American Astronautical Society

AAS 01-022

Deep Impact : 19 Gigajoules Can Make Quite an Impression

Thomas Bank, William Frazier

Ball Aerospace & Technologies Corp.

William Blume, Daniel Kubitschek, George Null, Nick Mastrodemos, Stephen Synnott

Jet Propulsion Laboratory
California Institute of Technology

24th ANNUAL AAS GUIDANCE AND CONTROL CONFERENCE

January 31 - February 4, 2001
Breckenridge, Colorado

Sponsored by
Rocky Mountain Section

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

DEEP IMPACT: 19 GIGAJOULES CAN MAKE QUITE AN IMPRESSION

Tom Bank and William Frazier
Ball Aerospace & Technologies Corporation

William Blume, Daniel Kubitschek, Nick Mastrodemos, George Null and Stephen Synnott
Jet Propulsion Laboratory
California Institute of Technology

Deep Impact will impact the comet Tempel-1 on July 4, 2005, with a 350 kg smart impactor, at a relative velocity of over 10 km/s. The impact energy of roughly 19 gigajoules is expected to excavate a crater of approximately 20 m deep and 100 m wide. The impact event will be clearly visible from small telescopes on Earth, especially in the IR bands. When combined with observations taken from the Flyby spacecraft, this science data set will provide unique insight into the materials and structure within the comet (the underlying relatively aged surface), and the strength of the surface. Secondary observations include coma dust environment, optical properties, and nucleus morphology.

Deep Impact will use autonomous optical navigation to guide its Impactor to Tempel-1, while its Flyby spacecraft uses identical software to determine its encounter geometry and slew profile for flyby imaging. The impact crater will be viewed from the Flyby spacecraft, collecting IR and high-resolution visible images of the ejecta and fully developed crater. The autonomous optical navigation software is an extension of the AutoNav software developed by JPL for the Deep Space-1 mission. The two spacecraft fly out to the comet mated to each other, and separate one day away from encounter. Following separation, the Flyby performs a trajectory control maneuver to affect a 500 km closest approach at 16 minutes after impact by the Impactor.

SCIENCE OBJECTIVES

Deep Impact will provide key insights into the interior of comets previously unavailable from other missions. This will lead to insights into the development of our solar system, and a better understanding comets in general; some of humankind's most ancient puzzles.

Cometary Materials

Our knowledge of comets is dominated by a number of paradoxes. For example: Comets contain perhaps the most pristine, accessible material from the early solar system, but where is it in the nucleus? Comets appear to become dormant, but does the ice become exhausted, or is sublimation inhibited somehow? Which dormant comets are masquerading as asteroids? Coma gas observations are widely used to infer ices in protoplanetary disks, but what is the composition of the nucleus? Comet nuclei have been observed to break apart under small stresses, but is there strength at any scale?

The present state of knowledge of cometary nuclei size and albedo are derived almost entirely from observations of comet Halley, as shown in **Figure 1**.

Cometary nuclear surfaces are thought to be aged by multiple processes. Aging processes while in the outermost solar system (Oort cloud) are limited to cosmic rays and "warming" by passing stars and supernovae but just beyond Neptune they also include collisions and accretion of debris. Perhaps more importantly, near perihelion, the surface is changed by relatively rapid solar heating, which causes outgassing, ruptures from gas pressure, migration of volatile ices, thermal stress fractures, and venting. These processes cause the surface layers to be dominated by lag and rubble layers that obscure observation of the mantle and pristine materials underneath. Various models show the depth of these outer coatings to range from one to many tens of meters, as shown in **Figure 2**.

Cratering

Cratering is a very effective and relatively simple method of exposing the nucleus mantel and pristine materials for observation. Observation of the crater development process also yields additional information about the mechanical properties of the materials. Scaling from terrestrial craters and hypervelocity impact experiments provides models of the DI crater depth, which yields a baseline prediction of approximately 120 m wide by 25 m deep, and an excavation time of about 200 sec. Sample simulated crater images, as seen by the DI instruments, are shown in **Figure 3**. These images cover the extreme range of expected elevation angles, and also indicate the expected crater shape and shadowing effects. The instrument suite developed to produce these images is presented in a subsequent section.

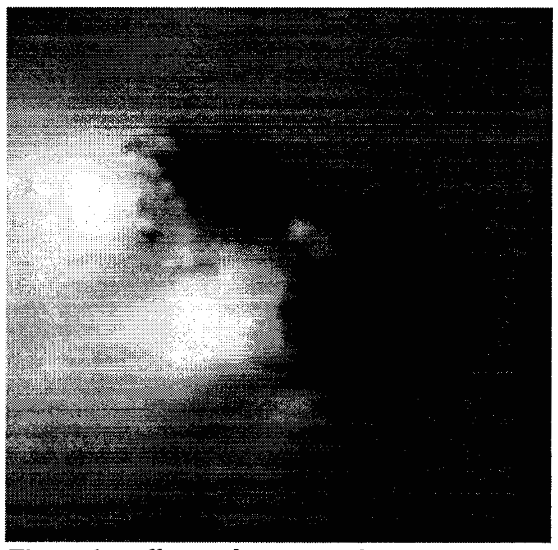

Figure 1. *Halley nucleus image from previous flyby provides basis for present knowledge of nuclei size and albedo.*

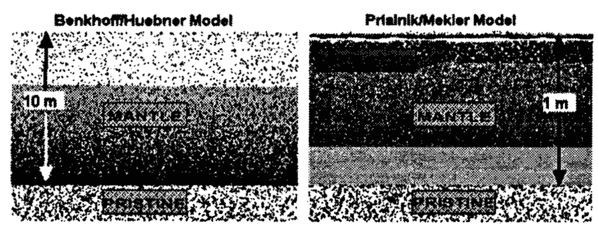

Figure 2. *Benkhoff/Huebner and Prialnik/Mekler comet surface models differ in the sign of the density gradient near the surface.*

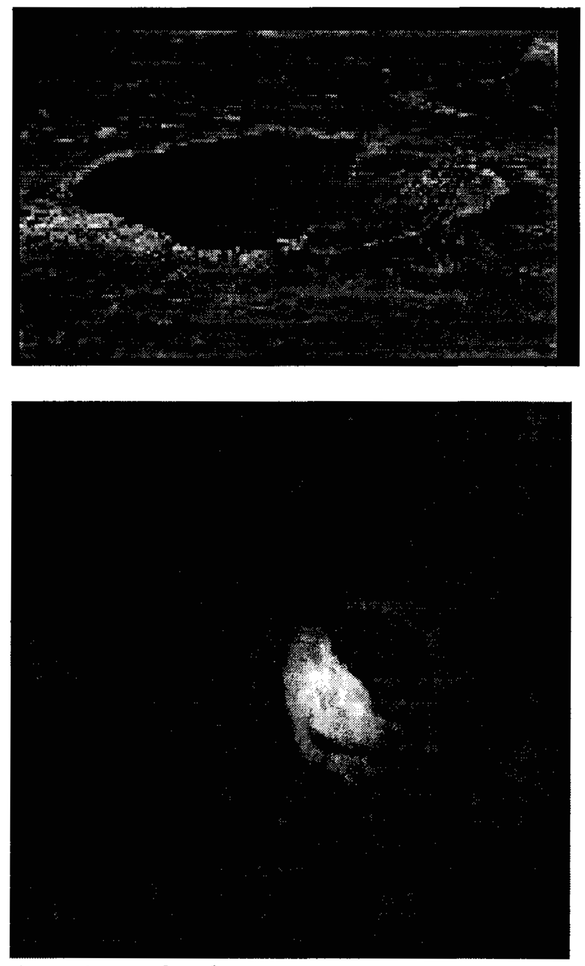

Figure 3. *Simulated crater images from extremes in expected range of elevation angles.*

Comet Environment Models

The very same unknowns that make comet exploration extremely rewarding, also make it technically challenging. The challenges include modeling the visible appearance of the nucleus, to aid in development of the autonomous-impacting navigation algorithms. The nucleus shape may be rather irregular due to accretion, which causes light and dark patches, depending on the solar phase angle. For visibility from the Flyby spacecraft, the Impactor must hit in a lighted area.

Ground-based observations of Tempel 1 have been made during the 2000 apparition using the UH 88-inch and Keck 10m telescopes to assist in characterizing the nucleus rotation period, albedo and phase function, and dust environment that DI will face during the next apparition in 2005. A visible image taken on Sept 9 at a range of 2.6 AU, 8 months after perihelion, is shown in **Figure 4**. This indicates a much dustier environment than previously expected, probably due to the presence of residual large dust particles ejected near perihelion. The current best estimate, based on a very preliminary analysis of the data from August 2000, is that the comet nucleus has dimensions of roughly 2.5 by 5 km.

Modeling of the dust particle size distribution is critical to the DI flight system design process, since it drives attitude control capabilities and shielding requirements. Curves of the currently-predicted dust flux are shown in **Figure 5**. The horizontal scale covers the time between Impactor impact, closest-approach by the Flyby spacecraft, and egress from the coma. The Impactor is expected to experience many dust collisions prior to hitting the nucleus, while the Flyby Spacecraft is expected to experience a relatively small number. Uncertainties in the data underlying these

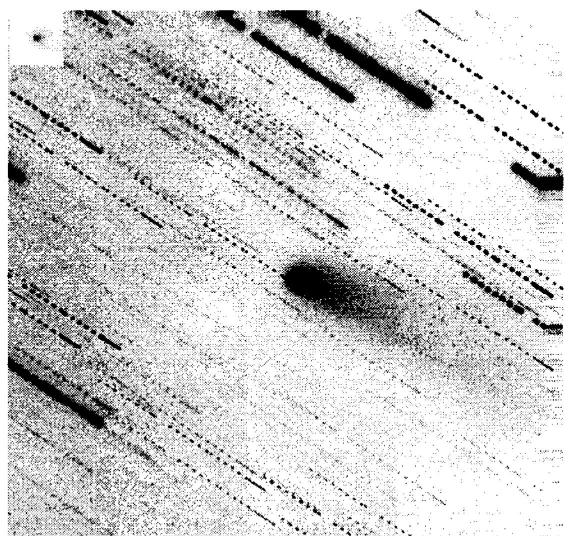

Figure 4. *Recent visible image from University of Hawaii 88-inch telescope shows high dust content.*

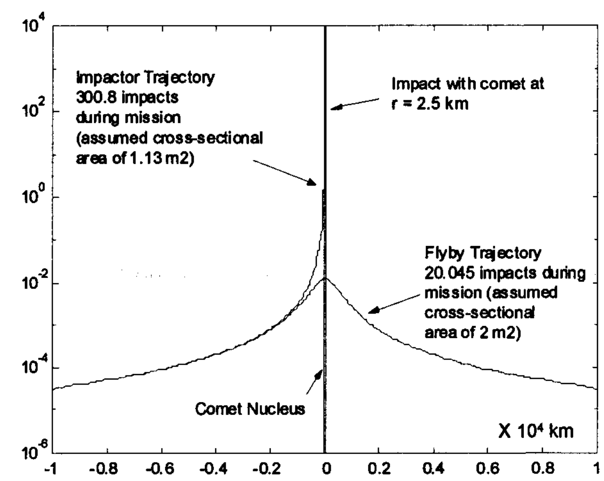

Figure 5. *Expected dust flux profile is highly nonlinear due to inverse-square density model (**TBR**).*

curves, and their associated statistical probabilities, create a range of flux that covers an order of magnitude. High-fidelity performance simulations of the flight system in this range of environments shows that the Flyby Spacecraft has a good probability of maintaining high-quality pointing control throughout the encounter, whereas the Impactor attitude control may be lost shortly prior to impact.

MISSION DESIGN

Earth-to-Earth Cruise Phase

The complete mission trajectory is shown in **Figure 6**. The Earth-to-Earth cruise phase provides over a year to fully characterize, calibrate, and test the FS. A swing-by of the Earth/moon system will occur in January 2005, allowing for calibration and test of the encounter software and imaging instrumentation.

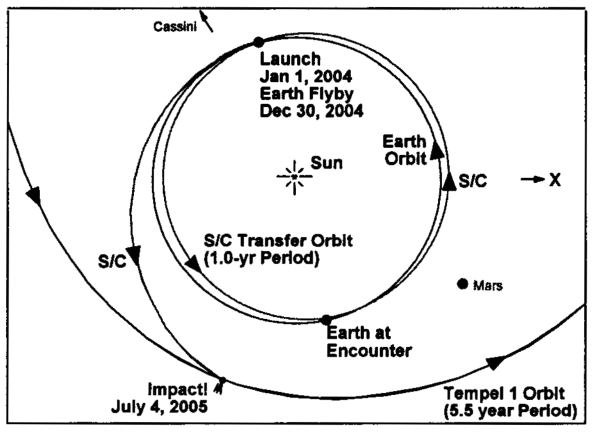

Figure 6. Mission trajectory includes launch, Earth-Earth Cruise, and encounter phases (TBR).

Encounter Phase

The encounter phase includes optical navigation prior to Impactor separation. Following separation, the Flyby spacecraft will slow itself relative to the Impactor by 106 m/s, which also includes a small cross-track component to provide the required 500-km flyby distance. The comet environment (primarily albedo and jets) will then be characterized by high-rate optical imagery downlinked in near-real-time, processed on the ground, and if necessary, uplinked to the Flyby Spacecraft and cross-linked to the Impactor. At the time of impact, the range to the comet from the Flyby will be approximately 8,700 km. The Flyby spacecraft instruments observe the impact event (crater and ejecta) temporally, spatially and spectrally. The long range at impact provides 16 minutes of imaging time, which provides a 200% margin over the predicted crater development time. At the end of the imaging sequence, the Flyby Spacecraft will have pitched 45 deg, and then be in a "shield-mode" attitude to enter the higher density dust region and for crossing the more hazardous orbital plane, as shown in **Figure 7**.

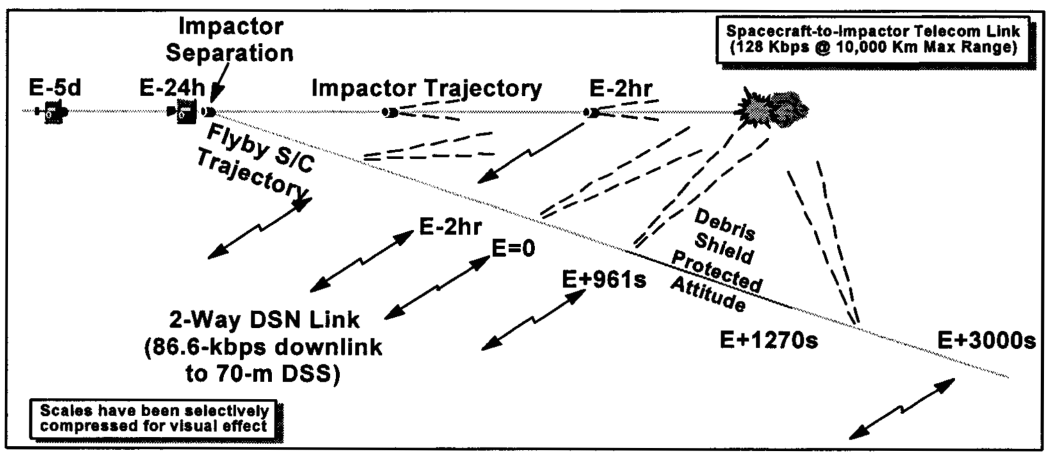

Figure 7. Encounter design supports imaging requirements with acceptable hazard to Flyby Spacecraft. (TBR)

FLIGHT SYSTEM

The DI Flight System is composed of the Instruments, the Impactor, and the Flyby Spacecraft.

Instruments

There are 3 primary instruments, two of which are shown in **Figure 8** and are accommodated by the Flyby Spacecraft. The High Resolution Instrument (HRI) uses a 30 cm aperture to support a Full Width Half-Max (FWHM) performance of 3.4m at closest approach. The visible CCD response spans 0.3 to 0.95 μm imaging, while the IR spectrometer spans 1 to 4.8 μm. The Medium-Resolution Instrument (MRI) design is similar to the HRI, although at 5 times lower spatial resolution, and supports optical navigation and provides functional redundancy to the HRI. The Impactor carries the third instrument, the Impactor Targeting Sensor (ITS), which to reduce cost and risk, is nearly identical to the MRI.

Impactor

An exploded view of the Impactor configuration is shown in **Figure 9**. It is designed to nestle within the Flyby spacecraft, and also carry the launch loads into the LV adapter. The Impactor will use the ITS and advanced JPL software to autonomously perform any course corrections required to assure impact in a lighted area. A cross-link capability is provided to transmit close-up images of the comet surface prior to impact, and also provides contingency commanding to the Impactor.

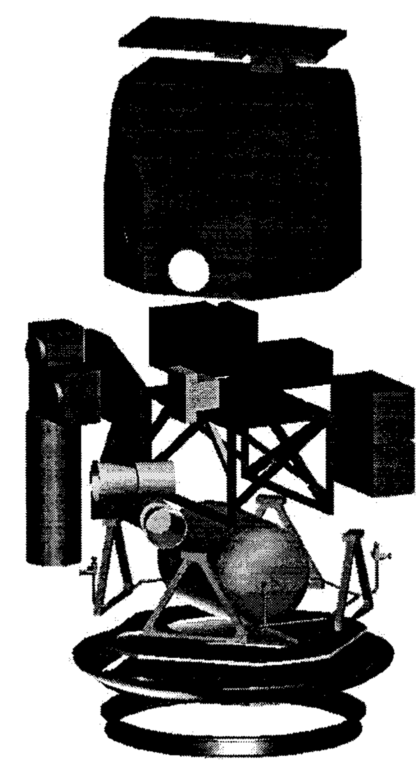

Figure 9*. Impactor exploded view shows FS interfaces.*

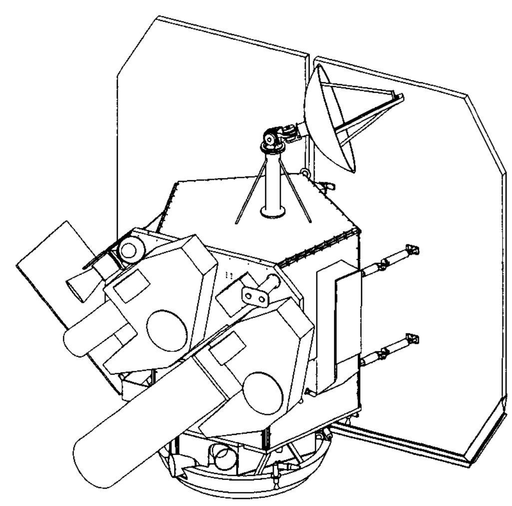

Figure 8*. Flight system cruise configuration.*

Flyby Spacecraft

The Flyby spacecraft configuration with solar arrays deployed during cruise, and including the Impactor inside, is shown in **Figure 8**. The instrument assembly can be seen mounted to the side of Flyby Spacecraft. Shielding is added to what appears to the "bottom" side of the spacecraft in this view, to survive the coma passage following closest approach and the end of imaging (this accounts for the 45-deg rotation of the instrument boresights relative to the vehicle figure axis). The Flyby spacecraft is entirely redundant, and features a very-high throughput RAD750 CPU and 1553 data bus-based avionics architecture, and a high-stability pointing control system.

Navigation Functionality

The Deep Impact navigation functions are based on radio measurements during cruise, until the comet nucleus can be resolved by the HRI at 30 to 60 days prior to encounter. After this time, the spacecraft position will be determined relative to the comet to increasingly greater accuracy using ground-based optical navigation techniques. The inertial Tempel 1 ephemeris will be estimated from ground-based observations of the comet, and interpretation of the comet against the star background observed by the HRI. During the encounter sequence, ground-based ephemeris will be loaded into the Autonomous Navigation flight software, and the remaining sequence will proceed automatically. This includes separation of the Impactor from the Flyby spacecraft, and the large "divert" burn by the Flyby. AutoNav software in the Impactor will then use optical navigation techniques to guide it to impact the Nucleus. Simultaneously, AutoNav software on the Flyby will use MRI measurements of the nucleus to provide desired pointing state information to the Flyby Attitude Determination and Control System (ADCS), to support imaging of the Impact event, crater evolution, and high-resolution crater imaging prior to closest approach.

IMPACTOR NAVIGATION

The conceptual approach to navigation and targeting of the Impactor spacecraft (s/c) consists of trajectory estimation based on optical observations of the center of brightness (CB), trajectory control maneuver (TCM) computation and execution to impact the nucleus in a desired (lighted) location, and impact site prediction following the last TCM for cross-link and Flyby s/c pointing control. The predicted impact location is important for crater imaging on the Flyby s/c since control errors preclude the Impactor from impacting at exactly the desired location. Pointing corrections based on cross-link information from the Impactor will be discussed in later sections. The AutoNav software system[2,3], developed at the Jet Propulsion Laboratory, California Institute of Technology, will be used to autonomously carryout these functions after the Flyby releases the Impactor, the tip-off rates are nulled, and the comet is acquired in the ITS. The baseline desired impact location is at the CB, however, the CB is not always lighted as will be shown in the next section.

Nominally, three crosstrack TCMs are scheduled during the last 100 minutes prior to impact. Each TCM is based on position and velocity estimated from a batch of optical observations (images processed every 15 seconds) and propagated to the planned time of maneuver execution. **Figure 10** shows a schematic of the Impactor control strategy. Table 1 lists the control sequence during the terminal phase. The large range of ΔV values shown for the last maneuver is a result of two things: 1) the selection of comet nucleus model and 2) the ΔV required to affect the same B-plane intercept location change increases as the distance to the B-plane decreases; the B-plane is defined as the plane containing the target body (Tempel-1 nucleus) that is normal to the spacecraft incoming asymptote. The two nucleus models under consideration will be discussed in the following section. **Figure 11** shows a set of three simulated images each corresponding to the last image processed in each of the three batch solutions prior to TCM execution.

Table 1
IMPACTOR CONTROL SEQUENCE

Time of Maneuver Execution (min)	*Image Processing/OD Time Before Impact (min)*	*ΔV (m/s)*
	100.0 ↓ 90.0	

I-88		2 – 7
	60.0 ↓ 50.0	
I-48		1 - 6
	44.0 ↓ 16.0	
I-15		0.5 - 11

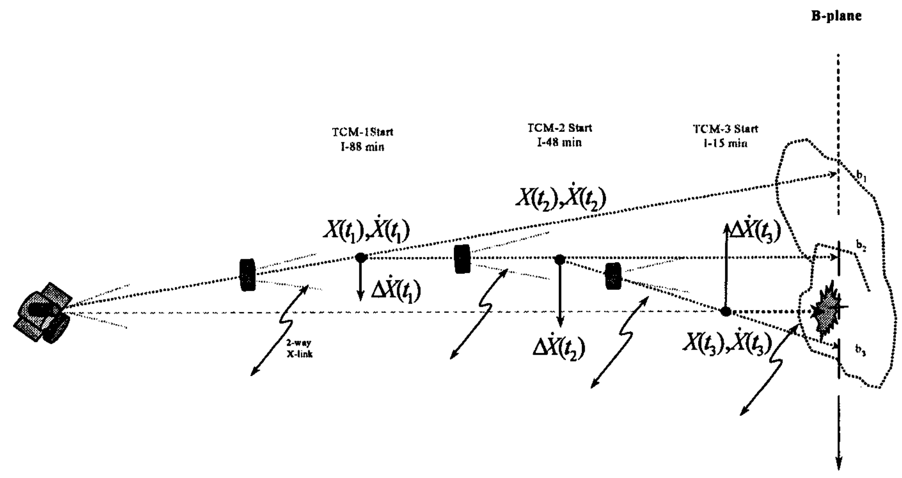

Figure 10. *Impactor control sequence schematic*

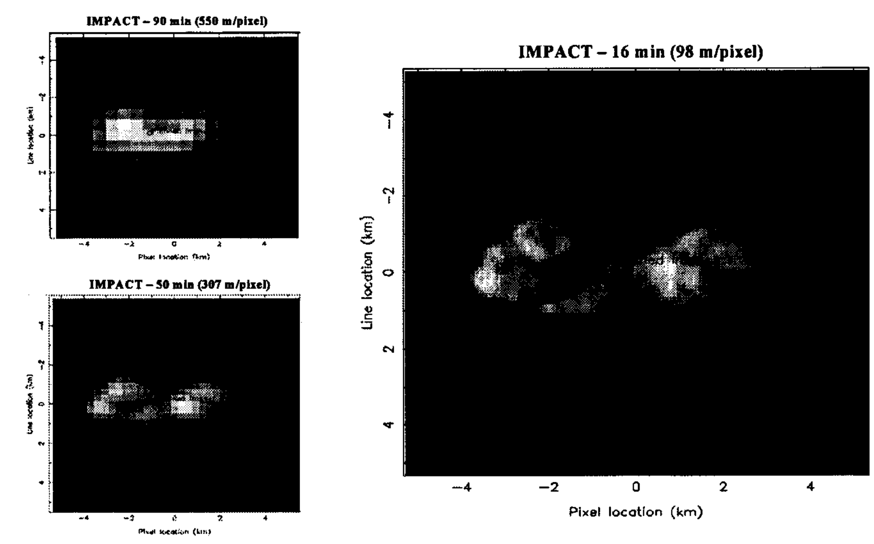

Figure 11. *Nucleus images corresponding to each orbit determination batch prior to TCM execution.*

Comet Nucleus Models

Two models have been defined by the Deep Impact Science Team and are currently being used to develop and test navigation algorithms. The baseline model was developed by Stooke and Abergel and based on Giotto and Vega s/c images of the comet Halley nucleus; this model is here referred to as the Halley-Stooke baseline model. The second model is based on accretion[1] and the worst-case expected axial ratio ranging from 3:1 to 5:1; this model is here referred to as the worst-likely case (WLC) comet model. During simulation, the rotation period of the nucleus is taken to be ~ 42 hours as defined by the Science Team and derived from Earth-based observations. Figures 12 and 13 show simulated images (software and analysis by Bob Gaskell, JPL) of the baseline and WLC comet nucleus models, respectively, 1 minute prior to impact.

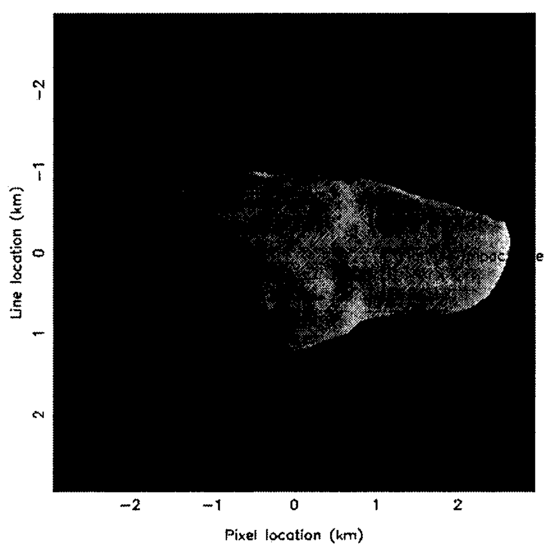

Figure 12. ***Baseline Halley-Stooke model***

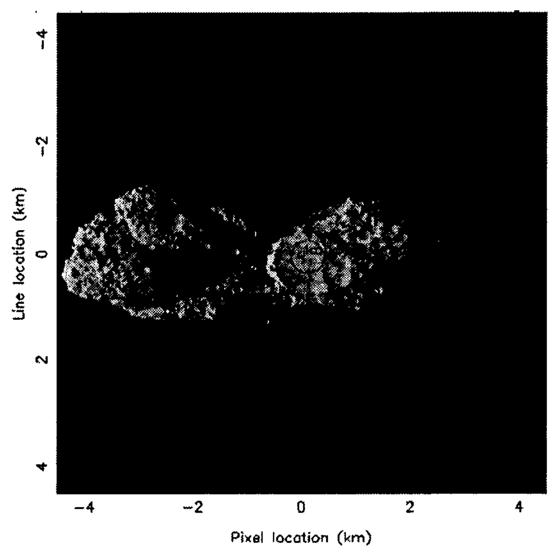

Figure 13. ***WLC comet model***

The approach here is to design targeting algorithms using the baseline Halley-Stooke comet model and test robustness of algorithms using the WLC comet model. Figures 12 and 13 illustrate an important point: the CB likely will not be lighted should the WLC comet model be encountered.

Scene Analysis

The alternative to targeting the CB is an algorithm based on scene analysis of the last image prior to computation and execution of the final TCM before impact. Scene analysis consists of the following steps:

1. Apply brightness discrimination cutoff to remove off-nucleus coma
2. Compute the center of brightness
3. Step thru pixels from outside of image toward center of image, marking each pixel below cutoff as off nucleus until the first pixel above cutoff is encountered
4. Locate on-nucleus dark pixels
5. Step thru on-nucleus pixels above cutoff and no less than the 3σ control error from the nucleus boundary
6. Compute the total number of lit pixels in the circular area of radius equal to the 3σ control error mapped to the nucleus surface surrounding each on-nucleus lit pixel

7. Select the pixel that has the largest number of lit pixels within the surrounding circular area of radius equal to the 3σ control error mapped to the nucleus surface
8. Compute the offset relative to the CB and use for targeting

The control error (3σ) mapped to the comet nucleus is represented by a circle of ~ 300 – 400 m in radius at the surface of the nucleus. These control errors consist of both pointing errors and the unmodeled variations in the CB relative to the nucleus center of mass due to the lighting phase angle and nucleus rotation, which cannot be modeled apriori. This procedure gives the Impactor additional capability and may prove essential to satisfying the requirement that the impact occur in a lighted area.

Algorithm Performance

To evaluate the performance of the targeting strategy and navigation algorithms a Monte-carlo analysis was performed. This analysis simulates the impact navigation sequence and evaluates the impact site selection and outcome. Images are simulated and processed in a batch-sequential filter to estimate the trajectory based on observation of the nucleus CB. The equations of motion are integrated to B-plane intercept. The intercept location is used to compute the lateral maneuver required for the desired B-plane intercept. The simulation computes and generates a "scorekeeping" image at I-30 sec to assess whether or not there is a lit impact. **Figure 14** shows a typical scorekeeping image for the WLC nucleus model. A successful impact is defined by hitting in a region with > 80% of the pixels in the impact site area (based on the assumed crater diameter of 100 m) above the brightness cutoff. Impact runs were performed for several different nucleus pole directions and 72 initial nucleus orientations for each pole direction. Tables 2 and 3 show the impact statistics for both the baseline model and the WLC nucleus model.

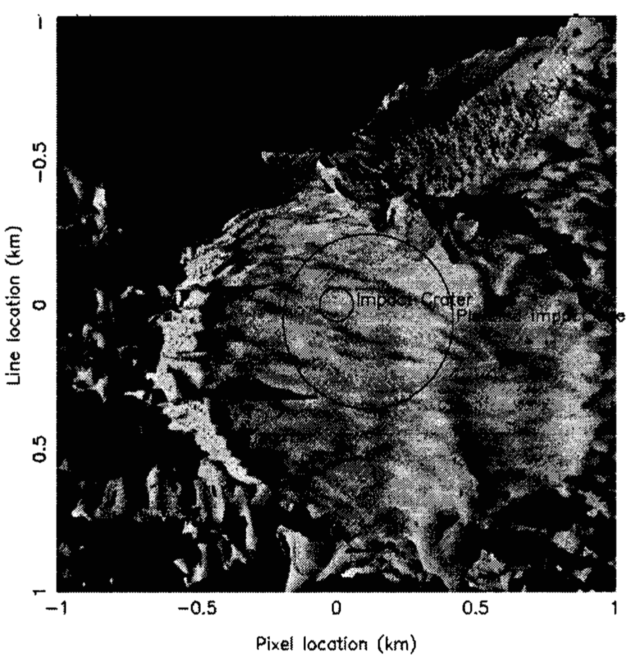

***Figure 14**. Impact scorekeeping frame*

Table 2
IMPACT STATISTICS FOR BASELINE MODEL
Results from 7 pole directions at 72 orientations (504 total) for each case

Targeting Mode	*Time of last image (min)*	*Lit Impacts (%)*	*Dark Impacts (%)*	*Misses (%)*
CB	16	100	0	0
	9	100	0	0
Scene Analysis	16	100	0	0
	9	100	0	0

Table 3
IMPACT STATISTICS FOR WLC MODEL
Results from 7 pole directions at 72 orientations (504 total) for each case

Targeting Mode	*Time of last image (min)*	*Lit Impacts (%)*	*Dark Impacts (%)*	*Misses (%)*
CB	16	63	37	0
	11	65	35	0
Scene Analysis	16	97	2	1
	11	99	0.8	0.2

The results shown in Tables 2 and 3 indicate that the probability of impact in a lit area for the baseline model is 100%. The probability of impact in a lit area for the WLC comet model ranges between 97 - 99% when scene analysis is used. Although scene analysis improves the results over CB targeting by 34%, it introduces the possibility of missing the comet nucleus altogether. Further analysis is being done to eliminate the possibility of a miss. The trade-off will be an increased number of dark impacts, which is the only acceptable alternative to impact in a lit area.

FLYBY NAVIGATION

Navigation of the Flyby s/c plays two important roles: 1) release of the Impactor on a collision course, and 2) computation of information necessary to control pointing of the imaging instruments during encounter for comet nucleus tracking and crater imaging. After release of the Impactor, the Flyby s/c will perform a divert maneuver that is designed to slow the s/c and control the close approach distance of 500 km to ± 50 km. Additionally, this maneuver provides time separation between the impact event and the time of last imaging, which will occur approximately 50 sec before time of closest approach (TCA). During the encounter sequence, the autonomous navigation (AutoNav) software will be processing images every 15 sec and performing an orbit determination batch solution once every minute. The equations of motion are then integrated through TCA using the updated initial conditions. The resulting ephemeris file is used, along with the target body ephemeris file, to compute a time series of the inertial pointing direction vector beyond TCA. Each time the comet-relative trajectory solution is updated, the inertial pointing time series is updated. The pointing direction vector is then adjusted using cross-link information from the Impactor spacecraft to provide inertial pointing to the estimated crater location.

Flyby Critical Sequence Pointing

After releasing the Impactor the Flyby executes a TCM to affect a 500 km nucleus fly by 17 minutes after impact. This provides the Flyby with a good viewing position for the impact and crater development.

After the TCM, the Flyby acquires the comet nucleus in the field of view of its two co-aligned science instruments, the High and Medium Resolution Imagers (HRI and MRI.) At separation the nucleus range is 881,280 km and the comet nucleus subtends 6 micro-radians or three pixels in the HRI and less than one in the MRI. For the next 24 hours the Flyby spacecraft maintains the comet nucleus in the science instrument FOVs collecting image data according to a stored command sequence.

At impact, 24 hours after release, the nucleus range is 8700 km and the comet nucleus subtends 600 micro-radians or about one third the HRI's field of view. A burst of high rate (10 Hz) imaging on the Flyby captures the impact plume and crater development. At 600 seconds after impact the range drops to 2550 km and the nucleus fills the HRI FOV. At 800 seconds after impact imaging ceases and the flyby spacecraft enters debris shield mode. The Flyby's closest approach to the nucleus is 500 km, at 850 seconds after impact.

AutoNav, the on-board optical navigation software, processes a science instrument image every 15 seconds providing the ADCS with predicted pointing commands to target the nucleus center of brightness and estimated time of impact.

Mosaic Imaging versus Direct Crater Targeting

The Deep Impact mission design relies on AutoNav pointing commands that directly target the impact crater. An alternative approach relying on a mosaic maneuver was considered and ultimately rejected on the following grounds. The primary science requirements call out crater imaging with 3.4 meter resolution. Analysis of the engagement geometry, the Flyby flight system and science instruments show 3.4 meter resolution images are attainable using the HRI only during the last 50 seconds of imaging.

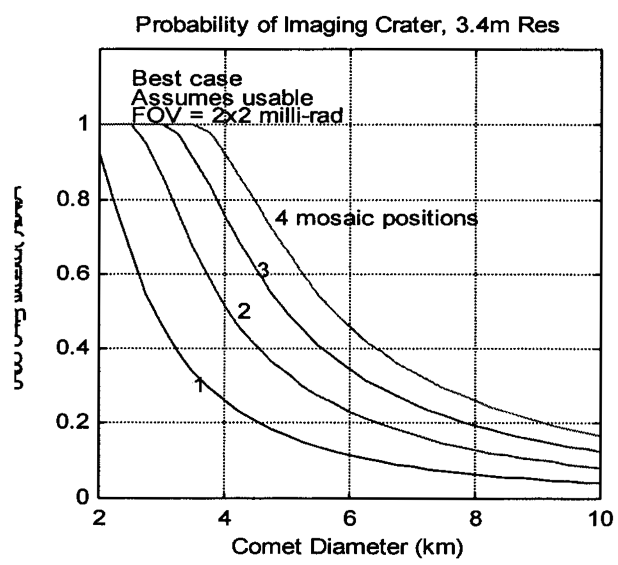

Figure 15. The available imaging time and ADCS authority limit the number of mosaic positions to 3, yielding unacceptable crater imaging probability for the expected 5km nucleus.

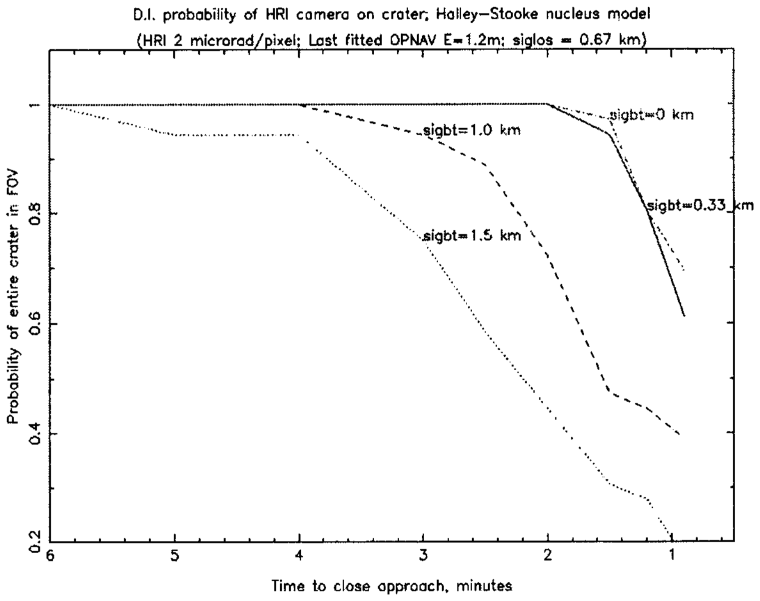

Figure 16. HRI probability of imaging crater in the presence of navigation errors, attitude errors, and uncertainties in the dimensions of the comet nucleus

ADCS simulation predicts 13 seconds to maneuver the spacecraft attitude one HRI FOV. Adding another 2 seconds for imaging each mosaic position requires 15 seconds allowing for no more than 3 mosaic positions and limiting imaging coverage to an unacceptable 50% of the nucleus, as shown in Figure 15. Also, images acquired during the 13 second slew are unacceptably smeared. The mosaic maneuver is unable to provide a comfortably high probability of imaging the crater and it was ruled out in favor of direct targeting. Therefore, crater imaging will be performed by receiving a predicted impact point from Impactor telemetry, and propagating that point forward on the Flyby spacecraft during the imaging sequence. The resulting probability of imaging the crater using this open-loop process is shown in Figure 16. The results shown in Figure 16 are based on a Monte-carlo simulation of the Flyby navigation and instrument pointing. Uncertainties in the downtrack position of the Flyby s/c, knowledge of the actual impact location (transmitted by the Impactor to the Flyby), errors associated with ADCS attitude knowledge and stability, and uncertainties in the comet nucleus dimensions combine to give the total pointing error. This error can be expressed in terms of the probability of maintaining the crater in the HRI FOV during the imaging sequence. In figure 16, *sigbt* represents the uncertainty in the crater plane-of-sky position. This uncertainty represents the error in the Impactor pre-impact location estimates. The larger the uncertainty, the lower the probability of having the crater in the HRI FOV at the time of last imaging and the sooner the possibility of losing the crater in the HRI FOV. Even in the case of perfect impact location knowledge from the Impactor s/c, the probability of keeping the crater in the HRI FOV all the way into the final imaging time is less than 100% since the navigation errors cannot be completely removed and since the projected comet nucleus dimension (third dimension) cannot be determined autonomously.

Attitude Determination and Control System

The ADCS design on the Flyby and Impactor are similar, only the hardware suites and internal gains distinguish them. Figure 17 shows the Flyby and Impactor ADCS block diagram, sensors and actuators.

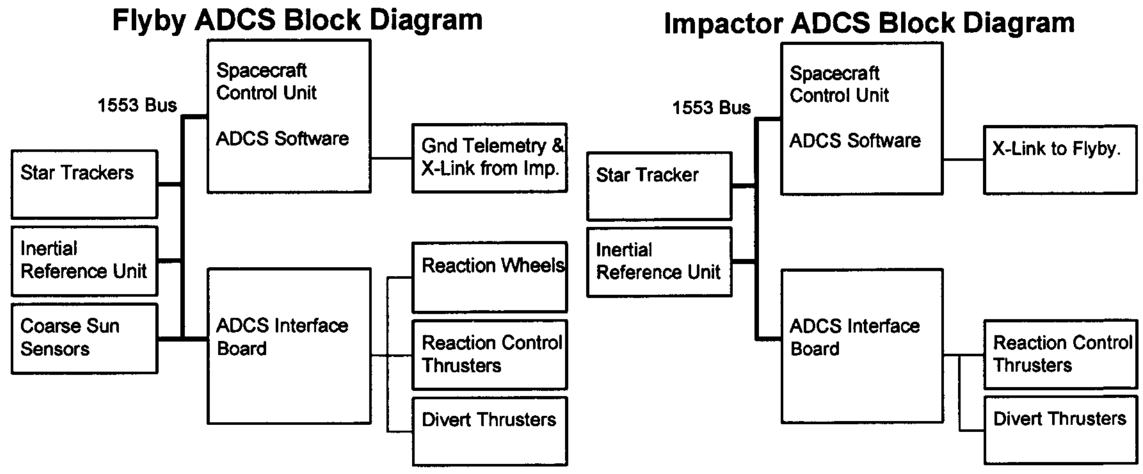

Figure 17*. The Flyby and Impactor ADCS block diagram showing the major sensors and actuators interfacing to the Spacecraft Control Unit where the ADCS software resides.*

Both systems rely on an IRU and Star Tracker combination for attitude estimation. The primary Flyby attitude actuators are four reaction wheels. Four 4-Newton reaction control thrusters are commanded by a background ADCS task to manage momentum. On the Impactor, attitude is maintained using four 1-Newton reaction control thrusters only.

For TCMs the Flyby and Impactor carry four 22-Newton divert thrusters. During TCM maneuvers, the divert thrusters are modulated to account for thruster misalignments and maintain spacecraft attitude. Figure 18 illustrates the ADCS flight software used on the Flyby and Impactor spacecraft. The ADCS software executes on the Spacecraft Control Unit with a cycle time of 0.1 seconds. The software is configurable by table uploads.

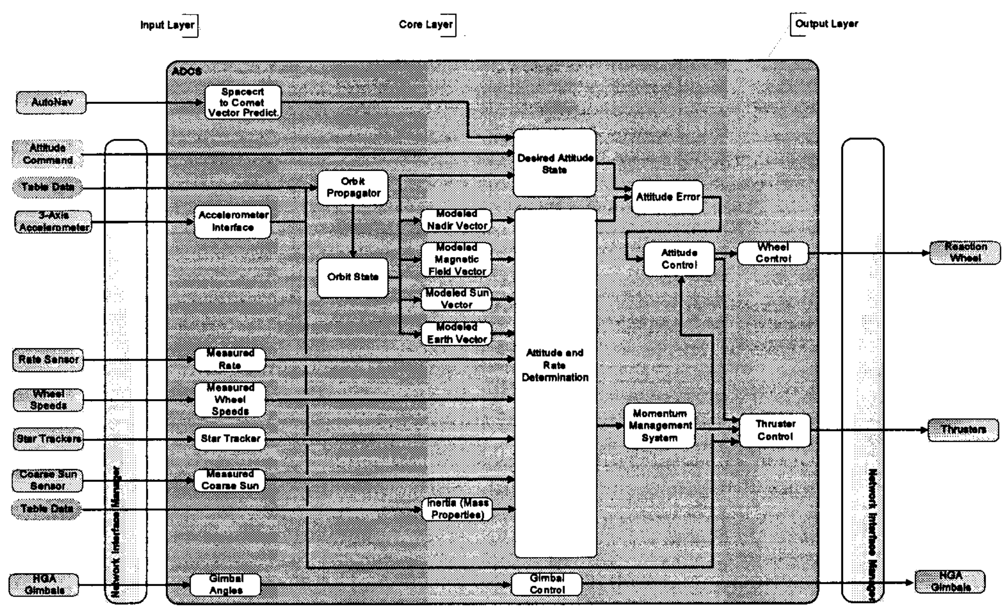

Figure 18*. Deep Impact ADCS flight software overview showing interfaces to sensors and actuators.*

The ADCS states summarized in Table 4 support all phases of the Deep Impact mission. The point/cruise state periodically rotates the spacecraft about the sun-line to avoid momentum build up due to solar pressure misalignment with the center of mass. The point/image performance requirements are closely coupled with imaging requirements and are the most demanding.

Table 4. Deep Impact ADCS States

ADCS State	Comment
Initialize	Initialize ADCS table parameters in preparation for operation.
Detumble	Null body rates using reaction control jets.
Acquire Sun	Point a selected body vector (specified in an ADCS parameter table) to the sun using sun sensors only.
Point/Cruise	Point a selected body vector to the sun and perform periodic about sun-line slews for momentum management.
Point/TCM	Point the divert thruster body vector in preparation for a trajectory change maneuver.
Point/Image	Point science instruments to comet nucleus while maximizing solar panel exposure.

ADCS Performance

The Deep Impact ADCS performance requirements vary with mission phase. During the 18 month cruise phase from Earth to comet Tempel-1, 3–axis pointing requirements for the mated Flyby and Impactor are driven by high gain antenna pointing requirements (1 degree, 3 sigma). Eight TCMs are scheduled over the 18 month cruise phase. Absolute 2-axis inertial pointing requirements during TCMs are 1.2 milli-radians, 3-sigma. Pointing requirements during the final engagement are tightest when DI ADCS and imaging requirements are tightly coupled. The science instruments are hard-mounted to the spacecraft bus, pointing and line-of-sight stabilization is accomplished using the ADCS. The Flyby ADCS pointing knowledge budget is 100 micro-radians, 3 sigma and the pointing stability (accomplished using the reaction wheels) is 1.5 micro-radians/sec, 3 sigma. The Impactor pointing knowledge is the same as the Flyby, 100 micro-radians, 3 sigma, the pointing stability achieved with the reaction control thrusters is 20 micro-radians/sec 3-sec.

ACKNOWLEDGEMENTS

The Deep Impact Project is funded by NASA under the Discovery Mission Program and managed at the Jet Propulsion Laboratory, California Institute of Technology.

REFERENCES

[1] Weidenshilling, Icarus Vol. 127, p. 290-306.

[2] Bhaskaran, S., et al. (1998). "Orbit Determination Performance Evaluation of the DEEP SPACE 1 Autonomous Navigation System." *AAS/AIAA Space Flight Mechanics Meeting*, Monterey, California, 9-11 February.

[3] Riedel, J.E., et al. (1997). "Navigation for the New Millennium: Autonomous Navigation for DEEP SPACE 1." *Proceedings of the 12th International Symposium on 'Space Flight Dynamics'*, ESOC, Darmstadt, Germany, 2-6 June.