AAS 06-081

Deep Impact Autonomous Navigation: The Trials of Targeting the Unknown

Daniel G. Kubitschek, Nickolaos Mastrodemos, Robert Werner, Brian Kennedy, Stephen P. Synnott, George Null, Shyam Bhaskaran, Joseph E. Riedel, and Andrew T. Vaughan

NASA – Jet Propulsion Laboratory
California Institute of Technology

29th ANNUAL AAS GUIDANCE AND CONTROL CONFERENCE

February 4-8, 2006
Breckenridge, Colorado

Sponsored by
Rocky Mountain Section

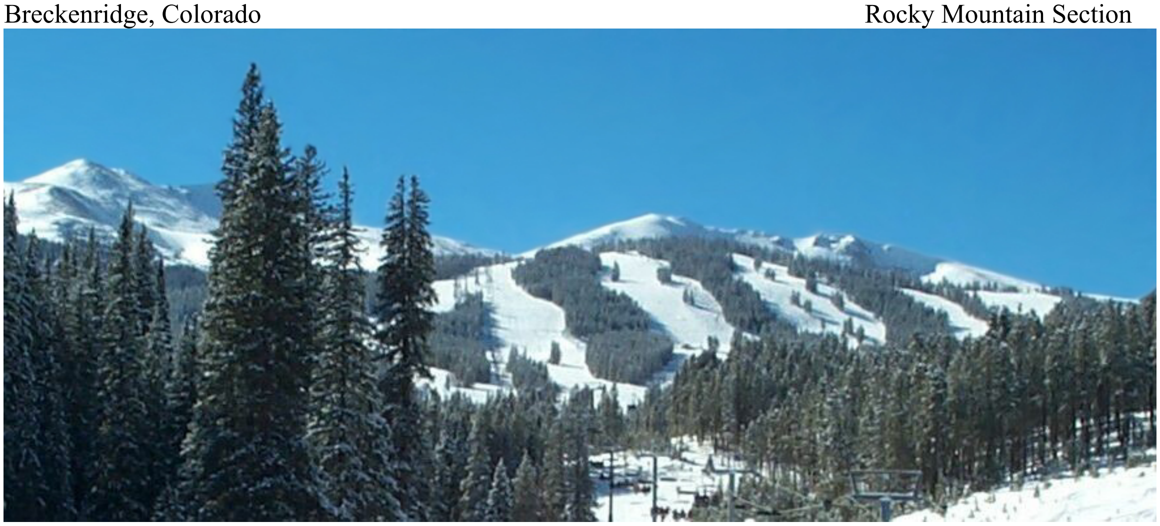

AAS Publications Office, P.O. Box 28130 - San Diego, California 92198

AAS 06-081

Deep Impact Autonomous Navigation: The Trials of Targeting the Unknown

Daniel G. Kubitschek†, Nickolaos Mastrodemos, Robert A. Werner, Brian M. Kennedy, Stephen P. Synnott, George W. Null, Shyam Bhaskaran, Joseph E. Riedel, Andrew T. Vaughan

Jet Propulsion Laboratory
California Institute of Technology

ABSTRACT

On July 4, 2005 at 05:44:34.2 UTC the Impactor Spacecraft (s/c) impacted comet Tempel 1 with a relative speed of 10.3 km/s capturing high-resolution images of the surface of a cometary nucleus just seconds before impact. Meanwhile, the Flyby s/c captured the impact event using both the Medium Resolution Imager (MRI) and the High Resolution Imager (HRI) and tracked the nucleus for the entire 800 sec period between impact and shield attitude transition. The objective of the Impactor s/c was to impact in an illuminated area viewable from the Flyby s/c and capture high-resolution context images of the impact site. This was accomplished by using autonomous navigation (AutoNav) algorithms and precise attitude information from the attitude determination and control subsystem (ADCS). The Flyby s/c had two primary objectives: 1) capture the impact event with the highest temporal resolution possible in order to observe the ejecta plume expansion dynamics; and 2) track the impact site for at least 800 sec to observe the crater formation and capture the highest resolution images possible of the fully developed crater. These two objectives were met by estimating the Flyby s/c trajectory relative to Tempel 1 using the same AutoNav algorithms along with precise attitude information from ADCS and independently selecting the best impact site. This paper describes the AutoNav system, what happened during the encounter with Tempel 1 and what could have happened.

INTRODUCTION

On July 3, 2005, NASA's Deep Impact Flyby spacecraft (s/c) released a small, 370 kg Impactor s/c 24 hrs before the planned time-of-impact (TOI). The Impactor s/c was designed to target comet Tempel 1, which was estimated to be 14 km x 5 km x 5 km in size at the time of release; the size, shape and orientation of the nucleus would not be known to either the Impactor or Flyby s/c for another 22 hrs. With a closing speed of approximately 10.3 km/s, the Impactor s/c autonomously guided itself to impact via 3 discrete propulsive targeting maneuvers and captured the highest resolution images ever of the surface of a cometary nucleus. The primary objective of the Impactor s/c was to impact in an illuminated area viewable from the Flyby s/c using autonomous navigation (AutoNav) algorithms and precise attitude information from the Attitude

† Corresponding Author. Jet Propulsion Laboratory, California Institute of Technology, NASA-Jet Propulsion Laboratory, M/S 301-150, 4800 Oak Grove Drive, Pasadena, CA 91109, 818.393.1101 (phone), 818.393.6388 (fax), daniel.kubitschek@jpl.nasa.gov

Determination and Control System (ADCS). The secondary objective was to capture high-resolution context images of the impact site for Science.

Meanwhile, the Flyby s/c autonomously acquired lock and tracked the nucleus center of brightness (CB) until 4 min before impact, after which time, the Flyby s/c began tracking the predicted impact site. The flash associated with the impact event was captured in both the medium resolution imager (MRI) and the high resolution imager (HRI). The Flyby s/c continued to track the impact site for the full 800 sec with the MRI. In addition, the impact site was captured in one of the last three HRI images taken just prior to shield attitude entry. The primary objective of the Flyby s/c was to capture the impact event in the MRI and HRI, observe the eject plume expansion dynamics and track the impact site until shield attitude entry to acquire the highest resolution images of the fully developed crater just before shield attitude entry.

Deep Impact Mission Overview

Deep Impact was a dual s/c mission launched on January 12, 2005 with the engineering goals of impacting comet Tempel 1 on July 4, 2005, observing the impact event and ejecta plume expansion, obtaining IR images of the ejecta and high resolution images of the fully developed crater using the Medium Resolution Imager (MRI) and the High Resolution Imager (HRI) on the Flyby s/c for the scientific purpose of exposing and understanding the interior composition of a comet nucleus.

After a brief 6-month cruise, the two spacecraft separated 24 hrs prior to the expected TOI. The encounter geometry resulted in an illumination phase angle of approximately 64° for the Tempel 1 nucleus. The Flyby s/c performed a slowing maneuver with a ΔV of approximately 102 m/s to provide 800 ± 20 sec of post-impact event imaging and control the flyby miss-distance to 500 ± 50 km. During the first 22 hrs following release, the Impactor s/c acquired and telemetered science and navigation reconstruction images to the ground using the Flyby s/c as a bent-pipe relay. The Flyby s/c also acquired and telemetered MRI and HRI visible and HRI infrared (IR) images of the nucleus and coma.

The autonomous phase of the encounter began at 120 min (2 hrs) before TOI. A critical sequence running on-board both the Impactor and Flyby s/c spawned science and navigation subsequences that issued Impactor Targeting Sensor (ITS) commands and MRI commands to produce navigation images at a 15 sec interval.

The Autonomous Navigation (AutoNav) software was originally developed and demonstrated during the Deep Space 1 (DS1) mission[1,2]. It is considered an enabling technology for high-speed encounter missions to small bodies. It has had a number of successes such as the DS1 encounter with comet Borrelly in September 2001 and the Stardust encounters with the asteroid Anne Frank, and finally, comet Wild-2 in January 2004. For these missions, the objective was to keep the target body in the instrument FOV. Deep Impact AutoNav would take the high speed encounter to a new level and that level would required a few new algorithms: 1) Biased Scene Analysis; 2) Time-of-Impact (TOI) and Time-of-Final-Imaging (TOFI) updates; and 3) Autonomous Coma Cutoff (ACC), to target a particular location on the comet nucleus, synchronize the timing of imaging sequences on two s/c, capture the impact event in an instrument with an effective FOV of 640 μrad at the time-of-impact, track that location until the very end with an instrument whose FOV spanned 2 mrad and finally provide the timing necessary to orient the Flyby s/c for passage thru the inner coma.

On Deep Impact, AutoNav was responsible for processing the navigation images to form observations for the purpose of trajectory determination (OD), computing Impactor targeting maneuvers (ITM) that were executed by the Attitude Determination and Control System (ADCS), using Scene Analysis on both s/c, in parallel, to target a particular location on the nucleus and to point the MRI and HRI instruments at the expected impact site, and providing updated timing information to optimize the start-time of Science imaging sequences. Figure 1 shows a schematic diagram of the encounter activities for both s/c.

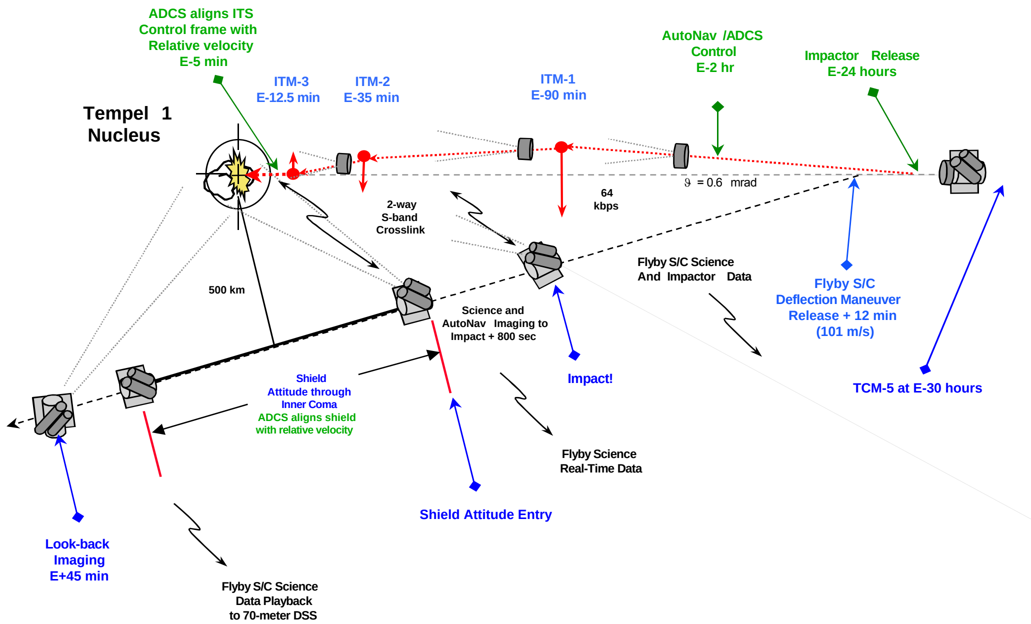

Figure 1 Schematic of Deep Impact Encounter with comet Tempel 1

Flyby Spacecraft Flight System

The Flyby s/c, shown in Figure 2, was designed and built at Ball Aerospace Technologies Corporation (BATC) and consisted of a split-panel, deployable solar array for power to the Flyby subsystem components and to the Impactor subsystem components while the two s/c were mated during cruise, a High Gain Antenna (HGA) for high-rate, ground uplink (U/L) and downlink (D/L), two Low Gain Antennae (LGA) for communication during periods when the HGA cannot be used, two redundant RAD750 processors, the MRI and HRI (visible and infrared) instruments, a S-band communications link to the Impactor s/c, 4 reaction wheel assemblies (RWA) for 3-axis attitude stabilized attitude control and reaction control system (RCS) thrusters for backup control, an ADCS system that estimates the International Earth Rotation Service (IERS) celestial reference frame (ICRF) attitude based on observations from 2 CT-633 StarTrackers and rates from two, redundant Northrup-Grumman Space Inertial Reference Units (SSIRU), which also provides linear acceleration measurements for autonomous trajectory integration. Communication between the ground and the Impactor s/c was established via hardline link between the Flyby s/c and the Impactor s/c before separation and via the S-Band transmitter/receiver link between the Flyby s/c and Impactor s/c following separation. The Flyby s/c was not configured to decommutate the signal coming from the Impactor. Instead the data was stored in non-volitile memory (NVM) on the Flyby s/c in the form of Impactor Spacecraft Interface (ISI) files, given a default priority and placed in the downlink stream for transmission to the ground.

The MRI camera had a 12 cm aperature (73.5 cm^2 collecting area with 35% obscuration), a focal length of 1.2 m and a 10 milliradian (mrad) field-of-view (FOV). The 1024x1024 pixel CCD is a split-frame transfer device with electronics that provides 14-bit digitization (16384 DN full-well). The MRI served a dual purpose: 1) provide autonomous navigation images and 2) provide high-rate images during the impact event for highest possible temporal resolution.

The HRI visible camera had a 30 cm aperature and a 2 milliradian (mrad) field-of-view (FOV). The 1024x1024 pixel CCD is a split-frame transfer device with electronics that provides 14-bit digitization (16384 DN full-well). The HRI served a dual purpose: 1) provide primary Scene Analysis autonomous tracking images and backup autonomous tracking images and 2) provide high-resolution images during the impact event and during final imaging of the fully developed crater to satisfy primary mission science objectives.

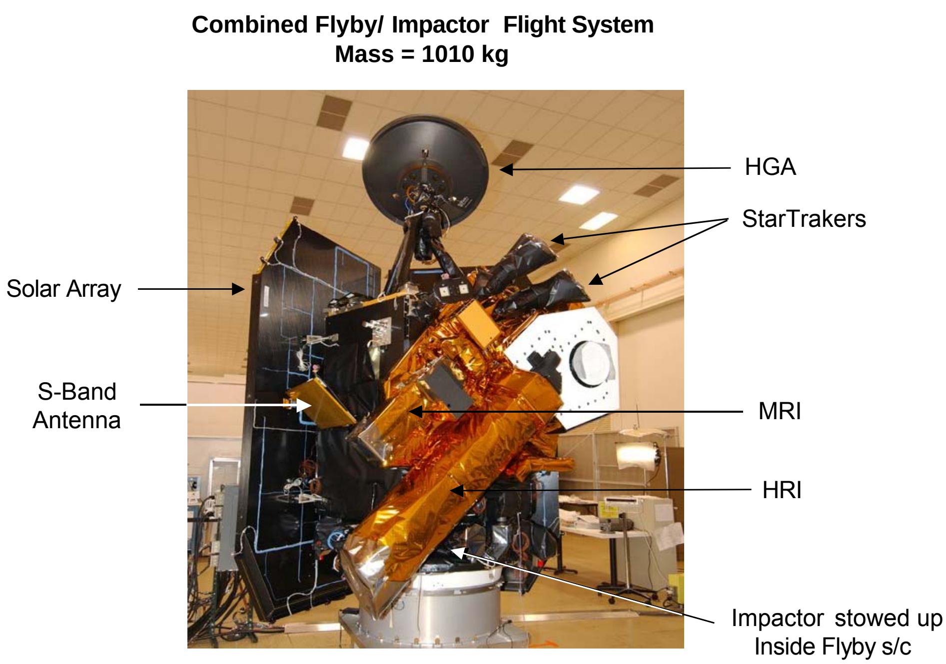

Figure 2 Flyby s/c Flight System configuration as seen during Integration & Testing

Impactor Spacecraft Flight System

The Impactor s/c, shown in Figure 3, consists of a battery for power after release, a RAD750 computer (SCU) for processing and command and data handling, an Impactor targeting sensor (ITS), which is a simple inverting telescope with a charge couple device (CCD) detector, a S-Band communications link to the Flyby s/c, a 3-axis stabilized attitude and rate control system (RCS), a 4 divert/4 RCS thruster hydrazine propulsion system with a ΔV capability of 25-30 m/s, and an ADCS system that estimates the International Earth Rotation Service (IERS) celestial reference frame (ICRF) attitude based on observations from a single CT-633 StarTracker and rates from a Northrup-Grumman Inertial Reference Unit (SSIRU), which also provides linear acceleration measurements for autonomous navigation. The mass of the s/c was approximately 370 kg with an all-copper fore-body cratering mass.

The ITS camera had a 12 cm aperature (73.5 cm^2 collecting area with 35% obscuration), a focal length of 1.2 m and a 10 milliradian (mrad) field-of-view (FOV). The 1024x1024 pixel CCD is a split-frame transfer device with electronics that provides 14-bit digitization (16384 DN full-well). The ITS served a dual purpose: 1) provide navigation images and 2) provide pre-impact high resolution (< 3 m) science images**.**

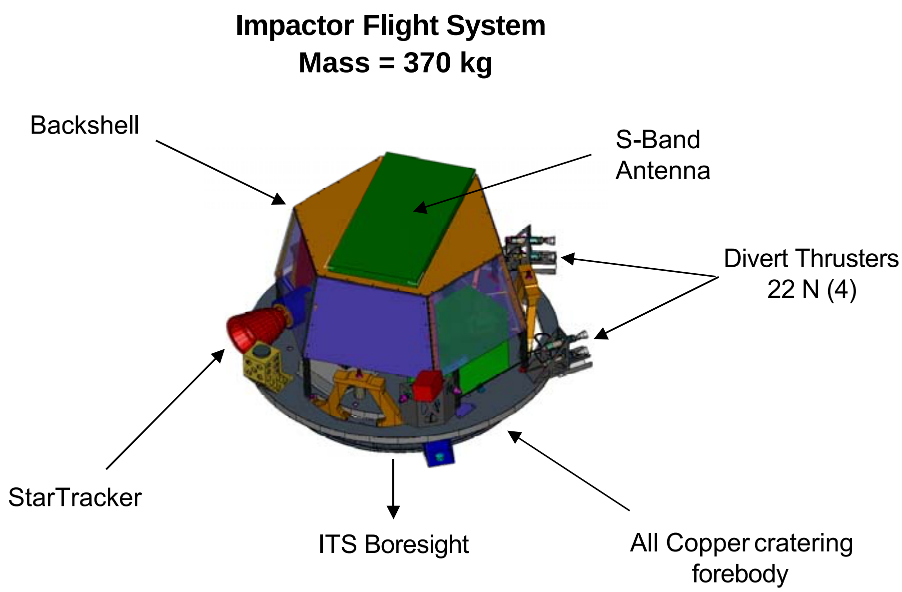

Figure 3 Impactor s/c flight system configuration[3]

Deep Impact Autonomous Navigation (AutoNav) System

The AutoNav system for Flyby s/c instrument pointing control and terminal guidance of the Impactor s/c relies on both the performance and interaction of the AutoNav and ADCS flight software and subsystems, MRI and the ITS camera. AutoNav consists of 3 distinct modules: 1) Image processing; 2) Orbit determination; and 3) Maneuver computation. AutoNav was originally developed to operate in two different modes: 1) 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 2) 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 was used based on the expected quality of the ADCS estimated attitudes. The combination of the CT-633 StarTracker(s) and SSIRU rate sensor was expected to provide an estimated attitude bias or attitude knowledge error (AKE) of no more that 150 μrad (3σ), bias stability of 50 μrad/hr (3σ), and estimated attitude noise of 60 μrad (3σ).

From a flight system perspective, there are 3 key interfaces between AutoNav and rest of the flight and ground system(s): 1) Flyby AutoNav-to-Ground interface; 2) AutoNav-to-ADCS interface; and 3) AutoNav-to-Sequence interface. Figure 4 shows the AutoNav-to-ground ground interface. Here, AutoNav provides ground navigation with snippeted approach images of Tempel 1 for ground-based orbit determination. In return, the ground-based navigation process provided the best available trajectory information (comet and spacecraft) for initializing AutoNav.

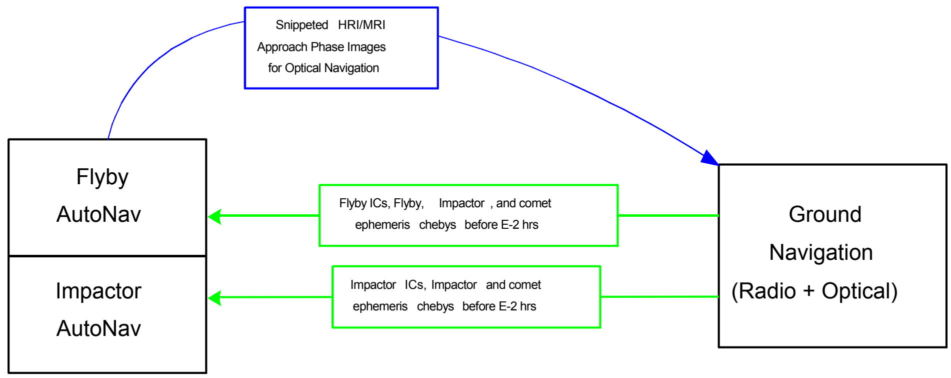

Figure 4 AutoNav interface with ground navigation team

The most important interface, in terms of AutoNav performance was the AutoNav/ADCS interface shown in Figure 5. When processing images in the starless AutoNav mode, AutoNav requires external s/c body attitude relative to ICRF. ADCS provides estimated attitude and attitude rate information to AutoNav in the form of a history buffer. When AutoNav receives a navigation image, the time of center of the image exposure is used to interpolate the ADCS attitude history at that time. The attitude quaternion is composed with the camera alignment quaterion to provide the inertial attitude of the navigation instrument, which is used by AutoNav to predict the pixel/line location of the target body in the instrument FOV. In addition, ADCS provides acceleration measurements to AutoNav for trajectory integration. This is particularly important for targeting and autonomous trajectory correction maneuvers on the Impactor s/c. The acceleration measurements are stored as accumulated delta-V (ΔV) in the non-gravitational acceleration history file. This file represents the actual inertial accelerations experienced by the Impactor s/c, primarily due to thruster pulses (RCS and divert). During autonomous operations on the Flyby s/c, there were no planned maneuvers and records written to the non-gravitational acceleration history file were negligible, although, they were included in the trajectory integration associated with each orbit determination update.

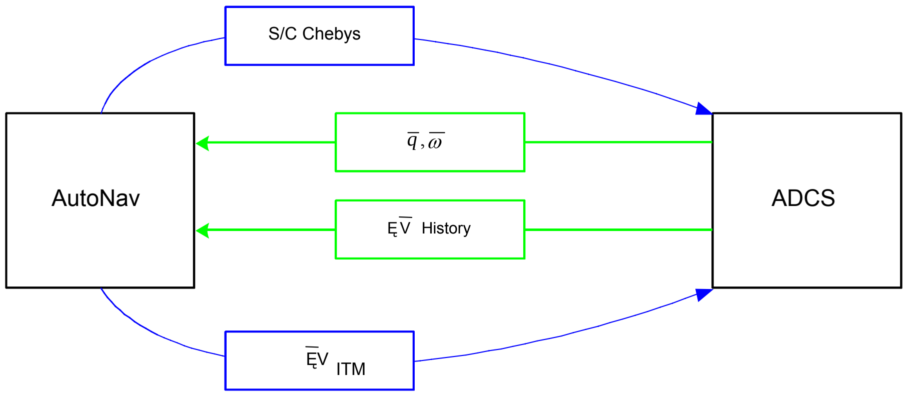

Figure 5 AutoNav interface with ADCS subsystem

The final important key AutoNav interface is related to optimizing the start time of imaging sequences during critical events such as TOI and Time-of-final-imaging (TOFI). Figure 6 gives a systems-level view of this AutoNav function. The algorithm is describing in the following section.

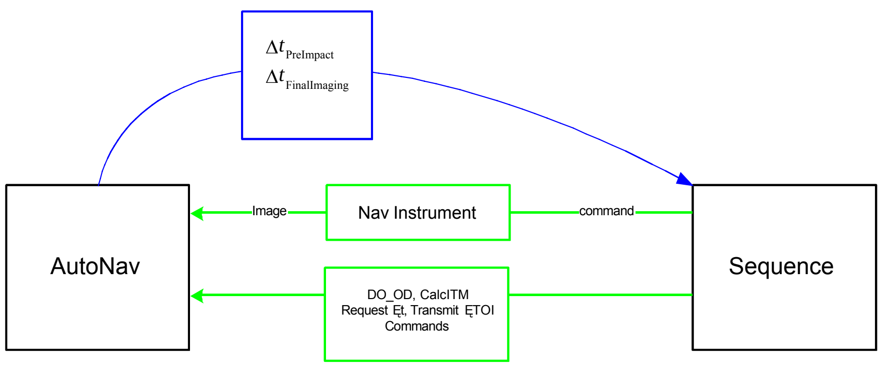

Figure 6 AutoNav interface with sequencing

AutoNav Based Time-of-Impact (TOI) Update

During AutoNav operation (last 2 hrs of encounter) the heliocentric position of the Flyby s/c is estimated in the OD process based on CB observations of the comet nucleus. The heliocentric ephemeris of the comet is assumed known. There are, however, errors in the comet's heliocentric position[4]. These errors propagate into the solution for the heliocentric position of the s/c, but the error in the comet-relative position and velocity is removed.

In order to determine the TOI on-board the Flyby s/c, the heliocentric ephemeris of the Impactor s/c, based on pre-release ground-based combined radio navigation and optical navigation observations (RadioNav/OpNav) and the expected change in velocity due to separation, detumble, and the heliocentric ephemeris of the Flyby s/c, based on post-deflection ground-based RadioNav/OpNav must reside on the Flyby in the form of Chebyshev coefficients that represent the position up to the impact event. AutoNav provided an updated heliocentric position of the Flyby s/c. Differencing the ground-based Flyby position and the AutoNav based Flyby position at any given time, gives the heliocentric correction to the a priori comet ephemeris. This inertial correction vector is added to the a priori heliocentric position of the comet and the resultant is differenced with the ground-based estimate of the heliocentric position of the Impactor to give an Impactor relative range, which when divided by the relative velocity of the Impactor, gives the time to impact.

AutoNav Based Time-of-Final-Imaging (TOFI) Update

The TOFI is important in that it is used to initiate the final Science imaging sequence based on when the s/c must transition to shield attitude for flight-system safety reasons. In order to determine the time of shield mode: time when s/c comet-relative range is 700 km, on-board the Flyby s/c, the heliocentric ephemeris of the target body (Tempel 1), and the heliocentric ephemeris of the Flyby s/c are used to compute the time of flight (TOF) from a given epoch time to the point of closest approach. In addition to the updated heliocentric position of the Flyby s/c, AutoNav provides the comet-relative state of the Flyby s/c in the following form:

$$\overline{X} = \begin{Bmatrix} b \cdot t \\ b \cdot r \\ TOF \\ s \cdot t \\ s \cdot r \\ V_\infty \end{Bmatrix} \tag{1}$$

Where, $b \cdot r$ and $b \cdot t$ are the components of the s/c position in the B-plane at intercept; $s \cdot r$ and $s \cdot t$ are the components of the s/c's downtrack position projected onto the B-plane; *TOF* is the time-to-go before B-plane intercept; and V_∞ is the comet-relative velocity, which is normal to the B-plane. The impact parameter or the B-vector magnitude is given by

$$B = \left|\overline{B}\right| = \sqrt{(b \cdot t)^2 + (b \cdot r)^2} \tag{2}$$

It was decided that shield mode should occur at a range of 700 km. Therefore, the downtrack location, relative to closest approach, will be (as shown in Figure 7).

$$S' = \sqrt{(700)^2 - B^2} \tag{3}$$

The absolute time of shield mode is then

$$t_{Shield} = t_{Epoch} + TOF - \left(\frac{S'}{V_\infty}\right) \tag{4}$$

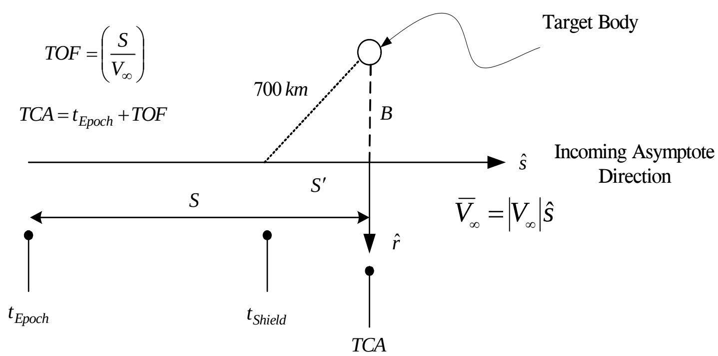

Figure 7 Downtrack distance from the B-plane

Sequence design requires that the nominal image sequence be specified based on the best available a priori ephemeris information (pre-release) in the event of a failure to provide AutoNav-updated sequence delta times. This delta time is combined with the sequence vehicle time code ($VTC_{derived} = VTC_{actual}$ + delta-time), in an additive sense to effectively move the time of all sequence events that follow. If the delta-time value is negative, this has the effect of delaying execution of sequence events (events will occur later); if the delta time value is positive, this effectively moves the VTC time forward, advancing the time of execution (events will occur earlier).

Flyby Spacecraft Autonomous Tracking

The steps involved in the Flyby autonomous tracking process were as follows:

1. Acquire MRI images of the comet nucleus, every 15 sec, starting 2 hrs before the expected time of impact
2. Process MRI 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 trajectory determination updates (OD), every 1 min, starting 1 hr 50 min before the expected time of impact (first OD arc had 40 observations)
5. Acquire 3 MRI images for computing an Scene Analysis-based offset, relative to observed CB, just prior to expected time-of-impact (TOI)
6. Transmit two (for redundancy) deltaTOI offsets to the Impactor s/c for optimizing the start of the final imaging sequence on the Impactor
7. Apply the deltaTOI offset on-board the Flyby s/c to optimize the start of the imaging sequences design to capture the impact event
8. Compute and apply the deltaTOFI offset on-board the Flyby s/c to optimize the remaining time for final imaging and transition to shield attitude

The MRI was designated the prime AutoNav instrument during autonomous operations, however, the sequences were designed to allow Fault Protection to change the prime select designation and default to using the HRI for AutoNav should a problem with the MRI be detected.

MRI image processing was accomplished using three different algorithms: 1) Blobber; 2) Centroid Box; and 3) Scene Analysis. The Blobber algorithm was used during the first 60 min of autonomous operations, since it scans the entire image looking for regions of contiguously lit pixels above the brightness threshold and therefore requires no a priori knowledge of where the target will be located in the MRI FOV. The remainder of the encounter was accomplished using the Centroid Box algorithm, which consists of a simple moment algorithm using those pixels within an NxN centroid box centered on the predicted location of the target body in the MRI FOV. For Deep Impact, a 400x400 centroid box was selected up to the time-of-impact. The sequence then had a built-in 3 min imaging gap to allow for the post-impact ejecta to dissipate and the remainder of the encounter used the entire MRI FOV as the centroid box due to the increasing size of the nucleus.

The image processing module had two means of applying a brightness cutoff to primarily remove the influence of off-nucleus coma on the Blobbing and Centroid Box processes: 1) a fixed brightness threshold below which all pixels brightness would be zeroed and 2) an autonomous coma cutoff (ACC) that dynamically determined the minimum allowable pixel brightness based on the average peak signal observed in each image. For the ACC algorithm, the brightest 5 pixels were thrown out of the process. The next 5 brightest pixels were averaged to give the average peak. Finally, 35%, a parameter specifiable in flight-software, was applied to the peak to arrive at the cutoff threshold. The 35% parameter was a factor of two greater than the peak nucleus brightness to peak coma brightness ratio recommended by the Science Team.

The purpose of the Scene Analysis algorithm was to allow the Impactor to target an illuminated area that would be viewable from the Flyby s/c and allow the Flyby s/c to independently arrive at that same location to track the impact site. The HRI had originally been designated as the prime Scene Analysis instrument, but an in-flight focus problem detected post-launch required that we abandon that approach and use the MRI for Scene Analysis as well. This will be discussed in more detail in the section dealing with the in-flight adjustments.

Impactor Spacecraft Autonomous Targeting

The Impactor s/c used a predictive guidance strategy and pulsed guidance system consisting of 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 (a priori comet Tempel 1 ephemeris) at the time of impact to compute the “zero effort” miss distance which is then used to compute the magnitude and direction of each ITM.

The a priori position of the target body had significant uncertainty prior to start of AutoNav operations, which was removed using optical navigation techniques. On the other hand, The dynamics were well-known, except that (an important exception) the nucleus rotational dynamics and solar phase angle combine to induce motion (translational velocity and acceleration) of the center of brightness (CB) with time, which can cause targeting errors on the surface of the nucleus via over-estimation of the lateral velocity. These were 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 suggested a predictive guidance strategy for Deep Impact and we selected a 20 min OD arc length.

The steps involved in the Impactor autonomous guidance process were as follows:

1. Acquire ITS images of the comet nucleus, every 15 sec, starting 2 hrs before the expected time of impact
2. Process ITS 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 trajectory determination updates (OD), every 1 min, starting 1 hr 50 min before the expected time of impact (first OD arc had 40 observations)
5. Perform 3 primary Impactor targeting maneuvers (ITMs) at 90 min (ITM-1), 35 min (ITM-2), and 12.5 min (ITM-3) during the terminal guidance phase
6. Acquire 3 ITS images for computing an 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. Perform the final targeting maneuver (ITM-3) 12.5 min before predicted time of impact
8. Align the ITS boresight with the AutoNav estimated comet-relative velocity vector starting 5 min prior to predicted time of impact to capture and transmit high-resolution images (< 3 m resolution) of the nucleus surface

The reason for selecting the AutoNav starting time at E-2 hrs was due to the need to correct for as much as 30 km of delivery error in the B-plane with ITM-1 and a 30 km maneuver at E-100 min would require ~5 m/s of delta-V, but the time of ITM-1 was moved for reasons that will be discussed in a later section. The Impactor s/c had a delta-V capability of 25 m/s allocated for targeting maneuvers. The remainder of the propellant was to be used by the RCS system for attitude control during the 24 hr free-flight. Selection of the 20 min OD arc length was a result of 10s of 1000s of Monte Carlo simulations with various nucleus models and model parameter assumptions. The OD arc length had to be long-enough to provide robustness in terms of the number of observations, but short-enough to allow the solution to respond to motion of the observed CB, described above. Placing ITM-1 at E-100 min and using a 20 min OD arc length put the start time at E-120 min.

THE CHALLENGES OF AN UNKNOWN ENVIRONMENT

The most difficult aspect of the Deep Impact Mission was that we had little or no definite information about what we could expect. The Science Team provided best estimates for the nucleus size: 3.2 km mean radius, axial ratio: 3:1, albedo: 0.04, rotation period: 42 hrs, peak brightness: ratio of peak jet brightness to peak nucleus brightness of 0.18, all of which influenced our targeting/tracking strategy and affected the performance of AutoNav. We had absolutely no information about what the initial orientation of the nucleus would be at encounter and what the large or small-scale topography would be. Considering the 64° phase angle, significant self-shadowing was a real possibility and could seriously affect the crater illumination. We had no information regarding the potential brightening due to the ejecta plume, which could have a significant affect on the observed CB and influence the Flyby's ability to track the impact site and capture the high-resolution images. In addition, the uncertainties in the heliocentric position of the nucleus were estimated to be ~1200 km (3σ) in the direction of the comet-relative velocity vector at encounter, which is a potentially dominating error source for Impactor targeting and required the use of AutoNav to update the imaging sequence start timing on both s/c. Because the two s/c were mated, we were unable to perform full functional tests of the Impactor s/c prior to separation at E-24 hrs, in particular, propulsion system and closed-loop AutoNav/ADCS tests were not possible. Finally, because there was no cross-link capability, the two s/c had to operate in an independent, parallel Scene Analysis targeting/tracking mode; another complexity that could allow the two systems to get out-of-synch.

Nucleus Modeling for AutoNav Encounter Simulation

The need for Scene Analysis became clear early in the development of AutoNav for Deep Impact. Initial testing focused on nucleus models based on the theory of accretion.[5] Figure 8 shows various accretion models in several different orientations. The upper right image illustrates the case where the computed CB may not necessarily be illuminated: targeting the CB could lead to impact in a dark location. Scene Analysis scaned the entire image, computed the CB location as a reference point, and proved a pixel/line offset, relative to the CB, that was targeted by the Impactor s/c and tracked by the Flyby s/c. At the beginning of Phase B development, neither DS1 had encountered comet Borrelly (September 2001), nor had Stardust encountered comet Wild-2 (January 2004). The best images, of any cometary nucleus, at that time were taken by the Giotto s/c during its encounter with comet Halley in 1986, but the nucleus spanned only a few pixels in those images. Nonetheless, the AutoNav team would make use of models based on Halley data[6] and shown in Figure 9 and would settle, based on Science Team recommendations, on a model that was a distorted representation of the Borrelly nucleus observed by DS1. In addition, the Stardust images of Wild-2 would be used to qualitatively assess the performance of the Biased Scene Analysis algorithm.

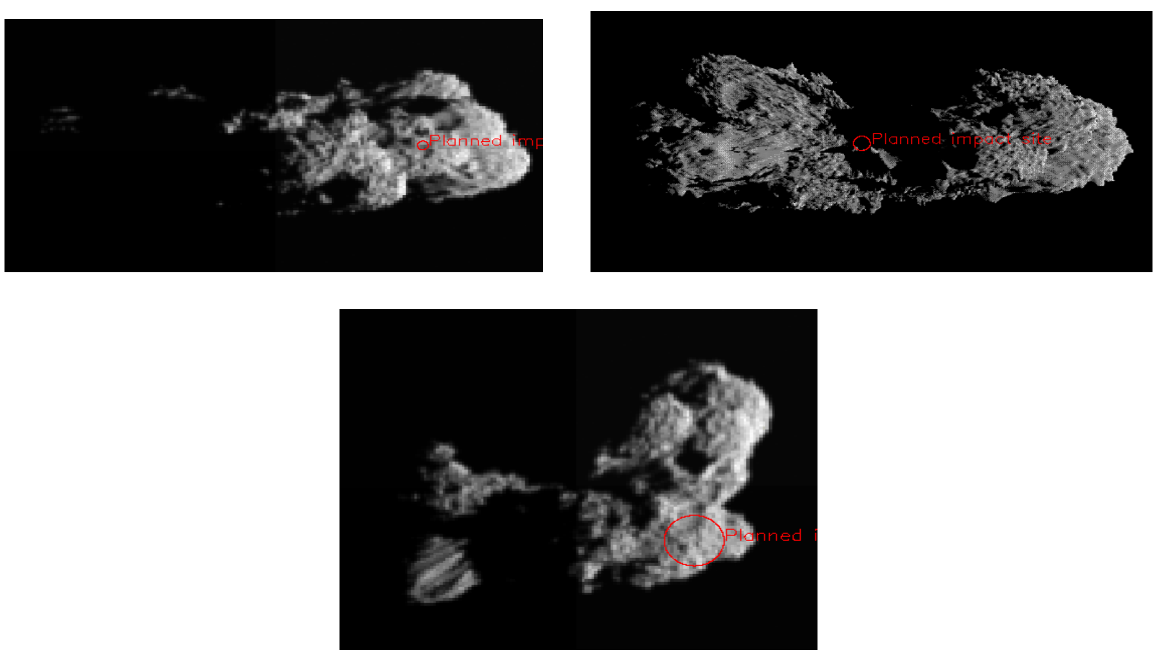

Figure 8 Theoretical nucleus models based on accretion theory

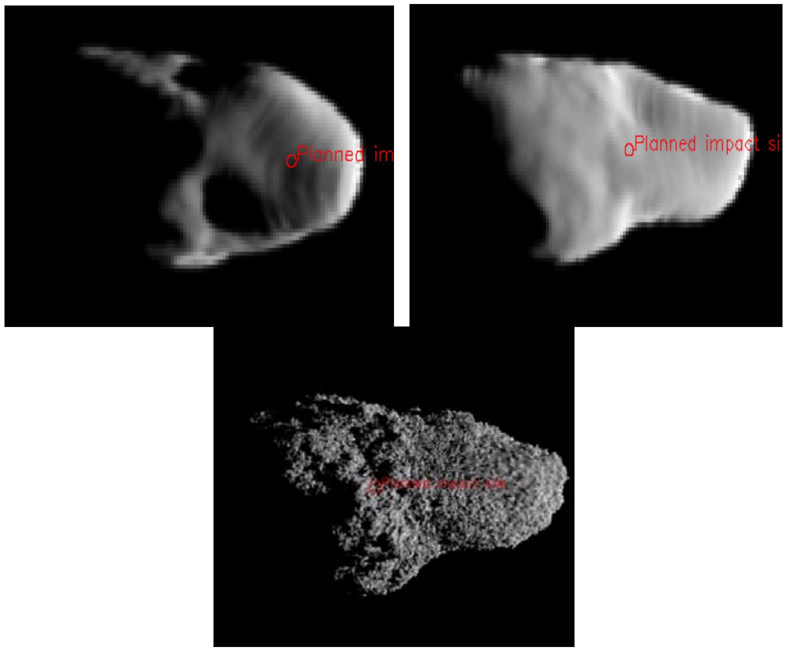

Figure 9 Theoretical Halley models based on shape models of Stooke

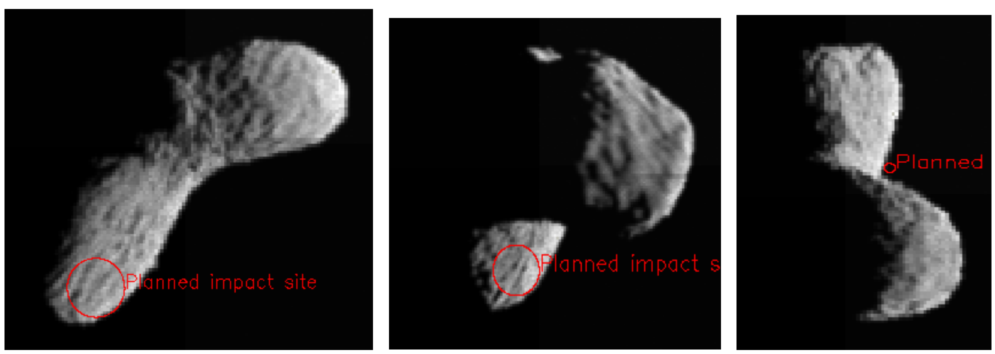

Figure 10 Borrlley-like models that were bent in the middle by up to 30°

In the case of the Borrelly-like models, shown in Figure 10, it was clear that not only could the CB be in the dark, it could be off-nucleus all together (far right image in Figure 10). And, the concave nature could lead to the impact site being obscured by the foreground lobe, not to mention the influence it could have on the Flyby's ability to track the impact site all the way to shield mode and the time of highest-resolution imaging. Figure 11 shows the case where the impact site is "over-the-hill".

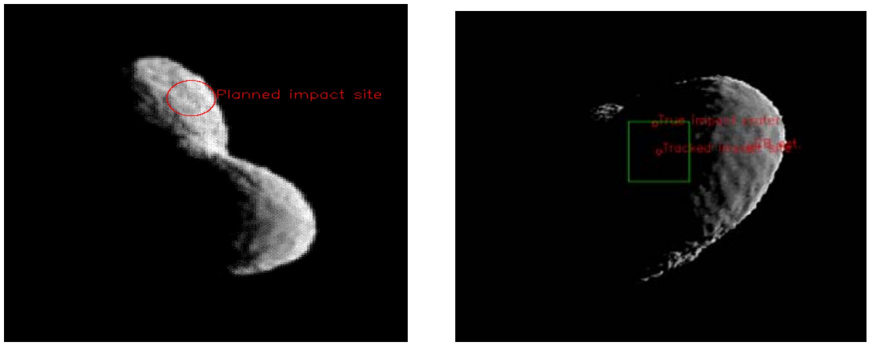

Figure 11 The left picture shows the site selected by the Impactor. The right picture shows the perspective from the Flyby at the time-of-final-imaging where the Impact site is obscured

Because of the potential for the impact site to be occulted by the foreground when dealing with nucleus models that have a large axial ratio and concave large-scale topography, Mastrodemos created a modified version of Scene Analysis that would "bias" the site selection to toward the Flyby s/c point of closest approach. The AutoNav algorithms were modified to make use of a priori knowledge of the Flyby trajectory on the Impactor s/c, such that selection of the best candidate site for Impactor targeting would have an additional criteria that would drive the solution to toward the Flyby s/c and increase the probability of good viewing from the Flyby s/c at the time of shield mode entry, without compromising either the primary objective: impacting the nucleus or the secondary objective: impacting in an illuminated area.

Center of Brightness Stability

Although it was decided that the CB would not be the target, it remained the reference relative to which the position and velocity of the Impactor and Flyby s/c were estimated. As a result, the stability or motion of the CB would have to be investigated. For the Impactor s/c, excessive CB motion is perceived as lateral motion and absorbed into AutoNav's estimate of the comet-relative velocity. The motion can lead to under or over-estimation of the comet-relative velocity used to compute the ITMs for impacting a particular site on the nucleus and thus increase the targeting errors. Figure 12 shows the motion of the CB relative to the center of mass of a simulated nucleus as it is rotated thru 360° and illuminated with a 64° phase angle. As previously mentioned, the best-estimate of the rotation period was 42 hrs. Over the OD arc length selected for Deep Impact: 20 min, the nucleus could rotate as much as 2.9°. The steepest slope on the curve in Figure 12 shows approximately 1 km of motion in 15° of rotation, which implies ~200 m of motion over the length of the OD arc, which then abruptly ends at around the 210° orientation and begins to reverse direction.

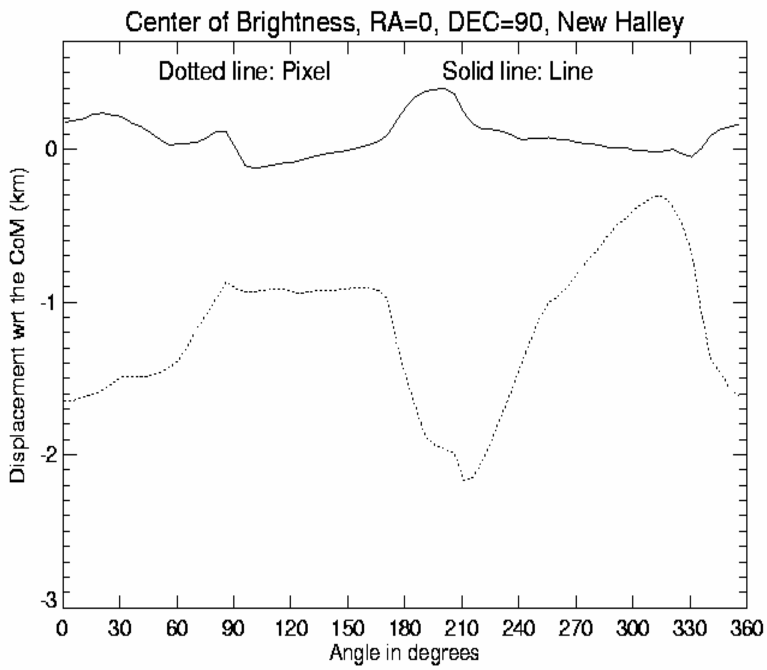

Figure 12 Simulated CB motion in the B-plane for a Halley nucleus model illuminated at a 64° phase angle

In addition to the CB motion that results from physical rotation of the nucleus, Figure 13 shows how the tracking is pulled away from the impact site (downward) due to the changing view angle that the Flyby s/c sees during encounter. On the left is the nucleus as seen by the Impactor s/c at ~E-11 min and on the right, the nucleus as seen from the Flyby s/c near the time of shield-mode entry.

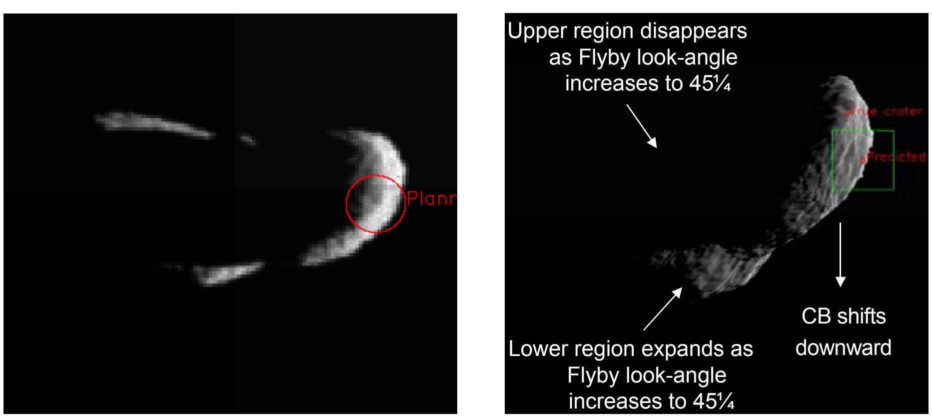

Figure 13 Tracking errors induced by concave nature of Borrelly nucleus model illuminated with a 64° phase angle

Finally, the post-impact ejecta plume brightening was thought to have the potential for significantly influencing the tracking by moving the observed CB. It did just that and the effect will be discussed in the section on encounter performance.

In general, studying the CB motion led to a counter-intuitive conclusion. The Science Team had been trying for years to understand the rotational dynamics (pole direction and rotation rate) and many had advocated that the mission adjust the s/c trajectory to time the impact event such that the "broad-side" of the nucleus would present itself during the last two hours of encounter to maximize the probability of impact; not necessarily and illuminated impact. While this might seem intuitively correct, it can be seen from the CB motion, that with a 64° phase angle, the "broad-side" would likely coincide with the steepest part of the curve and in turn lead to the largest CB motion-induced targeting errors. Combine this with the fact that the Project had agreed to use Biased Scene Analysis, and you might just have a potential for failure: increased targeting errors in the presence of an algorithm that necessarily drives the impact site selection to the edge of the nucleus (within a parameter specified distance).

Time-of-flight Errors

As was previously mentioned, the dynamics of the targeting problem were well-known, but the initial conditions: comet-relative position, could have significant uncertainty. In the days leading up to encounter a devoted effort was made to minimize the uncertainty in the ephemeris of comet Tempel 1, but the uncertainty in the time-of-flight direction could not be guaranteed to better than 1200 km (3σ)[4]. With an incoming speed of 10.3 km/s, the initial error in the TOI could be as much as 116 sec. Figure 14 shows the influence of TOF errors on the targeting.

Here,

t_{SA_OBS} ≡ time of Scene Analysis observation, which is simply an angular measurement, α
t_{MAN} ≡ time at which the targeting maneuver is executed to intercept at Δx
TOI_{nom} ≡ Expected time-of-impact
TOI_{actual} ≡ Actual time-of-impact (late)

The maneuver computation software uses Δx to compute the magnitude of the maneuver at t_{MAN}, which results in rotating the relative velocity vector by an angle, θ. As can be seen from Figure 14, this only works if

$$TOI_{nom} = TOI_{actual} \text{ or } t_{SA\_OBS} = t_{MAN} \qquad (7)$$

The first represents the TOF error and the second can never be the case due to the non-deterministic direction of the ITM, which requires that there be time between when the maneuver is computed and when the maneuver is executed to allow for the s/c to reconfigure and reorient for maneuver execution.

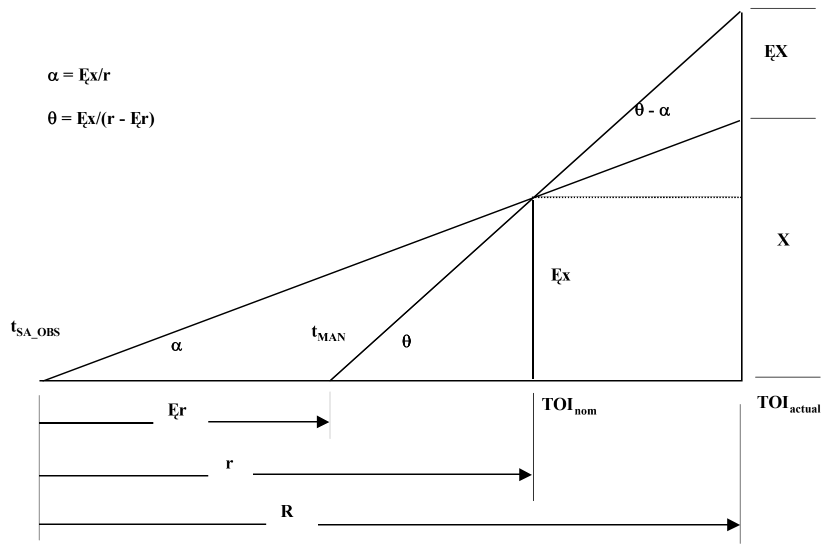

Figure 14 Targeting error due to uncertainty in the downtrack range to the target body

The relationship between the targeting error and the TOF error can be written as

$$\Delta X = (\theta - \alpha)*V_\infty*(TOI_{actual} - TOI_{nom}) = \Delta x[(1/(r-\Delta r) - 1/r)]*V_\infty*TOF_{Error} \quad (8)$$

where

$$\begin{aligned} &\Delta X \rightarrow 0 \text{ as } TOF_{Error} \rightarrow 0 \\ &\Delta X \rightarrow 0 \text{ as } \Delta r \rightarrow 0 \\ &\Delta X \rightarrow 0 \text{ as } \alpha \rightarrow 0 \end{aligned} \quad (9)$$

Figure 15 shows the family of curves representing the targeting error as a function of TOF_{Error} and ΔV magnitude for ITM-3 with the going-in strategy of commanding the Scene Analysis images at E-11 min and executing the finite maneuver such that it would be centered on E-7.5 min. It was clear that a 180 sec TOF error combined with the need for a 10 m/s ITM-3 could easily lead to a targeting error ~1 km. In the next section we discuss the necessary in-flight adjustments, which turned out to be fortuitous in that they reduced the targeting error due to TOF uncertainties.

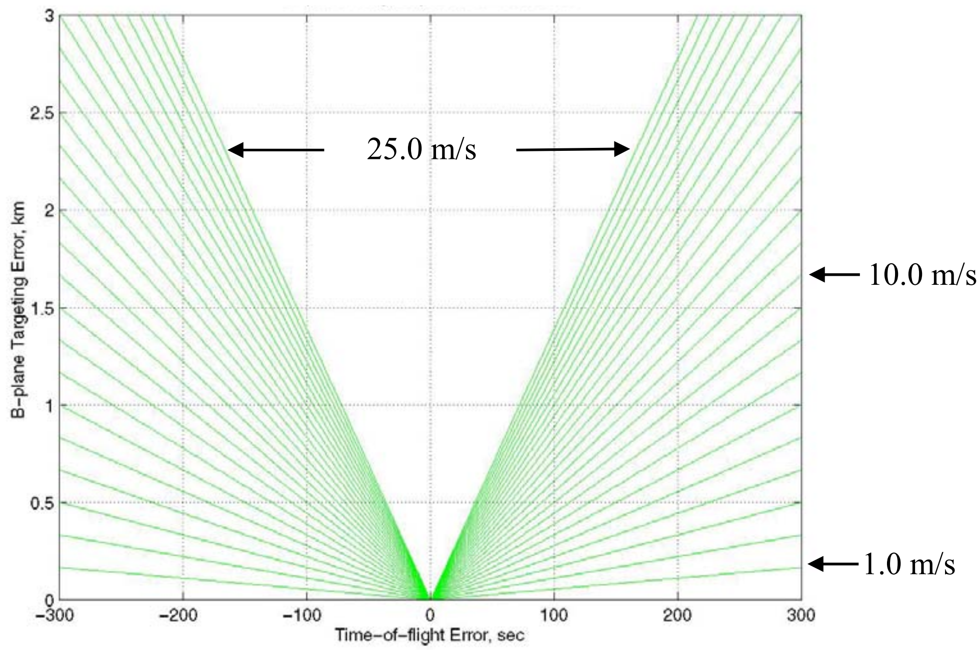

Figure 15 Targeting error due to TOF uncertainty for t_{SA_OBS} = 11 min before TOI_{nom} and t_{MAN} = 7.5 min before TOI_{nom}

IN-FLIGHT ADJUSTMENTS

In-flight Development of Cosmic Ray Model

In October 2004 (3 months before launch), we were completing a series of Mission Scenario Tests to demonstrate that both s/c could accomplish the nominal mission encounter. During the Impactor encounter scenario test, a parameter used to model the effects of cosmic rays in the images that were processed by AutoNav, was set incorrectly and resulted in an unrealistically large (length in pixels) cosmic ray, during Blobber image processing, that pulled the orbit solution off enough such that a loss-of-lock on the target body resulted when the transition was made to centroid box image processing. This resulted in failure to impact. Even though it could be shown using data from previous missions, that the modeled cosmic ray was unrealistic, this got everyone's attention and resulted in the following changes, which were good:

1. Blobber image processing would be extended to cover the entire first 60 min of AutoNav operations
2. The centroid box size would be increased to 400x400 pixels
3. The data weighting strategy used to de-weight bad observations would be closely scrutinized and ultimately shown to be the best defense against cosmic rays
4. The data editing strategy used to eliminate bad observations would be closely scrutinized
5. A cosmic ray model would be developed in-flight, using image data taken during cruise activities, for the purpose of monitoring and assessing the risk posed by cosmic rays and solar flare events

By mid-April (3 months after launch) a model for all three cameras had been developed and was being used for testing. By the end of May, it was demonstrated that the risks posed by cosmic rays were minimal, due in large part to the data weighting parameter settings within AutoNav: If a large, bright cosmic ray was ingested such that it had a significant influence on the CB observation, it would be de-weighted in the OD process thereby mitigating it's influence on the OD solution and subsequent targeting.

ITM Zig-Zag

After launch, the Project embarked on an intensive Robustness Test program designed to stress the Impactor and Flyby s/c during encounter operations. This consisted of pushing the limits of the following parameters in the software test bench simulations:

- Center of mass uncertainties
- Separation rates
- Uncertainties in initial position & velocity
- Instrument alignment errors
- ADCS component alignment errors and performance errors
- Nucleus size variations
- Nucleus brightness (albedo) variations (over/under exposure)
- Coma brightness variations
- Cosmic ray environment variations (frequency, length, brightness)

What came out of this effort was the discovery of a subtle interaction between ADCS attitude knowledge error (AKE) stability and AutoNav estimation of the comet-relative velocity. Although we understood the influence of AKE drift on targeting, particularly ITM-3, we did not consider its potential effect on ITM-1. The result was the ITM zig-zag in which AKE drifts lead to AutoNav velocity estimation errors that feed into the ITM-1 maneuver computation and execution. ITM-1 was originally placed at E-100 min to efficiently remove large (30 km) pre-release delivery errors. By the time ITM-2 comes along, at E-35 min, the same AKE drifts mapped to smaller spatial errors and thus smaller estimated velocity errors, so ITM-2 had to undo the trajectory error imparted by ITM-1. This zig-zag behavior was a concern for two reasons: 1) it wasted propellant that might be needed for ITM-3 at E-12.5 min; and 2) it would take the Impactor s/c off-course, in the presence of a perfect, pre-release delivery and increased the risk of failing to impact should we be unable to execute ITMs 2 and 3 for any reason. The response was to move ITM-1 a bit later, E-90 min, but not too late so as to give up the ability to efficiently remove a large pre-release delivery error and accept the waste of propellant. This placed the burden now on ADCS to perform within specification[7].

HRI Focus

In mid-March, a problem with the HRI focus was discovered. The Science Team began working to compensate using a post-processing method known as deconvolution of the point-spread-function (PSF)[8]. While that earned back the primary mission objective of high-resolution imaging of the fully developed crater, it would be no use to AutoNav operations. AutoNav relied on the use of the HRI for one primary function: Scene Analysis. The key to operating the two s/c independently and in parallel, such that they arrive at the same solution, depended on matching the resolution of the ITS Scene Analysis images and the HRI Scene Analysis images. The expected performance of the two instruments put the SA image time at E-11 min on the Impactor and at E-21 min on the Flyby. Now, the MRI would have to be used for Scene Analysis, which meant we would have to delay the MRI Scene Analysis images for as late as possible prior to impact and we would have to move the ITS Scene Analysis images back in time, so that they could be acquired a approximately the same range. The imaging sequences were modified to place the MRI SA images at E-4 min; a range of ~11,129 km and the ITS SA images at E-16.5 min; a range of ~10,197 km. Because of the TOF errors discussed in the previous section, the time of the maneuver execution was also moved back to E-12.5 min. For this selection the worst-case targeting error for a 10 m/s ITM-3 and a 120 sec TOF error was reduced to ~0.5 km.

Cometary Outburst

During the weeks leading up to encounter, the Optical Navigation (OpNav) Team at JPL embarked on an intensive imaging campaign. The OpNav data was combined with the radiometric tracking data to provide the best estimate of the trajectories of the Flyby s/c, pre-release, and comet Tempel 1. Those data revealed a previously unobserved phenomenon called "outburst", which is attributed to periodic illumination of active regions on the comet surface that resulted in significant outgassing[9]. Figure 16 shows two events where the integrated brightness increased by a factor of 2.

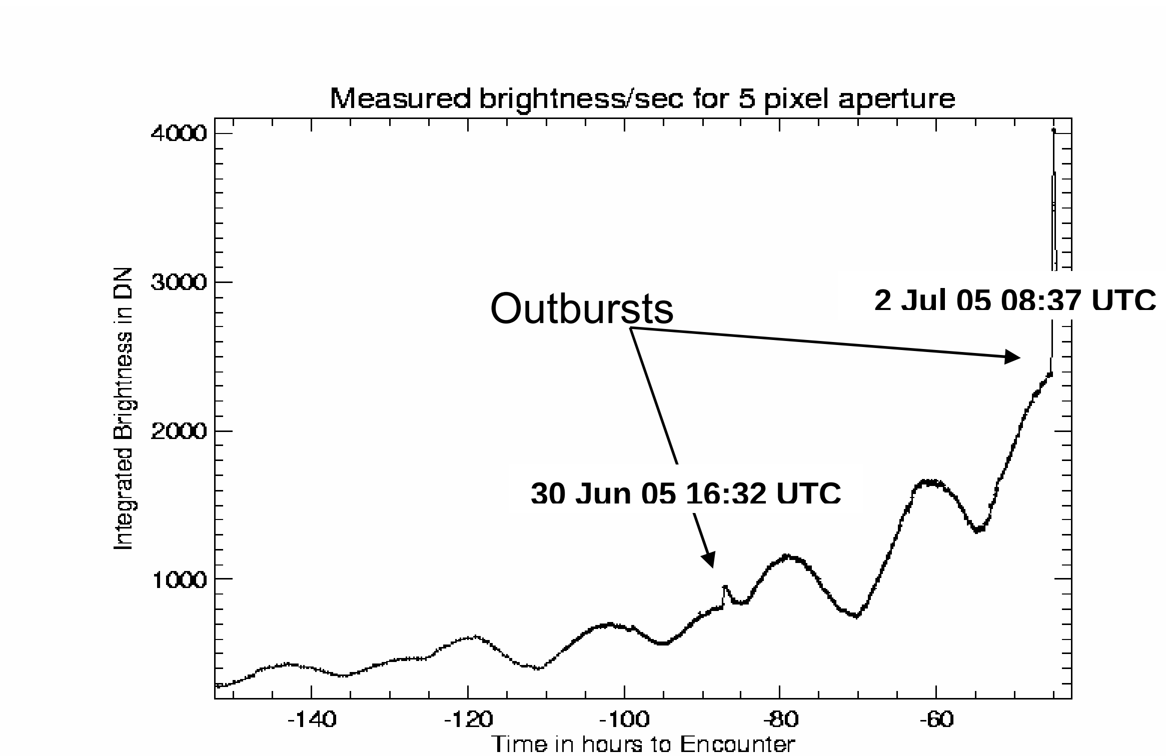

Figure 16 Outbursts observed OpNav data during approach to comet Tempel 1

The last outburst was observed at 22 hrs before release of the Impactor s/c. The rotational period between outbursts was determined to be ~ 40.8 hrs[9] meaning the next outburst would occur approximately 5 hrs before the expected TOI and 3 hrs before the start of AutoNav operations. An outburst during any 20 min OD arc during AutoNav operations was not expected to have a significant influence. Although the nucleus spanned 10s of pixels, it would depend on how collimated the jet was and how far from the nucleus surface the jet would extend. No outburst was observed in the final 5 hrs of encounter due to the reduced frequency of OpNav imaging. The persistence of the brightening due to the outbursts, was estimated to be on the order of just a few minutes for the rise in brightness followed by a dissipative period on the order of 1 hr. The outburst following the impact itself would be another matter that is described in the following section.

ENCOUNTER PERFORMANCE

The release of the Impactor s/c went very well. Both s/c saw low rates and 7 min later the Flyby s/c S-Band receiver locked-up on the S-Band signal from the Impactor s/c and telemetry was flowing. It was determined in the years before encounter that the worst-case separation event and rate-capture could result in linear accelerations that would perturb the trajectory by as much as 8 km. The post-release rate capture required RCS thrusters, which perturbed the pre-release trajectory by no more than 1 km; we had a chance.

Impactor Spacecraft AutoNav Targeting Performance

At the time of ITM-1 (E-90 min), the nucleus spanned ~10 pixels. AutoNav targeted the CB and computed a maneuver of 1.27 m/s. ITM-2, used for redundancy, also targeted the nucleus CB and at that time the nucleus spanned ~35 pixels. AutoNav computed a 2.26 m/s ITM-2. Finally, ITM-3 was computed and executed at E-12.5 min and based on a Scene Analysis offset of 1.85 km from the observed CB. The corresponding maneuver was computed to be 2.29 m/s. All told, the ITMs consumed a total of 5.82 m/s (23% of allocation). Figure 17 shows the pre-release and AutoNav targeting performance relative to nucleus as seen in an image taken at ~E-30 min. The pre-release delivery was superb. ITM-1, which was intended to efficiently remove large delivery errors (30 km) pulled the targeting away from the CB, however, ITM-2 brought the trajectory back in-line for intercept very near the observed CB. ITM-3 then applied the SA offset and Impact occurred at the desired location biased toward the Flyby s/c's point of closest approach.

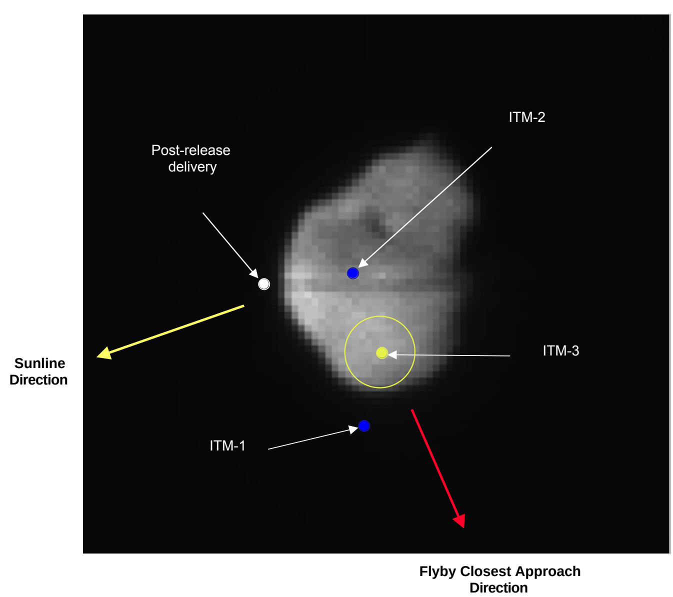

Figure 17 Performance of pre-release and AutoNav targeting during encounter with comet Tempel 1

Prior to impact, the Impactor s/c ADCS aligned the ITS with the AutoNav estimated comet-relative velocity vector to capture high-resolution images of the nucleus surface surrounding the impact site: the context imaging. Figure 18 shows a sequence of images, some full-frame, some sub-frame, taken with the ITS in the final hours of encounter. The Flyby s/c successfully transmitted the TOI offset to the Impactor s/c. The impact site, indicated by the arrow in Figure 18 was located relative to distinct features in the Flyby images showing the impact flash. The ITS pointing during the last 5 min resulted in the impact site being nearly centered in the ITS FOV as desired. The last image telemetered to the ground was taken 3.7 sec before impact. The range to the nucleus was < 40 km, which would give less than 40 cm/pixel resolution, the highest ever taken of the surface of a comet nucleus. Additional analysis being conducted by the Science Team, however, seems to indicate a distinct change in ITS performance in the last 30 sec[10] during which there were several attitude upsets as a result of particle impacts.

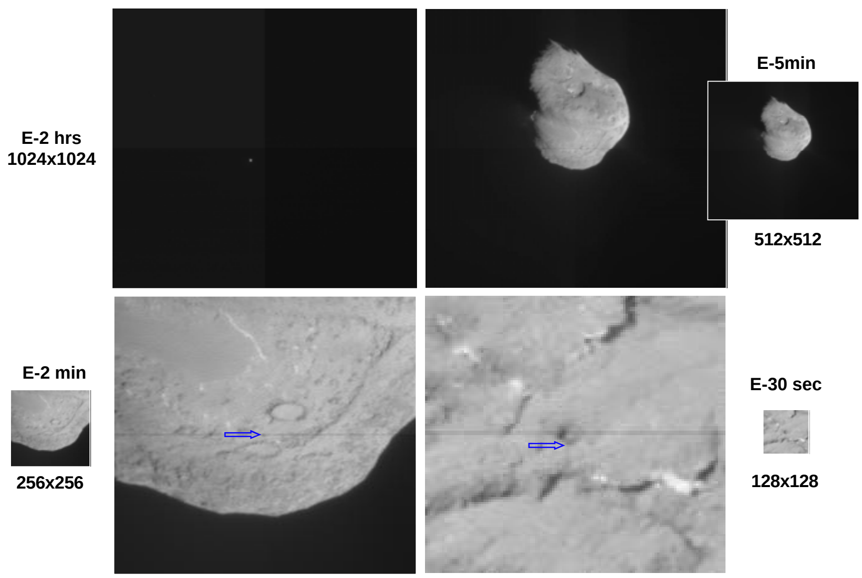

Figure 18 Pre-impact ITS images taken during the final 2 hrs prior to impact. Arrow indicates impact location

Flyby Spacecraft AutoNav Tracking & Timing Performance

The Flyby s/c combined AutoNav/ADCS performance at the two key epochs: TOI and TOFI can be summed up in two images. The flash shown in the center image of the bottom row in Figure 19 is evidence of two important performance indicators: 1) the image sequence timing updates provided by AutoNav were within the timing tolerance of 3 sec and allowed for high-temporal resolution of the impact event in the MRI and HRI instruments, and 2) the total pointing error (AutoNav tracking errors, ADCS pointing control errors, in-flight alignment errors) was just 61 μrad (31 pixels from the center of the HRI, which has a pixel scale of 2 μrad/pixel. The far right image in the bottom row clearly indicates that the impact site was out of the FOV of the HRI 11 sec before shield mode entry. Nonetheless, the impact site was captured in the first of the last 3

HRI images taken (not shown here) before shield mode due a 3-picture mosaic strategy that was implemented to account for possible AutoNav uncertainties at the end.

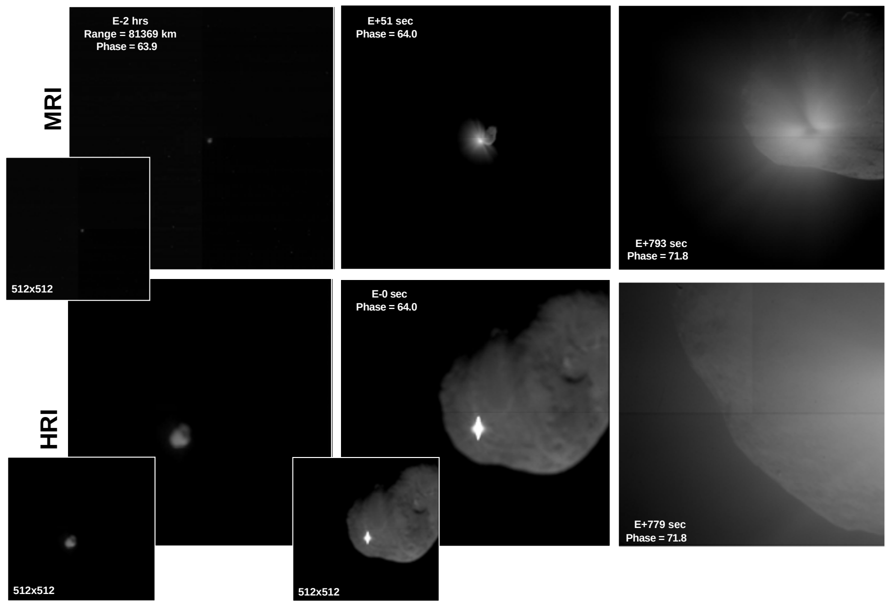

Figure 19 MRI (top row) & HRI (bottom row) images taken during the final 2 hrs of encounter with comet Tempel 1

Ejecta Plume Brightening

Figure 19 also illustrates the influence of the ejecta plume brightening on AutoNav post-impact tracking. The center image of the top row was taken ~51 sec after impact. The ejecta plume brightness is near saturation level in the MRI detector. The plume itself spans an area greater than the size of the nucleus. The pointing seen in the far right image of the bottom row of Figure 19 shows that the FOV is centered on a location near the limb of the nucleus and in the direction of influence due to the ejecta plume. Not only had the ejecta saturated the CCD, the brightness and extent of the plume persisted for the entire post-impact Flyby imaging campaign and obscured the view of the fully developed crater. No one could have predicted this outcome and it is not clear that we could have done anything to mitigate the effect. We had been concerned in the months leading up to encounter, but the Science Team could not provide us with any recommendations. In the end we built an MRI AutoNav imaging sequence that contained a mere 3 min data gap, post-impact, to allow for the plume brightening and dissipation. That was not sufficient, but in the end it did not matter and there is likely little we could have done to mitigate the influence, however, in hindsight analysis to try to bound the range of possibilities might have led to a solution that could have been implemented?

SUMMARY

The challenges for AutoNav on the Deep Impact Mission came from the unknown nature of the Tempel 1 nucleus; we had no way of knowing what to expect in advance and no previous s/c had flown past Tempel 1 to give a close-up look. The three cometary nuclei shown in Figure 20 have very little resemblance and highlight the fact that if you've seen one nucleus you definitely haven't seen them all. Although, characteristics of Borrelly were exaggerated and used to test the AutoNav system and Wild-2 images were used to assess the performance of Biased Scene Analysis.

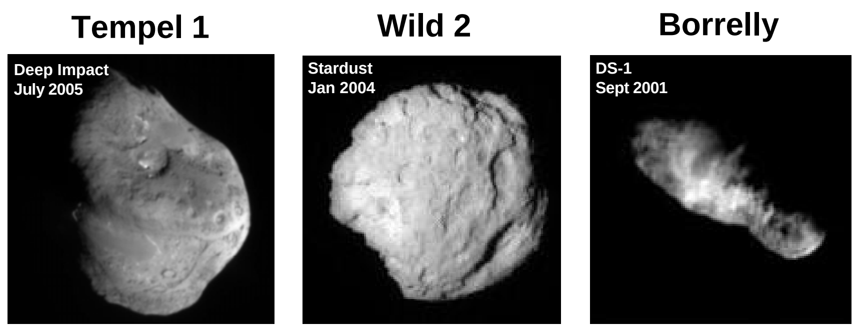

Figure 20 The three cometary nuclei observed in recent years using AutoNav

The AutoNav system performed extremely well on both spacecraft with all the objectives being met: 1) Impactor targeted and impacted in an illuminated area viewable from the Flyby s/c at the time-of-final-imaging; 2) Impactor s/c acquired and transmitted high-resolution images of the Impact site up to 3.7 sec before impact; 3) Flyby s/c AutoNav updated the TOI event sequence timing on both the Impact and Flyby and captured the impact event in both the MRI and HRI high-rate subframes; 4) Flyby s/c tracked the impact site for the entire post-impact observing period with the MRI; and 5) Flyby s/c AutoNav updated the shield-mode entry timing (final imaging timing) which allowed a high-resolution image of the impact site to be captured via use of a 3-image mosaic strategy.

In the end, all of those things that we were concerned about, and spent a great deal of effort to mitigate: cosmic rays, large-scale topography and significant self-shadowing, nucleus rotation, coma (coma jets), independent, parallel Scene Analysis operations, and time-of-flight errors, had little influence on the AutoNav performance. However, the one thing we were concerned about, but could do nothing about had the largest influence: eject plume brightening. Perhaps, this was the case only because we reduced or eliminated the effects of all the other concerns.

The challenges of autonomous targeting and tracking the unknown: comet Tempel 1, were overcome by exhaustive testing, open and independent review of the AutoNav system and the planned operations strategy, open dialog with the Science Team throughout the development phase, and in the end, a fairly benign presentation from Tempel 1.

AKNOWLEDGEMENTS

The authors would like to acknowledge Bob Gaskell, NASA-Jet Propulsion Laboratory, for his work on modeling cometary nuclei for simulation analysis. We would like to thank Miguel San Martin, NASA-Jet Propulsion Laboratory, for his effort and involvement to ensure the AutoNav sensitivity to ADCS performance on the Impactor s/c was completely understood. Finally, we thank Andrew Johnson, NASA-Jet Propulsion Laboratory, for his continually involvement in AutoNav issues and for developing the in-flight cosmic ray model for all of the Deep Impact instruments that allowed us to assess the influence of cosmic rays on Impactor targeting and

Flyby tracking performance. 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.

REFERENCES

1. Bhaskaran, S., Riedel, J. E., Synnott, S. P. (1996) "Autonomous Optical Navigation for Interplanetary Missions." *Space Sciencecraft Control and Tracking in the New Mellennium*, Proc. SPIE, pp. 32-34.
2. 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.
3. Phinney, A. (2001) "Mass on the Impactor." *BATC Systems Engineering Report DI-IMP-STR-004*.
4. Chesley, S. R. and Yeomans, D. K. (2006) "Comet 9P/Tempel 1 Ephemeris Development for the Deep Impact Mission." *16th AAS/AIAA Space Flight Mechanics Conference*, Tampa, FL.
5. Weidenschilling, S. J. (1997), *Icarus*, v. 127, p. 290.
6. Stooke, P. and Abergel, A. (2000), Halley Nucleus Shape Model, Personal Communication via M. Belton.
7. Hughes, M. and Schira, C. (2006) "Deep Impact Attitude Estimator Design and Flight Performance." AAS 06-084, *29th Annual AAS Guidance & Control Conference*, Breckenridge, CO.
8. Klaasen
9. A'Hearn, M. F., Belton, M. J. S., et al. (2005) "Deep Impact: Excavating Comet Tempel 1." *Science*, 10.1126, 8 September 2005.
10. Hampton