Impactor Spacecraft Encounter Sequence Design for the Deep Impact Mission

Daniel G. Kubitschek

NASA – Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

ABSTRACT

On July 4, 2005, another first in space exploration was achieved. NASA's Deep Impact spacecraft (s/c) released a small, 350 kg Impactor s/c designed to target comet Tempel 1, estimated to be 14 km x 5 km x 5 km in size at the time of release. With a closing speed of approximately 10.3 km/s, the Impactor s/c autonomously guided itself to impact and captured 40 cm resolution images, the highest resolution images ever of the surface of a cometary nucleus, just moments before the collision. The objective of the Impactor s/c was to impact in an illuminated area viewable from the Flyby s/c. This paper describes the Impactor encounter sequence design, execution and contingency planning that contributed to the successful outcome in which all objectives were met.

1.0 Introduction

On July 3, 2005, NASA's Deep Impact Flyby spacecraft (s/c) released a small, 350 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 the Impactor 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 with 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 Determination and Control System (ADCS). The secondary objective was to capture high-resolution context images of the impact site for Science.

1.1. 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. The baseline mission success criteria were:

- ♦ Target a short period comet understood to have a nuclear radius > 2 km.
- ♦ Deliver an impactor of mass > 350 kg to an impact on the cometary nucleus at a velocity > 10 km/s. The impact event and crater formation shall be visible from the flyby spacecraft and observable from Earth.
- ♦ Obtain pre-impact visible-wavelength images of the impact site including one with resolution < 3 m and FOV > 50 pixels.
- ♦ Obtain three visible-wavelength images, using at least two different filters, of the entire comet, pre-impact, with resolution < 50 m and average S/N > 50 for the illuminated portion of the nucleus.
- ♦ Obtain five visible-wavelength images containing the impact site with resolution < 50 m and showing the crater evolution from within 3 seconds of time of impact until full crater development (assumed to take less than 660 seconds).
- ♦ Obtain five visible-wavelength images of the ejecta cone, showing the ejecta cone evolution at a resolution < 50 m from within 1 second of impact until late in the cone evolution (assumed to take less than 60 seconds).
- ♦ Obtain five near-infrared (1.1 to 4.8 microns), long-slit spectra of the ejecta cone, showing the ejecta cone evolution with spectral resolving power > 200 from within 2 seconds of time of impact until late in the cone evolution (assumed to take less than 60 seconds).
- ♦ Obtain one image of the final crater with a resolution < 7 m.
- ♦ Obtain one near-infrared (1.1 to 4.8 microns), long-slit spectrum of the impact region pre-impact and one post impact, both with spectral resolving power > 200 and with noise-equivalent-surface-brightness < 150 kRayleigh per spectral resolution element at 3.5 microns.

- Obtain two near-infrared (2.0 to 4.8 microns), long-slit spectra of the coma, one before impact and one after formation of the crater (assumed to take < 660 sec), with spectral resolving power >200 and Noise-equivalent surface brightness <500 kRayleigh per spectral resolution element at 4.7 microns.
- Obtain at least three Earth-orbital or ground-based datasets of two different types of data complementary to the data from the spacecraft.

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 65° for the Tempel 1 nucleus, which induced self-shadowing as a result of the shape and orientation of the nucleus. The Flyby s/c performed a slowing maneuver with a ΔV of approximately 102 m/s to provided 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 begin at 120 min (2 hrs) before TOI. A critical sequence running on-board Impactor s/c spawned science and navigation subsequences that issued Impactor Targeting Sensor (ITS) 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, was responsible for processing these images to form observations for the purpose of trajectory determination (OD). OD updates were performed every minute. Three (3) Impactor targeting maneuvers (ITM) were computed by AutoNav and executed by the Attitude Determination and Control System (ADCS): ITM-1 at E-90 min (E- designates time before impact), ITM-2 at E-35 min, and ITM-3 at E-12.5 min. At E-5 min, the Impactor ADCS pointed the ITS along the AutoNav estimated comet-relative velocity vector to capture and telemeter high resolution ITS images of the impact site prior to impact. Figure 1 shows a schematic diagram of the encounter activities.

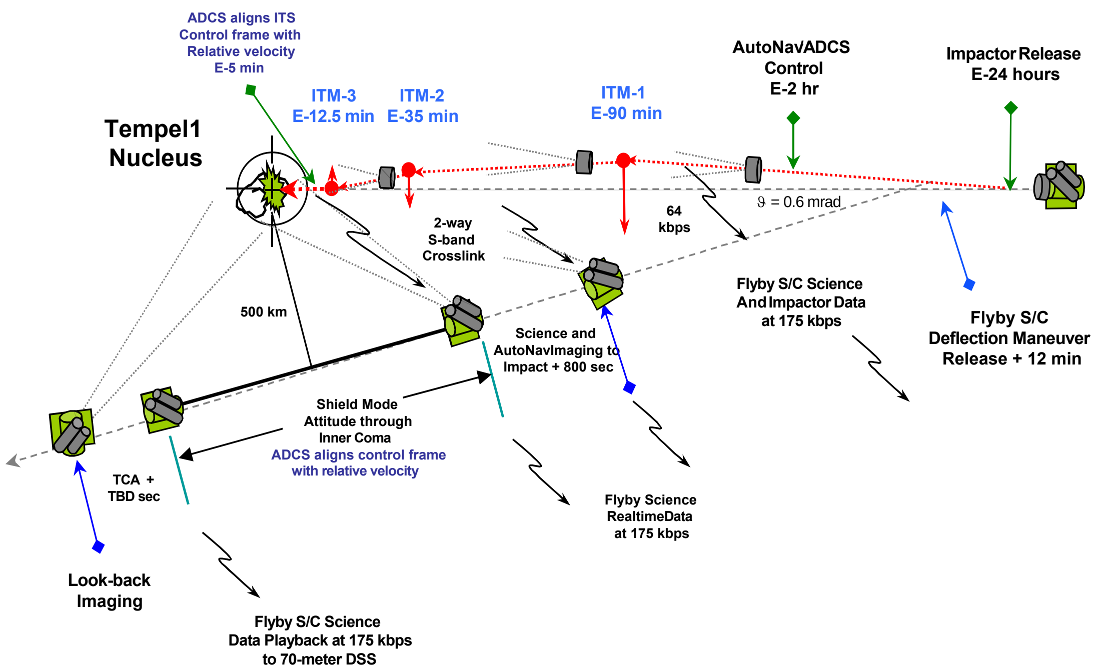

Figure 1. Schematic of Deep Impact Encounter with comet Tempel 1

1.2.Impactor Spacecraft Flight System

The Impactor s/c, shown in Figure 2, was designed and built at Ball Aerospace Technologies Corporation (BATC) and 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 350 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 serves a dual purpose: 1) provide navigation images and 2) provide pre-impact high resolution (< 3 m) science images**.**

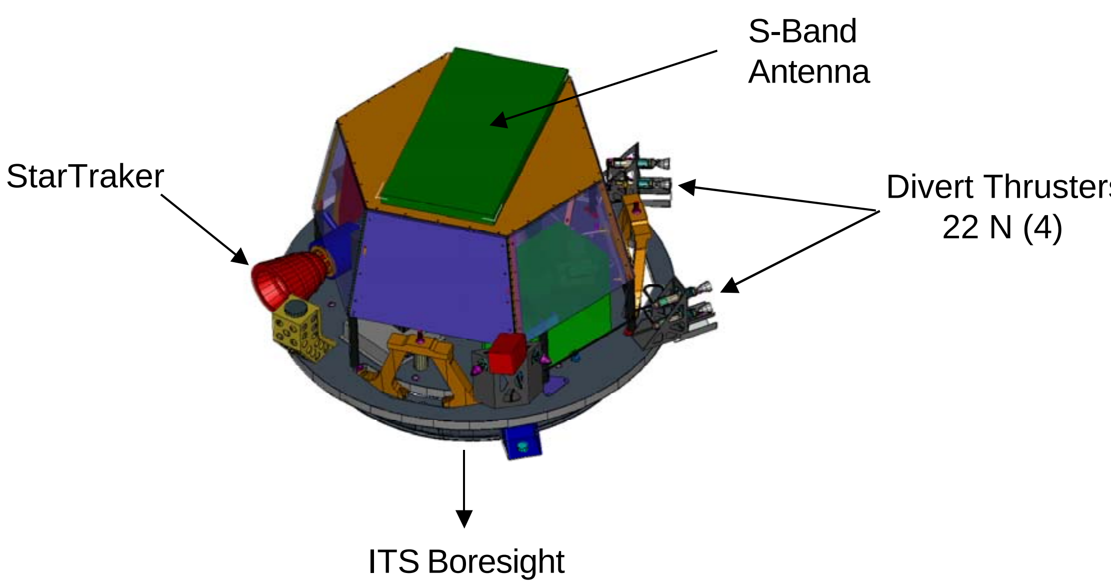

Figure 2. Impactor s/c flight system configuration[1]

1.3.Selection of Impactor Targeting Strategy

General and specific guidance and control laws have been developed for several interceptor designs[2]. There are two basic categories: 1) Tactical homing missiles that use “proportional navigation” or “augmented proportional navigation” techniques, and 2) Strategic interceptors that use either proportional navigation or “predictive guidance” techniques.

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

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

It should be noted that the best quality optical observations are obtained during non-thrust periods, which suggests a pulsed guidance system, as was selected for the Impactor s/c. Though the apriori position of the target body had significant uncertainty, which was removed using optical navigation techniques, the dynamics were well-known, except that (an important exception) the nucleus rotational dynamics and solar phase angle combine to induce motion (acceleration) of the center of brightness (CB) with time, which can cause targeting errors in the impact location 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.

In summary, the Deep Impact Impactor s/c uses 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 (apriori 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 to achieve impact.

1.4.Autonomous Navigation

The autonomous navigation system for terminal guidance of the Impactor s/c relies on both the performance and interaction of the AutoNav and ADCS flight software and subsystems and the ITS camera. As previously mentioned, the AutoNav software was developed for the New Millineum Deep Space 1 (DS1) mission[3,4]. 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 and SSIRU rate sensor provides an estimated attitude bias of no more that 150 μrad (3σ), bias stability of 50 μrad/hr (3σ), and estimated attitude noise of 60 μrad (3σ).

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. 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 of the Impactor. 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 must 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.

2.0 Impactor Encounter Sequence Design

The Impactor sequence design is shown in Figure 3. The critical sequence is the backbone that is initiated upon Structures and Mechanisms (SAM) subsystem separation detection and consists of absolute time commands. The critical sequence is responsible for spawning the engineering and imaging subsequences. There were six engineering subsequences to allow for mark and roll-back opportunities. The engineering subsequences issued commands related to AutoNav commanding, ADCS commanding, Flight System configuration commanding, and telecommunications commanding and log file downlink commanding. There were five ITS AutoNav and Science imaging subsequences corresponding to the last five engineering subsequences. The critical sequence allowed for one sequence time shift based on the updated TOI information received from the Flyby s/c prior to spawning the final engineering subsequence which in turn spawns the final ITS Science pre-impact imaging subsequence. The pre-ITM transition subsequence and the post-ITM transition subsequence are identical for all ITMs and were initiated by the Mode Manager subsystem when requested by ADCS. ADCS can only request an ITM if maneuver information is received from AutoNav. AutoNav will only send information to ADCS following successful completion of a maneuver computation.

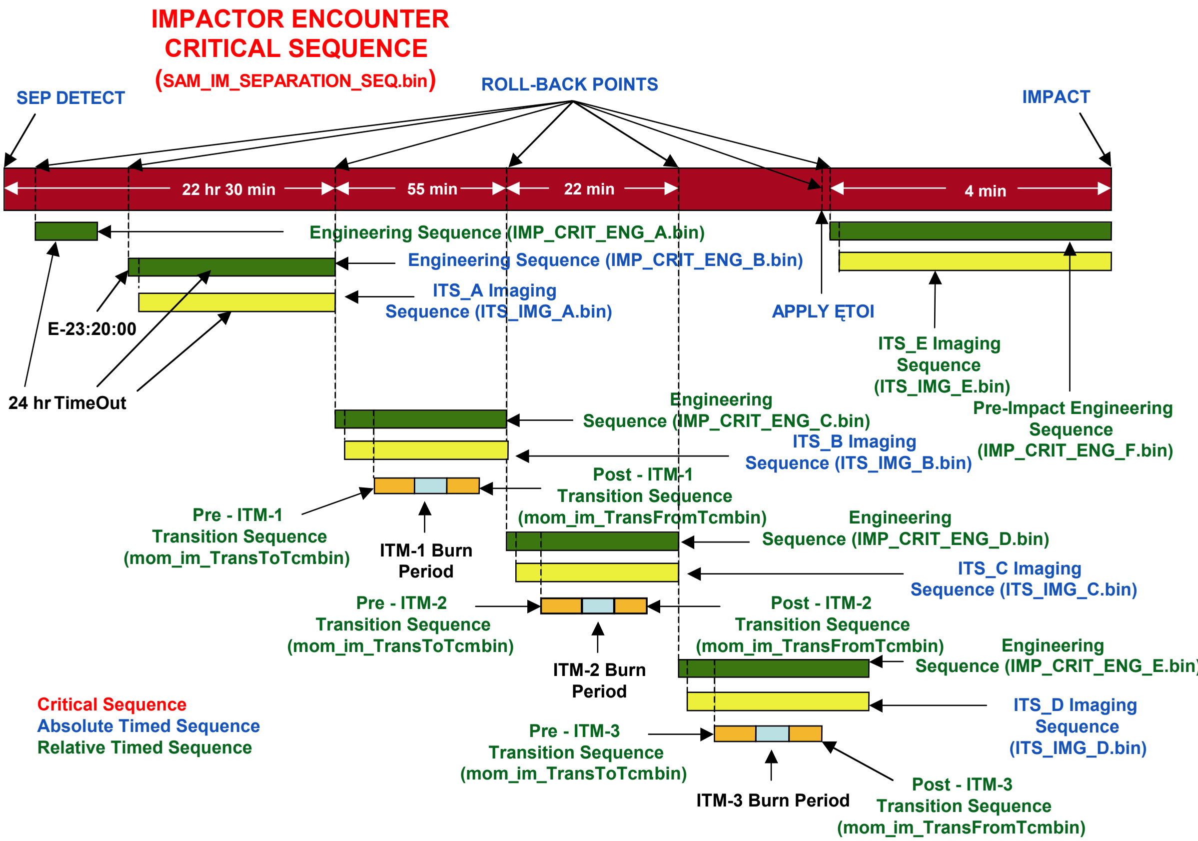

Figure 3. Impactor encounter sequence design schematic

2.1. Configuring the Impactor for Release

The Impactor pre-release configuration sequence was an absoluted-timed sequence that was uplinked to the Impactor s/c and activated several days before separation. This sequence was designed to begin issuing commands 6 hours before separation. The sequence consisted of a series of commands to configure Fault Protection; spawn sequences to configure the battery; configure the ITS; configure ADCS; configure the propulsion system by opening the latch valves and powering the RCS thrusters; configure the thermal control system; configure the S-band antenna by removing the attenuators, powering the receiver and switching the command source from hardline to S-band; powering on the hall-effect sensors for separation detection; and enabling telemetry data storing (TDS) to record engineering telemetry during the period between electrical disconnect and S-band transmitter power-on following the post-release attitude rate-capture.

In the days leading up to encounter, it was discovered that the Impactor encounter sequence design had a curious flaw that could leave the Impactor s/c in a state that would preclude it from receiving commands from the ground should a failed mechanical separation occur after a successful electrical disconnect. In this scenario, the Impactor s/c would await separation detection before spawning the post-release rate-capture sequence that would switch the command source from hardline to S-band to establish a communication link, however, there would be no separation detection if the two s/c fail to mechanically separate. A simple absolute-timed sequence was built, tested, uplinked and activated on the Impactor s/c to automatically switch the command source regardless of separation detection.

2.2.Impactor Targeting Maneuvers (ITMs)

ITMs were initiated via an AutoNav sequence command. These commands were issued from three of the six engineering subsequences that were running during encounter. When AutoNav receives the command, an impulsive maneuver (magnitude and direction) relative to the ICRF frame is computed for the time contained in the command packet and passed to ADCS in the form of a command issued by AutoNav. ADCS receives the maneuver ΔV information in the command packet, computes the finite duration burn start time, and populates the necessary flight software current value table (CVT). When the CVT is populated, ADCS waits until the proper time and issues a command to spawn a trajectory correction maneuver (TCM) transition sequence consisting of commands to disable AutoNav image processing, reorient the s/c to the burn attitude, clear the accumulated ΔV reported by the accelerometers, and set the ADCS state to ΔV in preparation for the burn. When the burn start time is reached, the burn is initiated. The ADCS flight software continually monitors the accumulated ΔV information based on incremental ΔV values measured by the SSIRU accelerometers and continuously adjusts the pointing and thruster duty cycle until the accumulated ΔV matches the desired maneuver ΔV, which results in burn termination. When the burn is complete, ADCS returns to instrument point mode and issues another command to spawn a second TCM transition sequence that re-enables AutoNav image processing following completion of the turn from burn attitude back to the nominal imaging attitude.

The need for TCM transition sequences and for disabling AutoNav image processing arises from the non-deterministic burn duration of each ITM. This allows instrument commands for AutoNav images to be issued every 15 sec without regard to when a particular maneuver will occur or how long the burn will last. Since the quality of the images acquired during the burn may be degraded, AutoNav simply ignored the navigation images received during these periods. The entire ITM process and the entire encounter sequence design was rigorously tested during the Mission Readiness Test (MRT) program through the Impactor Encounter Mission Scenario Test runs and during a dedicated robustness testing effort conducted in the months leading up to encounter.

2.3.ITS Science Imaging

The ITS camera was a split-frame transfer device that allowed four subframe modes and one full-frame mode: VISMODE_FF generated 1024x1024 pixel full-frame images; VISMODE_SF1 generated 512x512 pixel subframe centered on the CCD; VISMODE_SF2S generated 264x264 pixel subframed images centered on the CCD; VISMODE_SF3S generated 128x128 pixel subframed images centered on the CCD; and VISMODE_SF4NO allowed for high-rate acquisition of 64x64 pixel subframed images centered on the CCD. The purpose of the subframed modes was to allow control of the temporal resolution of images at the expense of spatial coverage.

The five ITS imaging sequences contained commands for ground-based optical navigation (OpNav), AutoNav and Science. Figure 4 shows the imaging sequence layout and lists the types of ITS commands contained in each sequence. The ITS_IMG_E sequence contained the high-rate commands to capture the context images prior to impact and extended 180 sec beyond the nominal TOI to account for the possibility of a later-than-expected TOI. The SF4NO images could be acquired at a frequency of 1 Hz. Since the Impactor s/c had a closing speed of 10.3 km/s, relative to Tempel 1, the ITS resolution improvement from image to image was on the order of 10 cm. For an image taken 4 sec prior to impact, the range to the surface would be approximately 40 km. The resolution of the ITS at that range was approximately 41 cm and that decreases by 10 cm with each subsequent 64x64 pixel image that was taken. In contrast, the maximum frequency of the SF2S images, for example, was 0.1 Hz. So, without the ability to subframe, the probability of capturing pre-impact images with resolution of < 40 cm decreases with increasing subframe size.

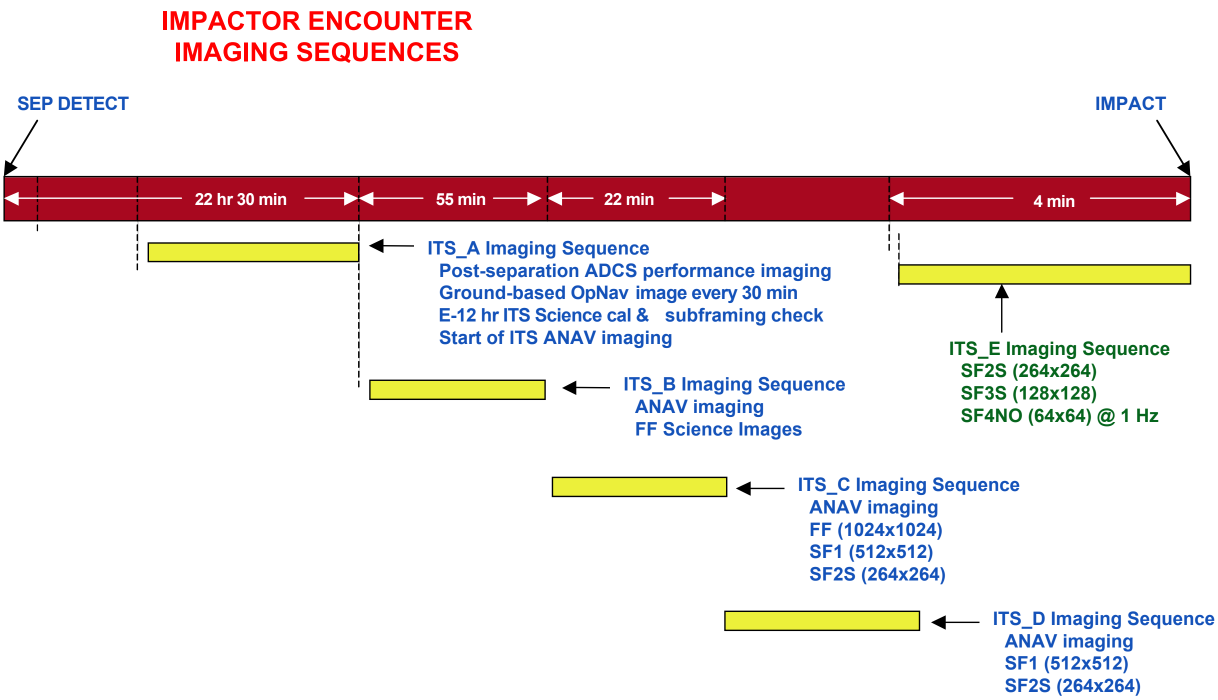

Figure 4. ITS imaging subsequence layout

2.4.Contingency Planning

For the Impactor s/c the contingency planning was limited due to the hands-off nature of the sequence design and the desire not to rely on the S-band antenna for command uplink following separation. The focus of the contingency planning effort dealt with the uncertain nature of the Tempel 1 nucleus. In the years prior to encounter, a best estimate of the nucleus albedo was determined by the Science Team and recommended to the Spacecraft Team. The importance of selecting the correct ITS exposure durations could not be overlooked as it directly affect AutoNav images and Science interpretation. Selecting exposure durations that were too short could lead to underexposure and insuffient signal for AutoNav target acquisition and Science interpretation. Selecting exposure durations that were too long could lead to overexposure and CCD pixel saturation that could distort the nucleus, influence the targeting algorithms and interfere with Science interpretation of the surface features.

In an effort to mitigate the risk of an incorrect nucleus albedo estimate, a multiple sequence strategy was proposed and adopted. This strategy consists of flying with three sets of Science observing and AutoNav imaging sequences. Since it was unlikely that there would be any improvement in the albedo estimate, beyond what is known over a year before encounter, three sets were generated, tested and uplinked to the Impactor before release. The switch to/from each sequence set could be activated anytime before separation as well as after separation should there be any late information that would conclusively indicate the current active sequence would be insufficient. The off-nominal sequences were designated as bright and dim. The bright sequences contained exposure settings that were a factor of 2 shorter than the nominal set and the dim sequences contained exposure settings that were a factor of 2 longer than the nominal set.

The decision to make a switch could only be reached via a formal process that consisted of a single albedo selection summit meeting that was held 30 hrs before the nominal TOI, between the Science Team, the Navigation Team, and Project Management, followed by a post-release contingency meeting opportunity at E-8 hrs. The Science Team presented the data based on the latest set of approach phase observations from the Flyby s/c. Their recommendation was presented and a unanimous decision was made by the selection committee to accept the recommendations of the Science and Navigation Team and stay with the nominal set of exposures.

The threshold or criterion for switching from the nominal sequence exposure settings was decided, several months before encounter, to be 1.5 times the nominal nucleus brightness. Should the actual observed nucleus brightness exceed that factor of the nominal brightness, a sequence set switch would be recommended. The process of determining the actual brightness will not be discussed here, but it depended on the ability to separate the optical signal of the coma from the optical signal of the nucleus to predict the peak brightness of the spatially resolved nucleus. The data used in the process came from MRI and HRI images acquired from the Flyby s/c in the several days leading up to encounter.

3.0 Impactor Encounter Performance

The performance of the Impactor s/c exceeded all expectations. The s/c had never flown under its own control until release, which was just 22 hrs before it was to perform critical autonomous operations. Following separation, the Impactor ADCS captured rates and oriented the s/c to point the ITS at comet Tempel 1 and began transmitting telemetry to the Flyby s/c. Confirmation of the S-band signal was received and all systems assessed to be healthy and ready for encounter.

The performance of the Impactor flight system and sequence design during encounter can be measured by looking at the performance of the closed-loop AutoNav/ADCS design and looking at the the Science data return from the Impactor s/c during the 24 hr free-flight.

3.1. Autonomous Navigation Performance

In the weeks leading up to encounter, the Impactor ADCS performance was the subject of intense scrutiny. The software in the CT-633 StarTracker was inducing unacceptably large discontinuities in the attitude quaternion information that was being used in the attitude estimation process. This resulted in larger than desired attitude instability, which influences AutoNav’s estimation of comet-relative velocity and in turn results in computed targeting maneuver errors. The ADCS engineers provided recommendations for attitude filter parameters that, in the end, led to extremely good ADCS estimated attitude stability.

The performance of AutoNav trajectory estimation depends directly on the performance of ADCS attitude estimation. During encounter, the ADCS estimated absolute attitude error was approximately 200 μrad; slightly larger than the 3σ requirement of 150 μrad, while the attitude stability was seen to be approximately 25 μrad/hr, a factor of two better than the 3σ requirement of 50 μrad/hr. Figure 5 shows the images used in the autonomous targeting process. At the start of AutoNav/ADCS closed-loop control the nucleus spanned approximately 10 pixels. Prior to ITM-2, the nucleus had grown to span nearly 50 pixels. At the end, 100 sec before impact, the ITS was pointed along the comet-relative velocity vector and the nucleus spanned more than 2 times the ITS field-of-view.

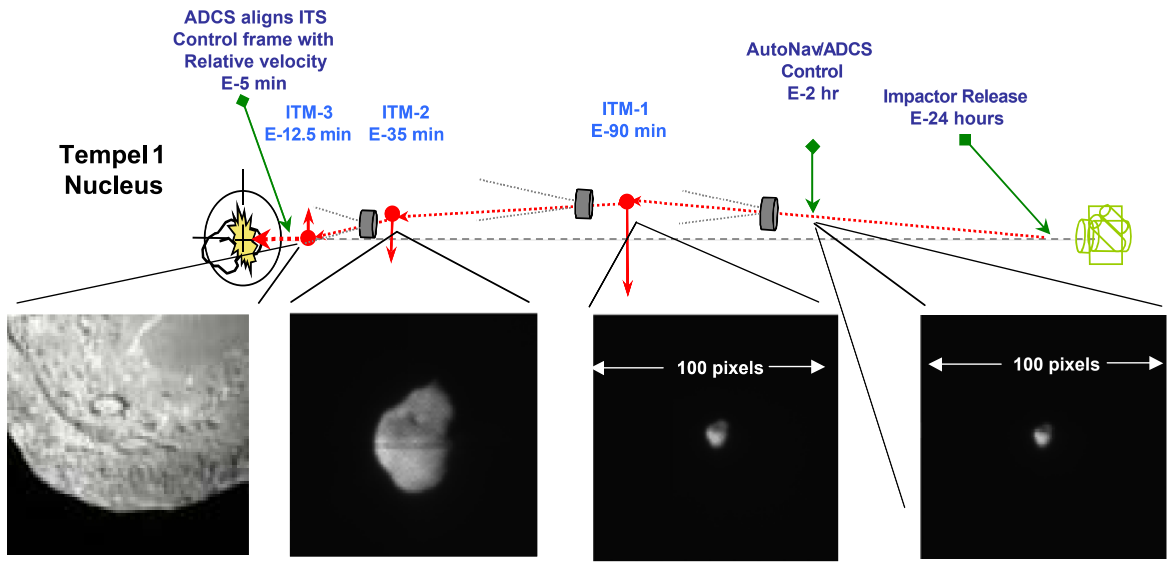

Figure 5. Sequence of images taking during the AutoNav/ADCS closed-loop operations

Figure 6 shows the B-plane intercept for the pre-release delivery and the performance of each ITM. The Flyby s/c delivered the Impactor s/c to < 2.5 km in the B-plane relative to the actual impact site that was selected and targeted by AutoNav. ITM-1 was based on 25 min of optical CB observations and targeted the CB. AutoNav commanded a maneuver of 1.27 m/s in the lateral (normal to the comet-relative velocity vector). As seen in figure 6, ITM-1 modified the delivered trajectory, but did not improve the targeting accuracy. The cause is still under investigation, however, the purpose of ITM-1 was to remove large initial delivery errors (as large as 30 km) without expending an unnecessary amount of propellant. Possible causes include: 1) ADCS attitude stability transients; 2) nucleus rotation and lighting conditions; 3) increasing nucleus size and improving resolution during ITM-1 OD arc.

ITM-2 also targeted the nucleus CB and AutoNav commanded a maneuver of 2.26 m/s. As seen in figure 6, the performance was exceptional and placed the Impactor s/c on a trajectory to impact within approximately 300 m of the observed CB.

ITM-3 targeted to an offset that was determined using the biased Scene Analysis algorithm. Scene Analysis essentially consists of computing the pixel/line location of the center of an illuminated circle with radius (parameter set in AutoNav and based on expected Flight System performance) of 1.17 km that is biased in the direction of the Flyby s/c point of closest approach for optimizing Flyby s/c viewing. The difference between the CB pixel/line location and the Scene Analysis selected pixel/line location was converted to an inertial offset and used to compute ITM-3 burn direction and magnitude of 2.29 m/s.

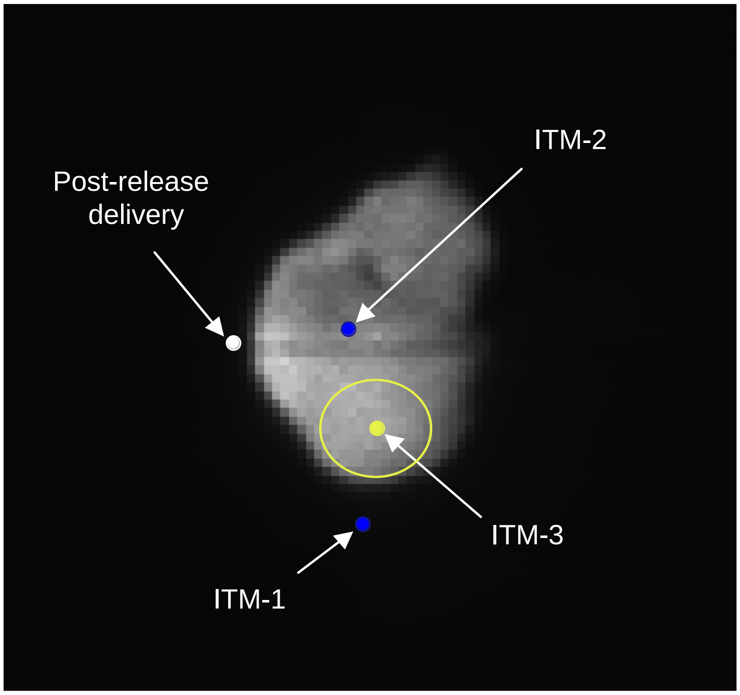

Figure 6. Performance of ITMs shown with respect to an AutoNav image taken just prior to ITM-2. Range is approximate 21,000 km. The CB of this image was located within a few pixels of the ITM-2 predicted intercept location

Figures 7 and 8 compare what was expected prior to encounter assuming nominal performance and what was observed during the actual encounter with comet Tempel 1. The images in figure 7 are simulated and based on comet Halley data. The images shown in figure 8 are Science images acquired during the last 2 hrs of encounter. The blue arrow points to the actual impact site as determined from Flyby images of the flash at the time of impact. As seen in figure 8, the closed-loop AutoNav/ADCS control maintained the ITS camera boresight within approximately 5 pixels of the actual impact site. The last high-rate (64x64 pixel subframe) Science image received on the ground was capture 3.7 sec before the nominal TOI. At that time, the range to the surface was 38 km, and the pixel scale was 38 cm/pixel. The 64x64 pixel subframe footprint spanned approximately 24 m. The Science Team and Instrument Engineering Team are currently investigating a noticeable change in the ITS camera characteristics in the last 30 sec. That fact, combined with at least two observed attitude upsets in the last 30 sec seem to indicate ITS degradation due to particle (dust) impacts within 300 km of the surface.

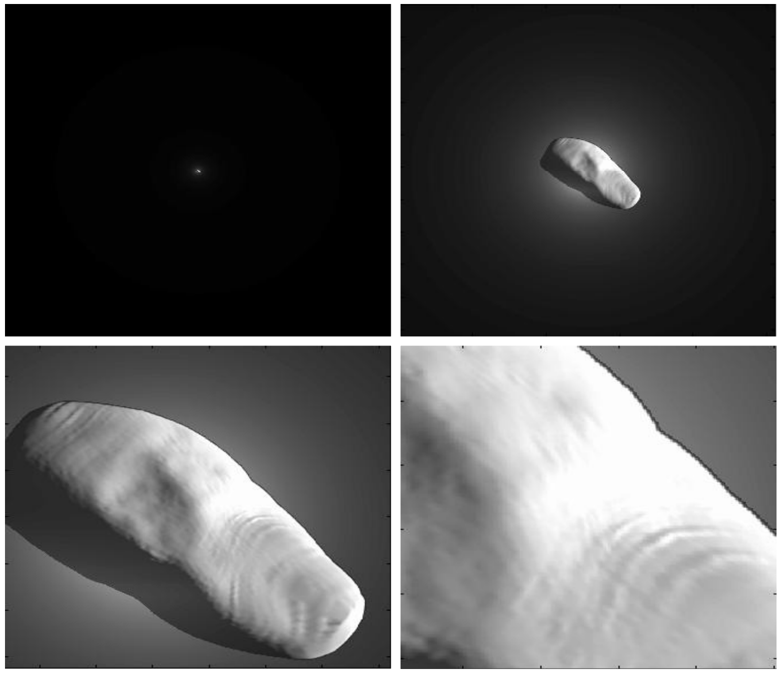

Figure 7. Simulated ITS images showing what could be expected. The images are based on the comet Halley model

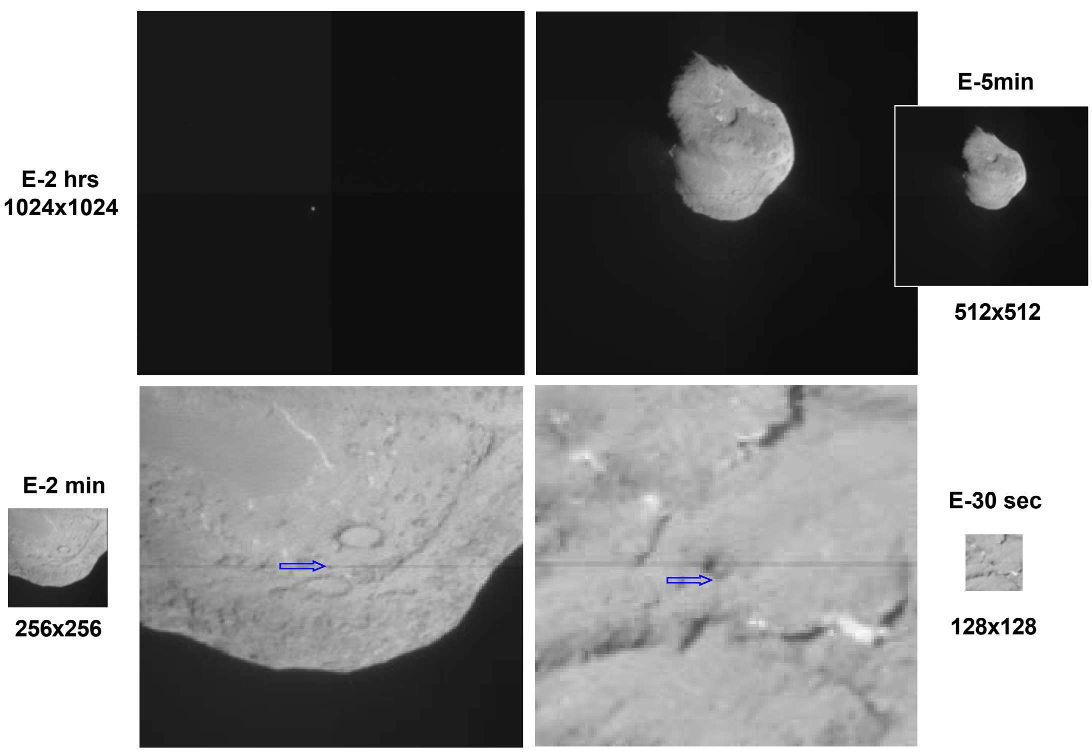

Figure 8. Actual ITS images of Tempel 1 acquired during encounter. The blue arrow indicates the actual impact site as determined from Flyby s/c images

4.0 Summary

The Deep Impact Mission to Comet Tempel 1 was a unique and successful mission. The science and engineering objectives for the Impactor s/c were met through the nominal execution of the critical sequence and each subsequence:

- Impactor targeting led to an illuminated impact that was viewable from the Flyby s/c
- Science imaging sequences captured images of the nucleus surface with the desired temporal and spatial resolution
- The Impactor s/c captured several high-resolution context images of the actual impact site, with the last being captured only 3.7 sec before impact providing the highest resolution images ever taken of the surface of a cometary nucleus
- The Impactor s/c transmitted all Science and Engineering telemetry needed for post-encounter reconstruction and Scientific interpretation

In the end, none of the contingency plans were needed and the Flight System performed as designed.

5.0 Acknowledgements

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

6.0 References

[1]Phinney, A. "Mass on the Impactor." *BATC Systems Engineering Report DI-IMP-STR-004*, 2001.

[2]Zarchan, P, *Tactical and Strategic Missile Guidance, 3rd Edition.* Progress in Astronautics and Aeronautics, v. 176, American Institute of Aeronautics and Astronautics, Reston, VA.

[3]Bhaskaran, S., Riedel, J. E., Synnott, S. P. "Autonomous Optical Navigation for Interplanetary Missions." *Space Sciencecraft Control and Tracking in the New Millennium*, Proc. SPIE, pp. 32-34, 1996.

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