

Kirk, R. L., *et al.*: 2004, *Icarus* **167**, 54.

Riedel, J. E., *et al.*: 2000, Autonomous Optical Navigation Technology Validation Final Report, *Deep Space 1 Technology Validation Symposium*, February 8–9, Pasadena, CA.

Russ, J. C.: 1999, *The Image Processing Handbook*, CRC and IEEE Press.

Stooke, P. and Abergel, A. 2000, Halley Nucleus Shape Model, Personal Communication via M. Belton.

Trochman, W.: 2001, Impactor Spacecraft Attitude Knowledge Performance, BATC System Engineering Report, DI-IMP-ACS-010.

Weidenschilling, S. J.: 1997, *Icarus* **127**, 290.

Zarchan, P.: 1997, *Tactical and Strategic Missile Guidance*, 3rd edn., *Progress in Astronautics and Aeronautics*, Vol. 176, AIAA, Reston, VA.

Zimpfer, D.: 2003, *26th Annual Guidance and Control Conference*, February 6–10, Breckenridge, CO.