DEEP IMPACT COMET ENCOUNTER:
DESIGN, DEVELOPMENT, AND OPERATIONS OF THE BIG EVENT AT TEMPEL 1

Steven Wissler, Jennifer Rocca, and Daniel Kubitschek
Deep Impact Flight Team
Jet Propulsion Laboratory, California Institute of Technology
JPL M/S 301-490
4800 Oak Grove Drive
Pasadena, CA 91109
Phone: (818)354-8649
Fax: (818)393-6871
Email: Jennifer.M.Rocca@jpl.nasa.gov

ABSTRACT

Deep Impact is NASA's eighth Discovery mission. This low-cost, focused planetary science investigation gathered the data necessary to help scientists unlock early secrets of our solar system. The comet encounter with Tempel 1 was a complex event—requiring extremely accurate timing, robustness to an unknown environment, and flight team adaptability. The mission operations and flight systems performance were spectacular for approach, impact, and lookback imaging on July 4, 2005.

The Encounter Phase began one week prior to impact with a strenuous imaging campaign, including mission critical optical navigation and instrument calibrations, plus imaging for coma characteristics and nucleus rotation. Until just prior to impact, the comet nucleus was unresolved spatially, making operations demanding in the final days. Flyby and Impactor separation occurred 24 hours before impact, beginning the most intense part of DI's mission. The independent, yet coordinated missions of the two spacecraft were accomplished with elegant sequence design, an ambitious fault protection and recovery strategy, and a battery of punishing tests performed with very few resources to retire risk for the encounter event.

Encounter Approach imaging and calibrations were successfully executed, giving the science and navigation teams an array of early images in preparation for impact. The final Trajectory Correction Maneuver, Autonomous Navigation, and Attitude Control operations achieved desired performance during encounter—enabling stunning images captured as the Impactor approached Tempel 1, and as the flyby witnessed the impact and subsequent departure of the comet. The key strategy for data management, "live for the moment," was employed as planned; all high priority images were transmitted to the ground prior to closest approach.

Despite engineering challenges, an unknown and potentially dangerous comet environment, and a short operations phase, Deep Impact was a dramatic success. This Discovery mission spent Independence Day with the world watching—what a sight!

1.0 INTRODUCTION

The Deep Impact mission explored the interior of Comet Tempel 1 by using a 364-kg Impactor to excavate a crater in the comet's surface and taking data on the newly-exposed cometary interior with a companion flyby spacecraft. Deep Impact is the eighth mission in NASA's Discovery Program, following NEAR, Pathfinder, Lunar Prospector, Stardust, Genesis, CONTOUR, and MESSENGER. The project was organized as a team between the principal investigator, Dr. Michael A'Hearn of the University of Maryland; the science team of eleven other prominent experts on comets, remote sensing, and impact physics; the industrial partner, Ball Aerospace & Technologies Corp.; and the Jet Propulsion Laboratory as the NASA lead center.[1]

This low-cost, focused planetary science investigation gathered the data necessary to help scientists unlock early secrets of our solar system. The comet encounter with Tempel 1 was a complex event—requiring extremely accurate timing, robustness to an unknown environment, and flight team adaptability. The mission operations and flight systems performance were spectacular for approach, impact, and lookback imaging on July 4, 2005.

1.1 Mission Synopsis

The Deep Impact mission conducted a scientific cratering experiment on the nucleus of comet Tempel 1 near the time of perihelion for its 2005 apparition. This was accomplished by launching two joined spacecraft January 12, 2005 to approach the comet in early July 2005. The short six month cruise phase was comprised of a battery of engineering checkout/calibration activities (commissioning), and in-flight demonstrations of autonomous navigation and encounter-like imaging. Two course maneuvers were conducted during cruise to refine the path to Tempel 1. Figure 1 shows the flight system.

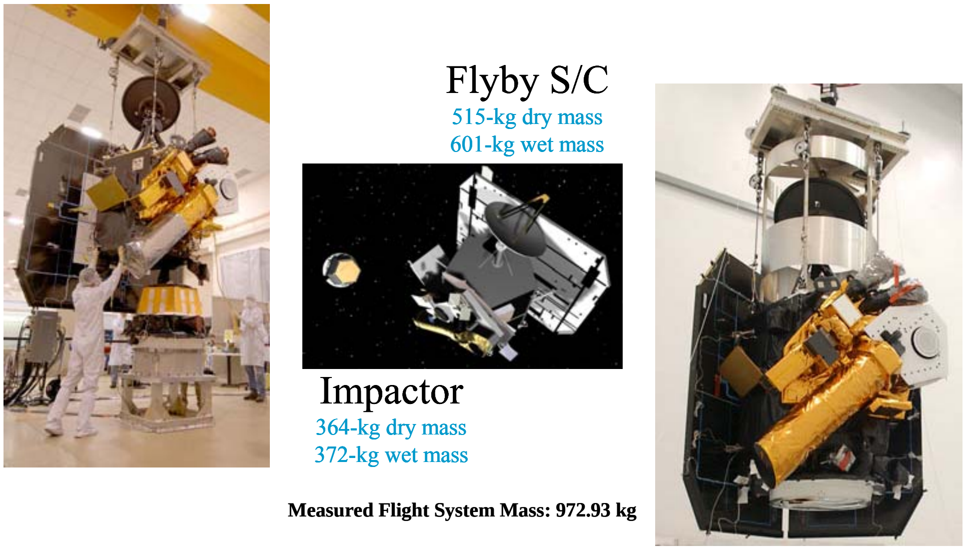

Figure 1: Deep Impact Flight System

Encounter Approach began one week prior to impact with a strenuous imaging campaign, including mission critical optical navigation and instrument calibrations, plus imaging for coma characteristics and nucleus rotation. Until just prior to impact, the comet nucleus was unresolved spatially, making operations demanding in the final days. Using spacecraft optical observations of the comet and conventional ground-based navigation techniques, the joined spacecraft were maneuvered as close as possible to a collision trajectory with the nucleus of Tempel 1, and the Impactor was released 24 hours before impact.

The Impactor observed the approaching nucleus with an optical camera and maneuvered itself to a collision course toward the lighted portion of the nucleus. After separation from the Impactor, the flyby spacecraft performed a maneuver to delay and deflect its flight path toward the nucleus so that it could observe the impact, ejecta, crater development, and crater interior during a 500-km flyby of the nucleus that occurred about 14 minutes after the impact. The flyby spacecraft carries a remote sensing payload of two instruments for imaging and infrared spectroscopy. Close-in observations of the nucleus by the Impactor camera were sent to the flyby spacecraft by an S-band radio link in the last seconds before impact. Simultaneous observations of the comet before, during, and after the impact were also conducted from ground and space-based observatories as an essential part of the total experiment. All scientific and supporting engineering data are archived for future use by the scientific community. Figure 2 shows an overview timeline for the mission.

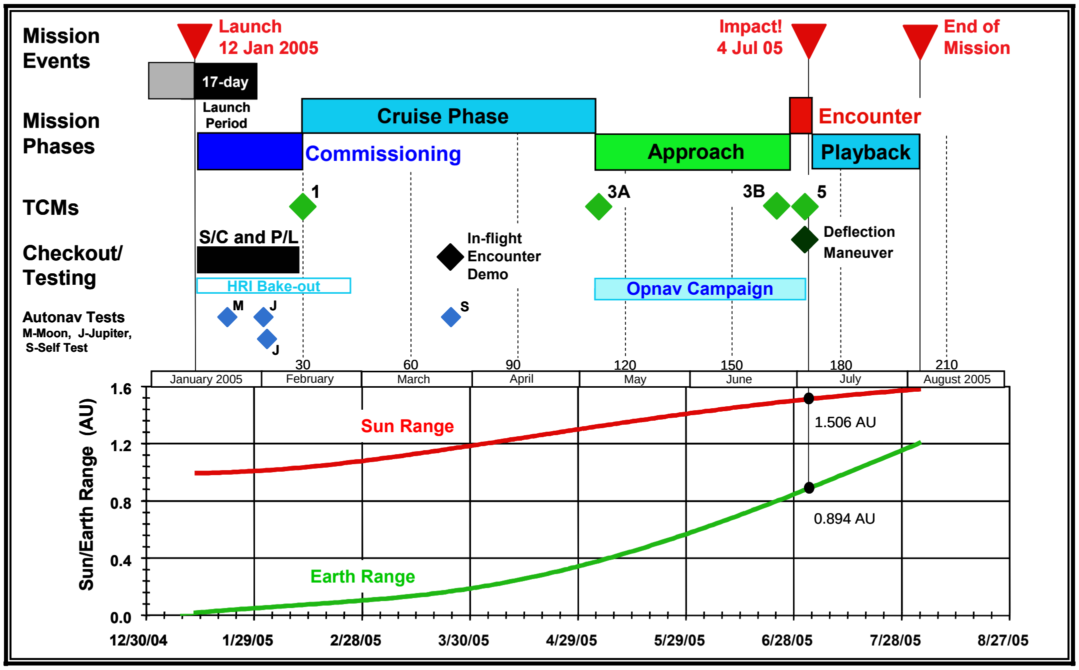

Figure 2: Deep Impact Mission Timeline

1.2 Mission Objectives

From the Discovery Program Plan, Deep Impact's Primary Mission Objectives:

- Dramatically improve the knowledge of key properties of a cometary nucleus and, for the first time, assess directly the interior of a cometary nucleus by means of a massive Impactor hitting the surface of the nucleus at high velocity.
- Determine properties of the surface layers such as density, porosity, strength, and composition from the resultant crater and its formation.
- Study the relationship between the surface layers of a cometary nucleus and the possibly pristine materials of the interior by comparison of the interior of the crater with the pre-impact surface.
- Improve our understanding of the evolution of cometary nuclei, particularly their approach to dormancy, from the comparison between interior and surface.

1.3 Comet Encounter Overview[2]

Critical comet encounter activities began 24 hrs prior to the expected time of impact (TOI), when the two spacecraft separated. The encounter geometry resulted in an illumination phase angle of approximately 65° for the Tempel 1 nucleus and induced some self-shadowing due to the shape and orientation of the 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. Figure 3 shows a schematic diagram of the encounter activities.

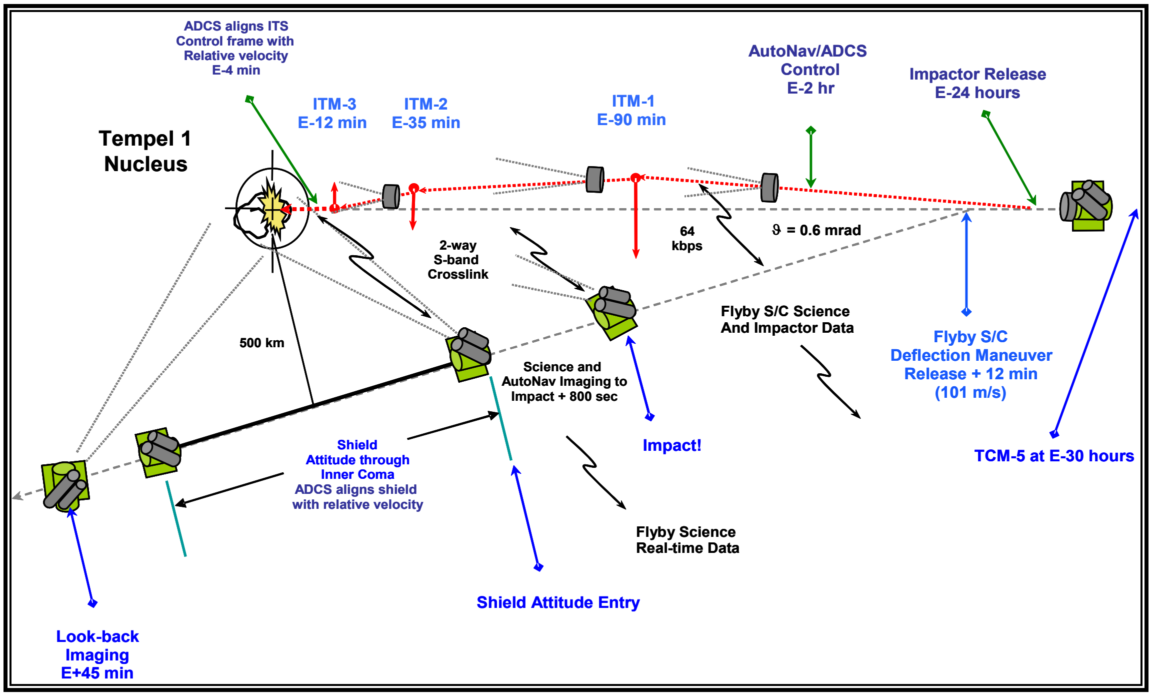

Figure 3: Comet Encounter Schematic

The autonomous phase of the encounter began at 120 min (2 hrs) before TOI. A critical sequence running on-board the Impactor spawned a science and navigation subsequence that issued Impactor Targeting Sensor (ITS) commands to produce navigation images at a 15 sec interval. The Autonomous Navigation (AutoNav) software (originally developed and demonstrated during the Deep Space 1 (DS1) mission) processed these images to form observations for the purpose of trajectory determination (OD). OD updates were performed every minute. Three Impactor Targeting Maneuvers (ITMs) were computed by AutoNav and executed by the Attitude Determination and Control System (ADCS): ITM-1 at E-90 min, ITM-2 at E-35 min, and ITM-3 at E-12 min. At E-2 min, the Impactor ADCS pointed the ITS along the AutoNav estimated comet-relative velocity vector to capture and telemeter high resolution (20 cm) ITS images of the impact site prior to impact. Meanwhile, the AutoNav software on the Flyby s/c was processing MRI images of the comet every 15 sec and updating the trajectory of the Flyby s/c every minute to continuously point the MRI and HRI instruments at the nucleus.

2.0 COMET ENCOUNTER DESIGN

The following sections will discuss the design approach for Deep Impact's encounter with Tempel 1, with a focus on the flyby spacecraft activity and sequence design.

2.1 Design Envelope

The Deep Impact encounter phase was very complex for a low-cost Discovery mission. The mission requirements and cost trades which drove the design of the encounter design and operations were:

- Live for the moment science return: There was a significant chance that the Flyby spacecraft would not survive the cometary dust environment at the closest approach to Tempel 1. The primary science data had to be transmitted by the spacecraft prior to the time of closest approach. This requirement drove the design of the telecom subsystem, requiring higher-power amplifiers and a larger high-gain antenna.
- High resolution crater imaging: This requirement increased the accuracy requirements for attitude control (precision and stability) and necessitated the use of on-board autonomous navigation on the Flyby spacecraft.
- 30km Impactor release accuracy: The requirement to release the Impactor at 24 hours prior to impact with a B-plane accuracy of less than 30 km required a trajectory correction maneuver (TCM) to be done 6 hours prior to Impactor separation and with a tracking data cut-off 11 hours prior to maneuver execution.

- On-board autonomous fault protection and critical sequence resumption. This level of fault protection and recovery was far more complex than on other discovery class missions and was difficult and costly to test.
- De-scopes to the flight system due to cost increases also drove up operational complexity. Some examples:
  - On-board file system space was cut in half. This increased the complexity and frequency of on-board file management. The backup flight computer file storage changed from being entirely redundant storage to storing some amount of unique images.
  - The communication link between the two flight computers was descoped, affecting the ability of fault protection software to autonomously recover from a flight computer swap. In the event that a swap occurred, the re-configuration of the spacecraft and re-starting the execution of the encounter sequence would have required ground-in-the-loop.
  - Downlink data rate was de-scoped from 400,000 bps to 200,000 bps. This descope increased the contention for bandwidth for science images, AutoNav images and Impactor telemetry files, significantly increasing the complexity of the sequence design and testing.

2.2 Encounter Approach

The Encounter Phase began with a week of Encounter Approach activities spanning from Impact minus eight days to Impact minus 28 hours. The last week prior to encounter was one of the most challenging of the mission due to the vast scope of activity that had to be accomplished so close to the impact event. Key elements of Encounter Approach included: activation of the critical encounter sequence, maneuvering to continuous comet-pointing attitude, extensive optical navigation and science imaging sequence execution, divert maneuver design and uplink, final science point spread function and infrared mini-calibrations, Trajectory Correction Maneuver design and uplink, and final preparation of the Impactor spacecraft for separation and its subsequent free-flight to Tempel 1. In addition to the major elements noted above, regular maintenance activities were also performed during Encounter Approach, including regular science and navigation image downlinks, ITS imaging and downlinks, delta-DOR, and Impactor engineering telemetry playback. Figure 4 depicts the phase timeline for Encounter Approach.

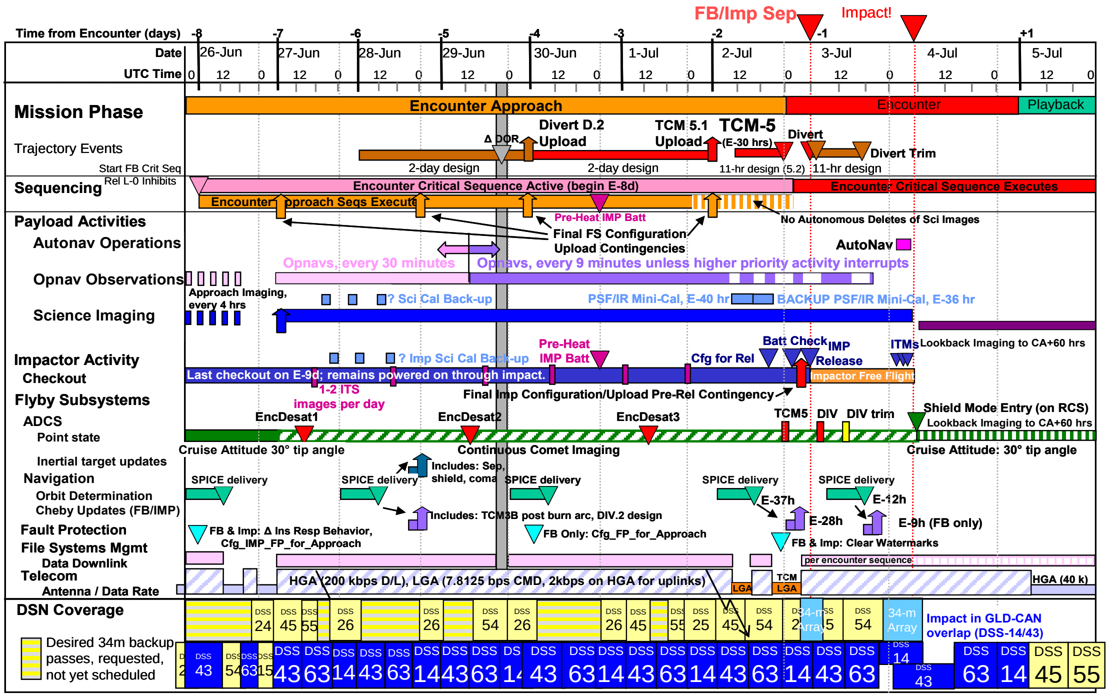

Figure 4: Encounter Phase Timeline

All encounter approach activities were designed to execute on-board via sequences, with the only ground commands issued to kick off the backbone sequence. Continuous 70-m antenna coverage and nearly continuous 34-m antenna back-up coverage provided by the Deep Space Network were required to support this period of Deep Impact's mission. During this final week, both spacecraft were fully powered, with all instruments and engineering subsystems in operation. This week represented the first extended period of continuous Impactor spacecraft operation outside of checkout/calibration/demonstration activities performed during cruise.

It was discovered during early commissioning that Deep Impact's High Resolution Imager was out of focus. The science team devised a plan to use deconvolution on images taken with this telescope. In order to perform successful deconvolution, a set of mini calibrations needed to be performed within 48 hours of impact. These calibration images were labeled as critical to mission success, as they would enable the science team's ability to post-process the HRI images and achieve improved resolution compared to the out-of-focus images.

Activity design for this mission phase was severely constrained. To begin the design effort, several guiding principles were employed: minimize real-time commanding, limit uplinks to pre-selected windows, design all sequences to be self-deleting, complete all downlinks prior to end of phase, and delete all images and sequences not needed for encounter at completion of phase. Fault protection design during this phase was similar to that used during cruise, with the addition of instrument fault protection. All sequences were non-critical, meaning that if fault protection interrupted them, they would not be autonomously restarted. Elements with time constraints were placed into the timeline first: TCMs, calibrations, Impactor preparations for separation, and uplinks. Following the placement of time critical activities, other required activities were placed on the timeline. The imaging campaign was the most challenging element to accommodate in the timeline. This imaging sequence, spanning seven days, was by far the largest sequence designed for Deep Impact. The data flow aspect of Encounter Approach was problematic with thousands of images being taken for both science and navigation. With very little design space, data flow modeling was essential. Much iteration with this model (and varied sequence timing) finally produced a sequence design that was predicted to meet all constraints.

2.3 Separation/Divert

The Impactor was separated from the Flyby 24 hours prior to impact. This event was similar to launching a spacecraft, then letting it fly alone for only one day prior to its main event. The Flyby spacecraft was targeted to impact the comet with a propulsive maneuver 6 hours before separation. The Flyby spacecraft needed to perform a propulsive maneuver of 101 meters/sec 12 minutes after Impactor separation in order to slow down and miss the comet by 500 kilometers as planned. This slow down was required in order for the Flyby to observe the impact and crater formation for 13 minutes prior to closest approach.

Table 1: Separation/Divert Event Table

Critical_Sequence_Preparation Sequence
☐ Configure spacecraft, instruments and fault protection for critical sequence operations.
Sep_Flyby_Preparation Sequence
☐ Bring Impactor battery on-line and disable Impactor bus
☐ Slew to the separation attitude.
☐ Power on the separation detection circuits (Hall Sensors)
☐ Turn on divert and RCS cat-bed heaters.
Sep_Flyby_Preparation2 Sequence
☐ Disable the Flyby to Impactor electrical connection.
Separation_Flyby Sequence
☐ Power on the separation circuits
☐ Configure and power on S-band radio.
☐ Separate Flyby to Impactor electrical connection
☐ Separate Flyby to Impactor mechanical connection
Sep_Flyby_Post Sequence
☐ Re-configure Spacecraft for post-separation mass properties and power off separation electronics
☐ Slew to divert attitude
☐ Turn on S-band transmitter
Divert Sequence
☐ Fire divert thrusters for divert maneuver
PostDivertTransition Sequence
☐ Re-configure spacecraft for imaging
☐ Slew to Temple-1 imaging attitude

Due to the critical nature of Impactor release, and its proximity to the actual encounter, the spacecraft had to be able to recover from a fault during separation and divert and re-start the on-board sequence. To implement this, the FSW had a sequence mark and rollback capability, where a sequence could be re-started, and rolled back to a previous event, or move forward to the next event, depending on how the sequence was constructed. The Impactor separation and divert sequences were a set of relative-timed sequences activated by the Flyby Critical Sequence. The separation and divert sequences were broken down into fragments which facilitated fault protection's ability to respond to a fault during separation and divert by interrupting the sequence, fixing the fault, and resuming the sequences. Activities performed during Impactor separation sequences are shown in Table 1.

2.4 Encounter and Look Back/Play Back

After separation and divert, the flyby periodically imaged the comet and sent these images and Impactor telemetry to the ground. Two hours prior to impact, the flyby began autonav operations. Autonav images were generated at a rate of four per minute, with orbit determination (OD) updates once per minute. The OD updates generated a new trajectory update to the attitude determination and control (ADCS) subsystem. This caused small corrections in the flyby's attitude to keep the instruments pointed at the comet. At 38 minutes before impact, the flyby's ADCS transitioned to an attitude mode (Shield Mode) which kept the solar array edge in the direction of the comet dust flow. This allowed the dust shield to protect the solar arrays and instrument electronics.

The a priori error in time of impact (TOI) when the imaging sequences were generated was +/- 3 minutes. The high-rate imaging sequence for impact needed to be started with an accuracy of approximately 3 seconds. The final imaging sequence for the high-resolution crater imaging had a similar timing requirement. The image sequences around impact and shield mode, were relative-timed sequences spawned from the critical sequence. Autonav could adjust when these sequence started by calculating a TOI and time of final imaging (TOFI) update. These updates were calculated by autonav, based on current OD, and sent to the command manager flight software for adjusting the sequence timing. The TOFI update also controlled the timing of the transition to shield mode. This transition was timed to occur when the Flyby spacecraft was 700 km from Tempel 1 and 13 minutes 22 seconds after nominal Impactor impact. The Flyby spacecraft remained in this attitude for 27 minutes, until it was through the worst of the dust environment. Subsequently, it turned back to image Tempel 1. In look back attitude, the High Gain Antenna couldn't be pointed at Earth. The look back imaging was performed on the Low Gain antenna at a downlink rate of 10 bits/second. The Flyby spacecraft was sequenced to periodically switch between an attitude which allowed for playback of stored images, and the look back imaging attitude.

2.5 Data Return Strategy

The Deep Impact Flyby Spacecraft stored engineering telemetry and instrument image files on a file system in non-volatile memory (NVM). There were two 256 MB NVM devices available. The first device had 2 dedicated drives for Flight Software and FSW parameter tables. This first device also had a general purpose data storage drive, which contains sequence and parameter files uplinked from the ground, all engineering telemetry files and "overflow" instrument image files. The second device was dedicated to instrument image files. Table 2 shows NVM Allocations for encounter.

Table 2: Encounter Flyby S/C NVM Allocations

Data Type	Allocation
Science	309 MBytes
Navigation	6.95 MBytes
Impactor Telemetry	30 MBytes

The Instruments could send images to either or both of the flight computers. The backup flight computer's file system was primarily used for redundant storage, however, due to the descope in file system space, some images were written non-redundantly to one or the other flight computer in order to preserve some of the original science data. This significantly complicated the sequencing and data return strategy. During playback, it was necessary to swap to the backup flight computer to retrieve images. Swapping telemetry streams was always a risky proposition because the backup didn't have an active 1553 interface so there was no engineering telemetry from the active spacecraft hardware.

Instrument image files could be queued for downlink to the ground in real-time (via correct parameter in imaging command), or sent to the ground later (via a file downlink command). All images prior to Impact minus 5 hours were sent to the ground in real-time. These images were periodically sequenced for deletion prior to ground confirmation of receipt in order to clear NVM for new image acquisitions. Within 4.5 hours of impact, most science images were kept in NVM storage until after the playback phase completed and all images were confirmed received on the ground. NVM was nearly filled at the time imaging was completed at shield mode entry. The look back imaging, which occurred after closest approach, used the remainder of the available NVM. Figure 5 depicts encounter data flow characteristics for the flyby spacecraft.

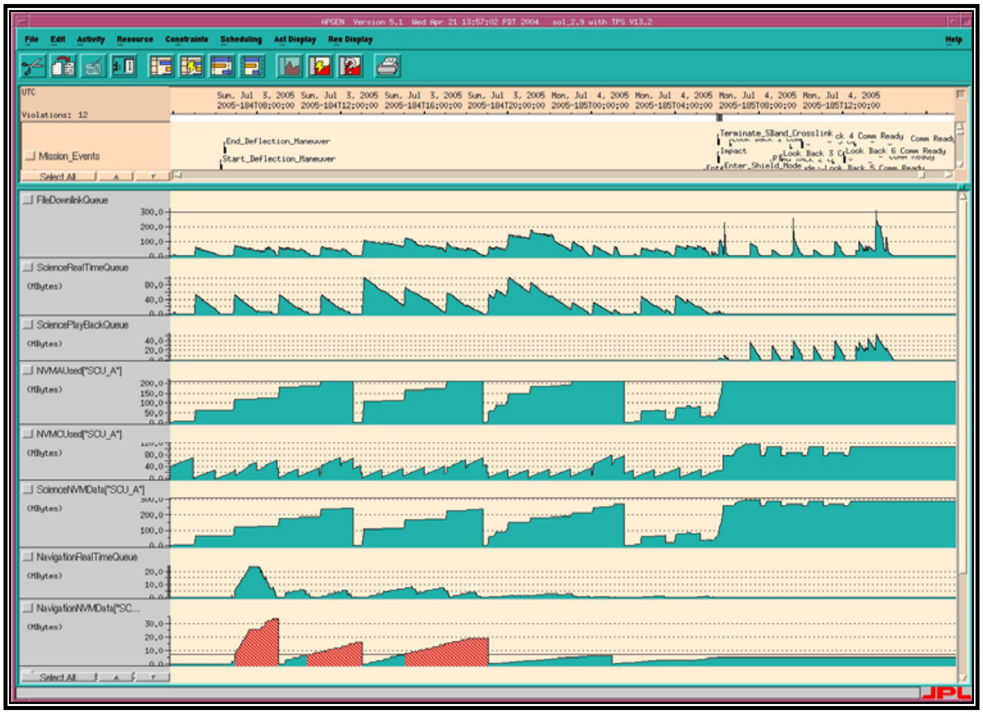

Figure 5: Encounter Data Flow Characteristics

C&DH throughput performance was a continual problem throughout most of the encounter test program, with critical impact and crator images lost on nearly every test run on the spacecraft pre-launch and testbed runs pre- and post-launch. The throughput performance fell off from predictions during periods of heavy imaging and file system usage. The performance was improved with Flight Software updates throughout the test program, and appeared adequate for flight in the final 2 months prior to encounter.

The Impactor spacecraft communicated with the Flyby spacecraft via an S-band radio. Commands and file transfers to the Impactor were processed on the ground into a file of bits to be radiated, then uplinked to the Flyby spacecraft as a file transfer to the Flyby's on-board file system. A command was then sent from the ground for the Flyby to send that file using the s-band link to the Impactor, which saw it as either a real-time command or a file transfer.

Telemetry data received on the Flyby from the Impactor S-Band was stored in 128 kB NVM files and automatically queued up for D/L. When the files were received on the ground, they were automatically read and processed into Impactor telemetry frames for real-time display and storage. The collection and storage of this data on the Flyby could be turned on or off by command. There was sufficient data allocation on the Flyby NVM to store 30 MB of Impactor telemetry. This equated to just over 1 hour of data at 64 Kbps. The Impactor data collection strategy was to turn on collection, collect data for 1 hour, turn off data collection and delete the files in NVM. This sequence was repeated until one hour before impact. The last hour was collected and remained in NVM post impact. This strategy ensured nearly continuous Impactor telemetry to the ground, but only stored the last hour of telemetry on the Flyby for post encounter playback. Any telemetry not received on the ground prior to 1 hour before impact was un-recoverable if initial transmission attempt failed. Due to descopes affecting downlink data rate, there was major contention between science image data and Impactor telemetry in the available downlink bandwidth. The sequences and downlink priorities had to be carefully orchestrated in order to avoid latency in either science images or Impactor telemetry.

Flyby S-Band RF reception was turned on at Separation plus 30 seconds. S-Band RF transmit was turned on at Separation plus 6 minutes to ensure the RF level was low enough to avoid damage to the Impactor S-Band receiver. There were 6076 Impactor telemetry files (about 760 MB) collected by the Flyby and relayed to the ground during the encounter critical sequence.

2.6 Flyby Encounter Imaging Campaign

The encounter imaging campaign consisted of a combination of 6012 science, 199 optical navigation and 936 autonomous navigation images. The transition between optical navigation and autonomous navigation occurred two hours before impact. Optical navigation images were used for determining the location of the comet relative to the spacecraft, in order to improve the Impactor delivery accuracy prior to the Impactor's autonomous navigation phase. The on-board autonomous navigation software took over 2 hours before impact and was used to keep the instruments pointed at the comet. The primary camera for AutoNav was the MRI, however, redundant AutoNav images were captured with the HRI as well. The redundant HRI images were discarded if MRI images were successfully acquired.

Early encounter science imaging (prior to E-2 hours), was primarily dedicated to periodic monitoring of the comet, definition of the rotation, and detection of the nucleus. The nucleus of the comet was resolved at around 82 hours prior to impact in the HRI and 16.5 hours prior to impact in the MRI.

Science and navigation shared use of the same instrument. This shared resource was a major driver in the complexity of the encounter sequence design, as there were conflicting requirements for when navigation and science images needed to be taken.

2.7 Autonomous Navigation

The autonomous navigation system for Deep Impact was designed to control pointing of the instruments on-board the Flyby. AutoNav also provided sequence timing updates required to support high-rate imaging at the time-of-impact (TOI) as well as initiate the transition to shield mode at the correct time. The system had two primary objectives: 1) Capture the impact event in both the MRI and HRI cameras and 2) track the impact site for at least 800 seconds following impact to monitor ejecta plume dynamics, crater formation, and capture the highest possible resolution images of the fully developed crater before shield attitude entry. AutoNav relied on the performance of and interaction with the Attitude Determination and Control System (ADCS) flight software and subsystems as well as the MRI camera, which served as both a Science instrument and the primary AutoNav camera.

The AutoNav software was developed for the New Millenium Deep Space 1 (DS1) mission[3,4]. AutoNav consists of three (3) distinct modules: 1) Image processing; 2) Orbit determination; and 3) Maneuver computation. On the Flyby, only the first two modules were used during the final 2 hrs of encounter with comet Tempel 1. 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 two CT-633 Star Trackers and SSIRU rate sensor provided 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 autonomous tracking process are as follows:

- Acquire MRI images of the comet nucleus, every 15 sec, starting 2 hrs before the expected time of impact
- Process MRI images to compute pixel/line location of the nucleus center of brightness (CB) using the Blobber algorithm for the first 60 min
- Use observed CB pixel/line locations to compute measurement residuals for comet-relative trajectory estimation
- Perform trajectory determination updates (OD), every 1 min, starting 1 hr 59 min before the expected TOI
- Produce and hold two (2) sequence time updates, with every OD update, to start the TOI imaging sequence and to control the timing of the final imaging sequence and entry into shield attitude for inner coma passage
- Perform Scene Analysis (SA), 4 min before impact using the MRI, to point instruments at the impact site
- Provide sequence time updates upon request (via command) from the critical sequence
- Provide comet-relative velocity information to align solar array edge in the direction of on-coming dust flow

2.7.1 Processing MRI Images

Image processing for the AutoNav system served the purpose of providing observations of the spacecraft comet-relative trajectory over time. When images were received by AutoNav, they could be processed in one of three ways: 1) Brightness centroiding of all pixels above a brightness threshold and within a predetermined pixel sub-region (Centroid Box); 2) Image blobbing to detect one or more contiguous regions of pixel brightness above a brightness threshold and compute the pixel/line location of each blob based on a simple moment algorithm (Blobber); or 3) Scene analysis to compute the pixel/line offset, relative to the CB location as determined from either 1) or 2) above, of the region most suitable for an illuminated impact.

Brightness centroiding is the process of determining the center of brightness using a moment algorithm. The MRI images consisted of an array of 1008x1008 pixels (when they arrived at AutoNav) with the brightness of each pixel expressed as a data number (DN). Based on the best estimate of the spacecraft trajectory relative to the nucleus and the MRI camera attitude at the time the image was taken, a predicted pixel/line location was computed. All pixels in an NxN pixel region surrounding this predicted location were used and the center of brightness determined by summing brightness values in the pixel and line direction. N was specified in the AutoNav image command parameter list. The brightness threshold could be a fixed value specified in the parameters file or based on a percentage of the peak pixel brightness. Deep Impact used N=400 to cover uncertainties in the ADCS estimated attitude and take 35% of the peak pixel brightness as the brightness cutoff. The brightness cutoff removed both instrument noise that is not gain dependent as well as signal from the dust or coma surrounding the nucleus.

The Blobber algorithm is well-known and makes use of brightness centroiding, but is different in that it first seeks and detects regions of contiguously lit pixels, then computes and returns the CB pixel/line location for each “blob” that is detected above the brightness threshold. This removed the dependence on AutoNav’s ability to predict where the nucleus will be located in the image array.

Scene Analysis was first developed by George Null and enhanced, specifically for the Deep Impact mission, by Nickolaos Mastrodemos at the Jet Propulsion Laboratory. Mastrodemos discovered that Scene Analysis, while improving the number of illuminated impacts, could drive the impact site to a region that is obscured from view as the Flyby passes beneath the nucleus depending on the orientation, size and shape of the nucleus. Mastrodemos’ version provided AutoNav with the ability to target a specific illuminated location on the nucleus in a biased fashion. The biasing drove the site selection to a location that was “biased” toward the Flyby’s point of closest approach, thus enhancing the ability of the Flyby to image the fully developed crater at the time of highest resolution. Here, images were first processed using either the Centroid Box approach or the Blobber algorithm to provide a reference CB location. Scene analysis was applied and the pixel/line location of the selected site returned. The difference between the CB reference location and the selected site was computed and converted to an inertial correction vector that was stored and passed to ADCS for pointing control. If no suitable site was found, then no correction vector was returned which resulted in a default to tracking the CB.

2.7.2 Trajectory Determination

For the Impactor, the trajectory was estimated and updated every minute during the last 2 hrs of encounter. The trajectory determination software supported both ADCS attitude control and sequence timing.

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

Following completion of each OD, the trajectory was updated and a representation of the estimated trajectory generated in the form of Chebyshev polynomial coefficients (Chebys). These Chebys represented a time series of the predicted Flyby positions relative to the nucleus and were passed to the ADCS flight software for evaluation and pointing control. ADCS evaluated the Chebys, computed the relative position vector, applied the SA offset and aligned the HRI camera boresights with the comet-relative position vector to center the nucleus in the instrument FOV.

2.8 Robustness

Several flight software features were used for encounter making it a robust, fault tolerant activity. The key fault protection design impact was in the sequence architecture. Mark and Rollback is a feature of the FSW Command Manager CSC which allows for an absolute-timed critical sequence to be restarted automatically after a fault has occurred. Each command in a critical sequence is a mark point. Once the next command in a sequence has been executed by Command Manager, the sequence cannot roll back to the previous command. Each command in a critical sequence can have a timeout window associated with it. This defines the extent of time over which a rollback can occur. Once this timeout window has expired, the sequence will roll forward and wait for the next command in the sequence.

The critical sequence for encounter was an absolute-timed sequence and contained spawns and calls to sub-sequences. These sub-sequences were relative-timed and contained spawns to absolute and relative-timed sequences containing the spacecraft and instrument commanding. The spawns/calls in the critical sequence to the sub-sequences had timeout windows which expired at the next command in the critical sequence. Of primary importance in critical sequence design was the selection of the mark points. Mark points were chosen to occur at the completion of a critical event, such as Separation, so that in the event of a fault, an important activity would be completed prior to rolling forward to the next task. In addition, some mark points were chosen to accommodate contingency plans, ie. swapping in alternate imaging sequences in the face of unexpected cometary environments.

The strategy for autonomous fault protection during encounter was challenging due to the time critical and unknown environment aspects of the event. At the start of the critical encounter phase, the set of FP enabled was similar to the cruise setting, with additional instrument, attitude control (point state), and telecom (HGA) recovery enabled. As impact time approached, more and more fault protection was disabled to lessen the chance of interference with encounter. At impact minus two hours, spacecraft system safing was disabled (recovery from safing would not have allowed recovery in time for successful encounter operations after this time), leaving system responses (major hardware swaps, basic hardware protections, etc.) enabled. At impact minus seven minutes all system responses were disabled. From that time until impact, fault protection was in place for only low level hardware protection, and processor reboots.

In addition to on-board autonomous fault protection, several contingency plans were prepared in advance of the encounter activity. All sequences to support these plans were loaded into on-board memory with the nominal encounter sequence upload. The major encounter contingency plans were as follows[5]:

- ☐ Unexpected Comet Albedo – If Tempel 1 was significantly brighter or more dim than predicted, alternate imaging sequences would be executed (exposure durations were dependant on comet albedo).
- ☐ Late Release – If in preparation for separation the Impactor Spacecraft experienced anomaly, the encounter timeline would be adjusted to perform separation at E-10 hours instead of the nominal E-24 hours
- ☐ Flight System Impact– If the Flyby and Impactor were unable to separate, this contingency plan would execute encounter with both craft still joined.

3.0 COMET ENCOUNTER TESTING

Developing the Tempel 1 encounter activity presented a complex, multi-disciplinary set of challenges. Following the design phase of encounter, it was necessary to develop and execute a test program that would verify and validate the design. The encounter test program was developed to test, in a flight-like manner, all aspects of the encounter activity. This included nominal sequence performance testing, off-nominal faulted testing, stress testing, and operations simulation testing. The program spanned two years, and had elements that were executed on all test venues, including the software test benches, the flight system test bench, and the flight system itself.

3.1 Encounter Test Program Objectives

Encounter testing, in line with the overall project Verification and Validation program, was focused on the following major ‘readiness’ objectives:

- Ensure that the system is robust under all expected encounter operating conditions
- Ensure that normal encounter operation does not trigger Fault Protection
- Ensure that system responses are predictable, and the results are acceptable
- Ensure that the system operates in off-nominal (within-specification) encounter conditions
- Push the boundary conditions of the system to investigate the ability to operate outside of specification, and understand when overstress occurs and how it is handled

3.2 Encounter Test Program Architecture

Each test conducted in this program was categorized as a Mission Readiness Test (MRT). The set of MRTs was composed of several types of system-level tests, defined below:

- Mission Scenario Test (MST): MSTs are tests using flight-like, MOS/GDS-generated sequences in tests targeted to verify and validate key mission scenarios (such as encounter) that will be executed in flight. MSTs are intended to provide flight-like operations to validate FSW design for the DI mission. All critical encounter activities, including Separation/Divert and Final Imaging/Impact, were tested on a software test bench (SWTB), a flight system test bench (FSTB), and the flight system prior to execution on the flying spacecraft.
- Faulted Mission Scenario Test (Faulted MST): Faulted MSTs are tests that inject specific faults into a nominal MST run, and let the system respond. The objective of these tests is to prove appropriate system response to known off-nominal conditions, and to validate system recovery in the presence of faults.
- Robustness Test: Robustness tests were intended to vary key performance parameters (ie. dust model, center of mass, etc.) by ±3σ and ±6σ to validate system robustness, and successful comet encounter in the event of stressors outside of specified performance.
- Operation Readiness Test (ORT): An Operational Readiness Test uses flight-like, MOS/GDS-generated sequences/tools/processes in tests targeted to exercise system performance while training the flight team. Encounter ORTs were conducted in the flight Mission Support Areas with full flight team participation, including Flyby and Impactor teams for the individual spacecraft. ORTs are the ‘rehearsals’ before the in-flight events.
- Incompressible Encounter Test: Incompressible encounter tests are a subset of the full test set. These tests are designated as mandatory, meaning that the in-flight event would be waived off if even one incompressible test was not successfully executed. In the case of Deep Impact Encounter, there was not an opportunity to waive off the event and try again later. The incompressible encounter tests included several *key* MSTs, Faulted MSTs, Robustness Tests, and ORTs. All were successfully executed prior to encounter.

Over the span of two years, many Flyby encounter MRTs were executed. The start of the test program included basic Flyby encounter MSTs to begin sequence validation and verification of nominal performance requirements. The Flyby encounter MSTs were of three different varieties. First, the early portion of encounter was broken out into an individual test, Flyby Encounter: Separation/Divert, which started prior to Flyby/Impactor Separation and ended after the Flyby Divert Burn. Next, the late portion of encounter called Flyby Encounter: Final Imaging MST, which began at two hours prior to impact and ended with the lookback/playback cycles post shield mode. The most comprehensive Flyby encounter MST was the Full Encounter case, beginning prior to separation and extending to the end of lookback/playback. These MSTs were executed in excess of 200 times on the test benches, and ran on the flight system prior to launch in support of final C&DH subsystem V&V. After the nominal encounter cases achieved stability, the test program continued with the off-nominal Faulted MSTs and Robustness Tests. ORTs were conducted every few months (and more frequently post-launch) to train the team for the encounter event in flight. An incompressible test list (ITL) was generated for the encounter event. There were a few of each type of MRT on this list to ensure complete coverage of nominal and off-nominal conditions in preparation for encounter—the mandatory nature of this set of tests enabled full verification and validation prior to execution in flight. Table 3 shows a summary for the numerous test cases run for Flyby encounter.

Table 3: Flyby Encounter Test Summary

TEST TYPE	NUMBER OF CASES	TOTAL # EXECUTED	# ON ITL	TOTAL HOURS
MSTs	3	279	3	5160
Faulted MSTs	16	25	8	300
Robustness Tests	11	35	11	288
ORTs	7	7	5	400
Totals	**37**	**346**	**27**	**6148**

3.3 Key Encounter Testing Challenges

Conducting the comprehensive test program for encounter was complicated by several major challenges.

The Encounter Approach sequences, spanning more than seven days, were far too large to be tested in their entirety. A combination of extensive data flow modeling/simulation and test bench execution of critical parts of the sequence proved readiness for the in-flight event. Taking more than 24 hours for each test execution, and still only testing a small fraction of the full set of sequences, this encounter-enabling set of events was tested the least of any critical Deep Impact event.

The majority of encounter tests took a minimum of 12 hours to execute. Most took in excess of 24 hours to run. With a small test team, and many runs of each test case, this guaranteed numerous executions per week per tester. The rapid nature of encounter testing required precise test procedures, excellent documentation of test data, and rapid review/revision prior to the subsequent run. The volume of test cases attempted in the short time between launch and encounter resulted in overstress of the test venue schedules, test personnel, and subsystem reviewers.

Sheer complexity of the encounter event, likened to hitting a bullet with another bullet, was one of the biggest hurdles. System performance was stretched in almost every subsystem: images were taken at maximum speed, NVM was filled to the maximum, ADCS and AutoNav processing was at its peak, and telemetry bandwidth was maximized. These factors, coupled with the unknowns of the Tempel 1 environment (albedo, dust environment, size of the nucleus, rotation rate, etc.), and precise timing synchronization between both craft, demanded correct settings for thousands of parameters. Any one of those parameters, if wrong, could have compromised the mission.

3.4 Performance Predictions

Upon successful completion of the Encounter Incompressible Tests, the flight team expected good performance from the flight system. Encounter Approach was predicted to successfully capture and downlink prior to deletion 1769 navigation images, 31903 ISI files, and 2616 science images, including those taken during the mission critical mini calibrations. A few navigation images (4) and Impactor ISI files (11) were predicted to be deleted prior to downlink to preserve the timely start of the mini calibrations. Bandwidth usage expectations were favorable for the entire week of continuous flyby and Impactor operations, showing no dropped image packets or interruptions in telemetry. Science NVM usage was predicted to fit within the 309 MB allocation, and the downlink queue never exceeded the 300 file maximum in the final modeled run. Figure 6 shows the fully modeled encounter approach predicted data flow performance.

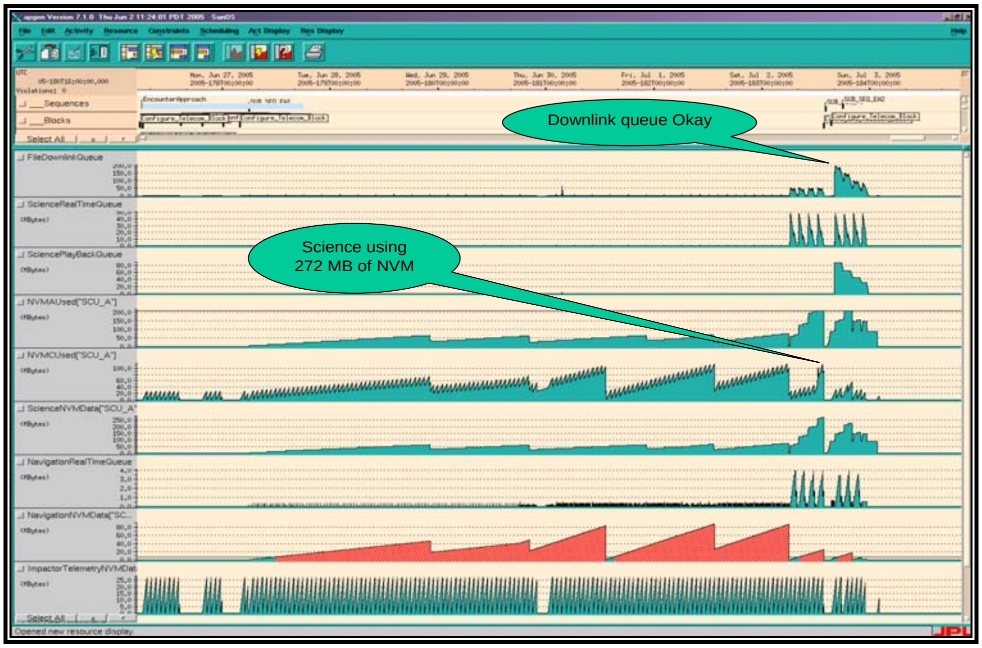

Figure 6: Encounter Approach Predicted Data Flow Performance

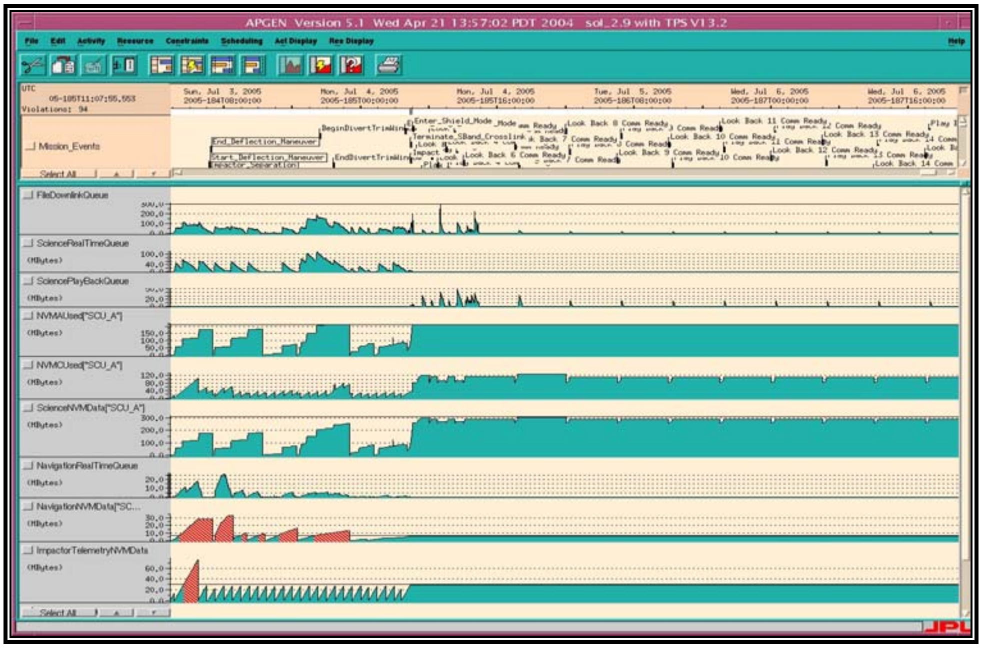

Figure 7: Flyby Encounter Predicted Data Flow Performance

The critical encounter sequence was predicted to also perform flawlessly. Of the 6012 science, 199 optical navigation and 936 autonomous navigation images taken during the encounter sequence, predictions indicated all would successfully reach the ground prior to deletion. All Impactor images, including those closest to impact time (~3 seconds from impact) were also expected to reach the ground successfully. The "live for the moment" strategy to get all high priority images down prior to closest approach of the flyby to Tempel 1 was also predicted to be executed without incident. Science NVM usage was expected to fall below the 309 MB allocation, and the downlink queue was not predicted to exceed the 300 file maximum. During ORTs, executed on the test benches, there were significant delays and losses of Impactor telemetry during critical Impactor AutoNav operations and

Impactor ITMs. Modeled runs also predicted some Impactor data loss, but modeling showed outages that represented less than three minutes of Impactor telemetry. The performance of the Command and Data Handling Subsystem (C&DH) on the test benches was an issue throughout testing of the encounter sequences, causing loss of critical science images of impact and the crater formation. Images were essentially being taken faster than the C&DH test bench hardware could process/store them. Modeled runs did not echo this poor performance in C&DH—dropped images were seen only on the test benches. Final performance predictions were based on the modeled runs due to the differences in performance between the flight system and the test benches. Figure 7 shows the final predicted performance for encounter data flow.

Final test executions showed excellent performance of all flight systems throughout encounter. TCM test runs were in line with earlier in-flight performance, and showed that TCM 5 should execute perfectly. ADCS and AutoNav were predicted to successfully execute their required functions, including delivering the Impactor onto an impact trajectory, and choosing an impact spot that the flyby could image at impact. The final runs of incompressible tests all resulted in Impactor "hits," even in the case of some off-nominal performance.

In-flight demonstrations were also an indicator of future performance at encounter. The Saturn Watch test, executed on the flight system during cruise, showed outstanding AutoNav performance in tracking Saturn and picking an impact location on that planet. The in-flight demonstration performance was as expected for the real Tempel 1 event. In addition, the in-flight encounter demonstration was quite successful. This demonstration executed the hardest part of the Flyby's sequence, including a demanding slew just after the impact images were captured, and demonstrated the end-to-end data flow from spacecraft to data archive. This demonstration included cooperation with the Deep Space Network to test the array of antennas schedule for encounter. The demonstration went very well—all test objectives were met. All test venues, from software simulations to the real spacecraft, showed predictions of a successful encounter with Tempel 1.

4.0 ENCOUNTER EXECUTION & PERFORMANCE

The stunning images seen around the world on July 4, 2005 proved without a doubt that Deep Impact was a success. The following sections detail the flight system performance of the flyby spacecraft for this historic event.

4.1 Encounter Approach Performance

The week prior to encounter was filled with intense imaging, maneuvers, and calibrations. Figure 8 shows representative science images taken during the Encounter Approach phase. The images, from left to right, span 4 days of approach. The last image was taken July 1, the flight system was 2,446,529 km and less than three days out from Tempel 1.

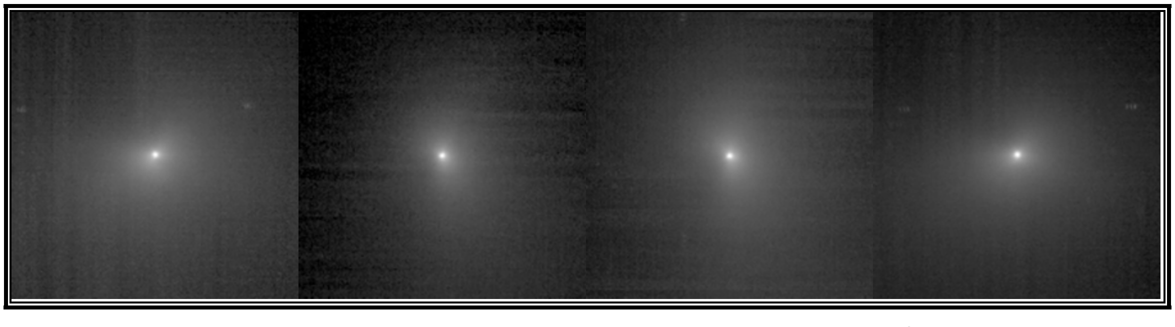

Figure 8: Encounter Approach MRI Clear Filter Images[6]

Of the 4389 science and navigation images taken during encounter approach, 127 did not reach the ground. All 127 images lost did not impact mission success. The lost images were deleted prior to downlink by the on-board sequence. It was found that a single over-subscription in real-time telemetry was the cause for the bandwidth overflow. This over-subscription fully accounts for the mismatch in performance to predictions.

The mini-calibrations performed for the science team to use for deconvolution of the HRI images could not have been more successful. The beautiful, high-resolution HRI images shown in the following sections are a direct result of the rapid turn-around from the Point Spread Function and Infrared Mini Calibrations.

Throughout the Deep Impact mission, Trajectory Correction Maneuvers were quite accurate. TCM 5 executed within 2% of desired performance. This placed the flight system with 1 km of an impact trajectory. The flight system encounter activity was designed to accommodate up to 30km of Impactor delivery error, so the 1km performance was outstanding.

The final responsibility of encounter approach was to clean up NVM for the critical encounter event. At the completion of the encounter approach sequence, all drives were clean saving only the files required to execute encounter. All performance metrics for each subsystem, the flight team, and the science team, were met or exceeded for the final week prior to impact. This week was a fitting prelude to the spectacular comet encounter.

4.2 Separation/Divert

The separation sequence executed nearly flawlessly. There was an excessive attitude error fault protection symptom raised at Impactor separation. Fortunately, fault responses were still disabled at this point, so no fault response ran. The error was due to a change in the timing of re-enabling attitude control after the separation event, and a lack of fidelity in testbed simulation to trigger the response during testing. The divert sequence also ran flawlessly. The divert performance resulted in approximately 1 second timing error in the time of impact, and 2 km error in closest approach distance.

4.3 Autonomous Navigation

The AutoNav system performed beyond expectations. Figure 9 is a compilation of images taken with both the MRI and HRI instruments at various times during AutoNav operations.

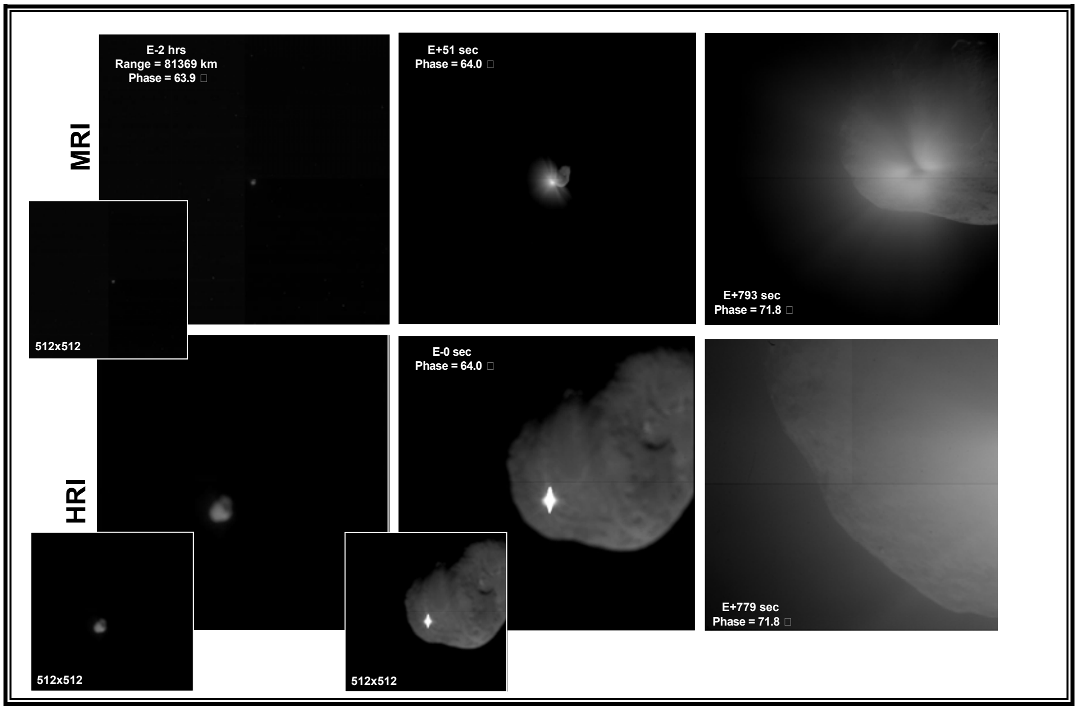

Figure 9: Set of MRI (top) and HRI (bottom) images taking during AutoNav operations[6]

All of the primary Flyby AutoNav objectives can be summarized in two images. Figure 10 shows the HRI image capture at the moment of impact. The saturated pixels are located within 31 pixels of the center of the HRI FOV. Each HRI pixel is 2 microradians and therefore indicates that the total of all error sources was less than 62 microradians at TOI. This includes errors associated with: 1) AutoNav estimation; 2) ADCS attitude estimation; 3) HRI camera alignment errors; 4) Independent Scene Analysis induced errors; and 5) Errors associated with updated estimates of TOI.

Figure 11 shows a full-frame MRI image of the impact site just 7 sec before shield attitude entry. While the impact site was lost for the HRI FOV nearly 13 sec earlier, the 3 picture mosaic (scan) that was implemented to increase the probability of capturing the crater in at least one of the last 3 images taken before shield mode, did just that and has given the Science Team an opportunity to quantify the size of the fully developed crater.

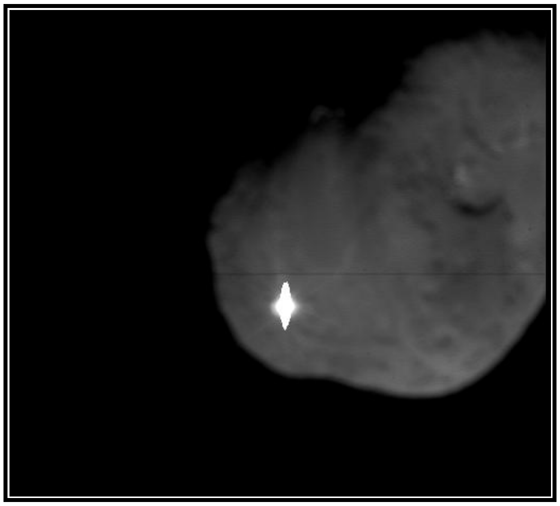

Figure 10: HRI 512x512 pixel subframe image showing initial flash at the moment of impact[6]

Figure 11: MRI Full-frame image of the impact site (center of ejecta cone) at 793 sec after impact[6]

4.4 Encounter Imaging and Data Return

The Flyby encounter Imaging campaign performed exceptionally well. Of the 6012 science images sequenced during the encounter phase, all of the pre-shield mode images taken after Impact minus 15 hours were received on the ground. Approximately 20 images from before Impact minus 15 hours were lost due to a failure in the DSN transmitter at one 70 meter ground station. The drop in gain from using a 34 meter station for uplink caused the loss of some transfer frames containing image file data. The downlink Signal-to-Noise Ratio (SNR) was improved by turning ranging modulation off at the 34 meter station. None of the critical impact or crater images were lost due to C&DH performance. All but 66 of the over 6000 Impactor telemetry files were received without error on the ground and processed into Impactor telemetry displays. 63 (representing less than 15 minutes of Impactor telemetry) of these occurred early in the encounter period during the DSN transmitter failure. All critical Impactor telemetry was delivered to the Flight Team in a timely manner. Figure 12 shows encounter data flow performance.

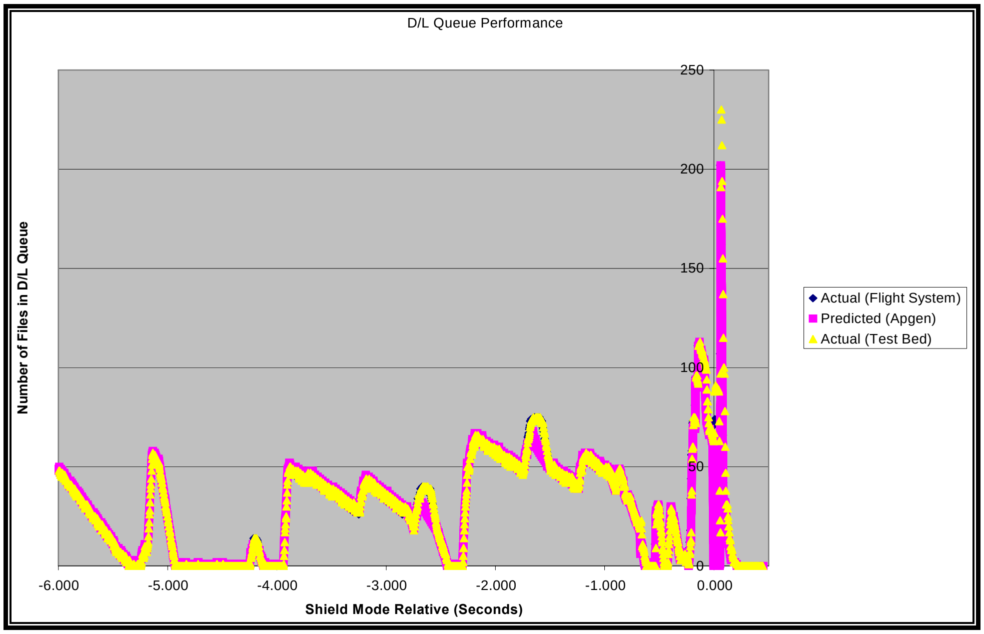

Figure 12: Encounter Data Flow Performance

Images downlinked from the Flyby during encounter were viewed in near real-time in the Mission Support Area. In addition to the engineering telemetry being tracked by each member of the flight team, the pictures of Tempel 1 just kept getting better and better. The well-scripted encounter event was captured in stunning detail—these images show the undeniable success of Deep Impact's date with comet Tempel 1.

Figure 13: HRI Image of Tempel 1, 67 Seconds After Impact[6]

This spectacular image of comet Tempel 1 (Figure 13) was taken 67 seconds after it obliterated Deep Impact's Impactor spacecraft. Scattered light from the collision saturated the camera's detector, creating the bright splash seen here. Linear spokes of light radiate away from the impact site, while reflected sunlight illuminates most of the comet surface. The image reveals topographic features, including ridges, scalloped edges and possibly impact craters formed long ago.

Figure 14: HRI Image of Tempel 1, False Color, 50 Minutes After Impact[6]

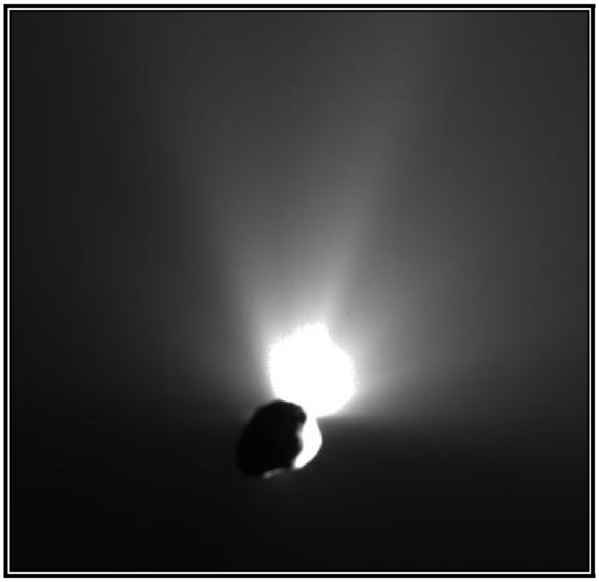

Figure 15: HRI Image of Tempel 1, 50 Minutes After Impact[6]

Figure 14 shows comet Tempel 1 about 50 minutes after Deep Impact's probe smashed into its surface. The impact site is located on the far side of the comet in this view. The image was taken by the mission's flyby spacecraft as it turned back to face the comet for one last photo opportunity. The colors represent brightness, with white indicating the brightest materials and black showing the faintest materials. This brightness is a measure of reflected sunlight. Because the sunlit portion of the comet is brighter, it appears white. The comet's nucleus is silhouetted against the light reflected from surrounding dust. The large plume of dust that was kicked up upon impact can be seen as the colorful, drop-shaped object. This plume was very bright, indicating that the comet's surface material must be very fine, like talcum powder. The blue speck in the upper left corner is a star.

Figure 15 shows the view from Deep Impact's flyby spacecraft as it turned back to look at comet Tempel 1. Fifty minutes earlier, the spacecraft's probe had been run over by the comet. That collision kicked up plumes of ejected material, seen here streaming away from the back side of the comet.[7]

4.5 Timing

One of the biggest concerns about the unknown trajectory of Tempel 1 was the uncertainty in the time of impact and its effect on sequence timing. The original uncertainty estimate at the time the sequences were generated was 3 minutes. The timing updates generated by AutoNav for updates to the impact and high-resolution crater imaging sequences were on the order of 1 second on the day of encounter. The Navigation Team did an exceptional job determining Tempel 1's trajectory early enough to avoid large timing uncertainties at encounter.

4.6 Tempel 1 Environment

Throughout design, development, and flight, the question repeated most often was, "what's it like at Tempel 1?" No mission prior to Deep Impact experienced the environment at this particular comet, yet two small spacecraft were destined to come closer than any had before. In the final days leading up to encounter, Tempel 1 was a bit more active than anticipated, with several "outburst" events that could have had implications on attitude control and AutoNav performance at encounter. The major elements of environmental concern were the dust environment and comet brightness (albedo).

The flyby spacecraft was designed to accommodate the science team's "best guess" for the dust environment at Tempel 1. The shielding used when the spacecraft was at its closest approach to the comet could protect the flyby from dust particles up to 2 mm in diameter. Despite concerns about potentially damaging larger dust particles, on the day of encounter, the flyby spacecraft navigated the dust field without mishap. No major damage was sustained, and the flyby craft went on to produce the full set of planned look back images. Post encounter, the Flyby successfully executed a final TCM to place it on a trajectory that could support an extended mission if approved.

One of the major contingency plans developed for encounter included alternate sets of imaging sequences in the event that Tempel 1 was significantly brighter or more dim than anticipated. Several key decision meetings in the hours prior to final imaging, gave the science and navigation teams the opportunity to use the alternate sequences. For the Navigation team, the brightness of the comet had a direct impact on the autonomous navigation system to correctly compute the impact location. For the scientists, the images needed to have appropriate exposure durations to capture the detail required for mission success—exposure duration was highly dependent on comet albedo. On July 3-4, 2005, comet Tempel 1 had albedo in the nominal range expected for encounter. The flyby and Impactor spacecraft executed the baseline imaging sequences, yielding valuable images for the science team and nominal conditions for navigation.

5.0 SUMMARY

Despite engineering challenges, an unknown and potentially dangerous comet environment, and a short operations phase, Deep Impact was a dramatic success. This Discovery mission spent Independence Day with the world watching—what a sight! Figure 16 shows the first definitive impact image transmitted in near real-time, and displayed in the Mission Support Area at JPL. The 15 minute encounter with Tempel 1 was viewed by astronomers across the globe and space-based observatories.

Following years of complicated design work and a rigorous test program, the virtually flawless performance at encounter was well-earned. All systems performed at or beyond expectations, resulting in the purposeful obliteration of one small spacecraft, and images that will unlock secrets of the solar system so long held inside comets.

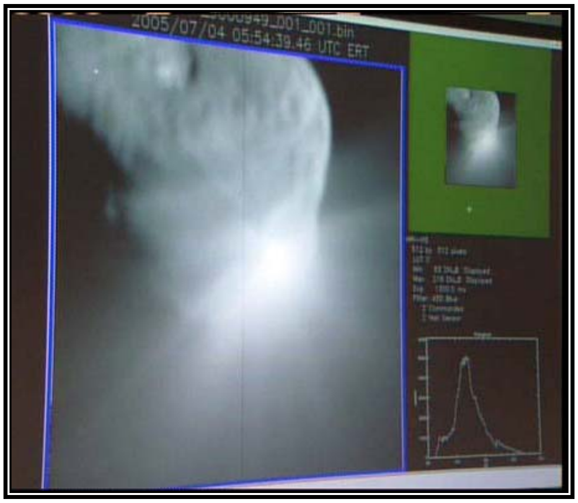

Figure 16: The Moment We Knew[6]

Figure 17: Encounter Flight Team[8]

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] Blume, W.H, and Wang, K.C, "Deep Impact Project Mission Plan," 2004.

[2] Wissler, S, "Deep Impact Flyby Spacecraft Encounter Design," 2005.

[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.

[5] Sierchio, M, et al, "Deep Impact Encounter Contingency Plan Notebook," 2005.

[6] Image Credit: NASA/JPL-Caltech/UMD

[7] Deep Impact Web Site, http://deepimpact.jpl.nasa.gov

[8] Image Credit: Reuters