(Preprint) AAS 11-217

OPTICAL NAVIGATION FOR THE EPOXI MISSION

Brian P. Rush, William M. Owen, Jr., Shyam Bhaskaran
and Stephen P. Synnott[*]

The Deep Impact spacecraft flew by comet Hartley 2 on November 4, 2010 as part of its extended mission called EPOXI. Successful navigation depended critically on the quality and timing of optical navigation data processing, since pictures of the comet provided the most precise comet-relative position of the spacecraft. This paper describes the planning, including the picture timing and pointing; the methods used to determine the center of the comet image in each picture; and the optical navigation results, which provided the necessary information to allow the cameras to accurately target the comet for science imaging at encounter.

INTRODUCTION

The EPOXI mission is the extended mission using the Deep Impact spacecraft, consisting of two investigations: EPOCh, or Extra-solar Planet Observation and Characterization, and DIXI, or Deep-Impact eXtended Investigation. The latter consists of a flyby and observations of the Hartley 2 comet, and the optical navigation planning and operations required to make that a success are described in this paper.

Optical navigation as used on the EPOXI mission has proven to be extremely helpful, if not absolutely necessary for virtually all high-precision small-body mission in recent years. In this process, the use of an onboard camera allows narrow-field astrometry to be applied to spacecraft orbit determination. Pictures are taken against a background of reference stars which are used to calculate the attitude of the spacecraft at the time of each picture. The location of the target body in the picture is then used, along with the spacecraft's heliocentric trajectory as determined from radio navigation[†] and the comet's *a priori* heliocentric ephemeris as determined by Earth-based astrometry, to calculate the apparent inertial coordinates of the target body as seen by the spacecraft. This provides the only direct measurement of the relative position of the target body and spacecraft.

For EPOXI, this knowledge was then used to precisely orient the spacecraft during the encounter to position the comet in the field of view of the cameras in such a way as to provide maximum science return. The opnav imaging used in this process began at 60 days before en-

---

[*] Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena CA 91109-8099.

[†] Several references are made in this paper to other work done by the EPOXI navigation team, including radio navigation analysis, orbit determination, comet-ephemeris generation, and autonomous navigation (Autonav). Publications describing each of these areas will be presented in the near future, but no specific references are available at this time.

counter, and were processed on the ground. During the final 26 hours before encounter, the opnav process was automated onboard. This paper concentrates on the ground operations.

THE CAMERA

The Deep Impact spacecraft carries two cameras, the Medium Resolution Imager (MRI) and the High Resolution Imager (HRI).[1] The HRI has two detectors, one for visual and one for infrared wavelengths. The HRI-VIS was originally intended to be the primary camera for optical navigation, but the first post-launch tests showed it to be significantly out of focus,[2] which would smear the dim stars that were to be used for reference. Therefore the MRI was used, both during the original Deep Impact mission and for EPOXI. As was the case for the Deep Impact mission, the HRI-VIS was available as a backup in case the MRI failed or had other serious problems, but fortunately that was never necessary.

The MRI has a 12 cm diameter primary mirror and a nominal focal length of 2.1 meters. The detector is a 1024x1024 framing CCD with 12 micro radian pixels, giving a resolution of about 9.7 micro radians per pixel and a field of view of about 0.57 degrees. The detector carries four read amplifiers, one in each corner of the chip, to allow the four quadrants of an image to be read in parallel. Each quadrant has its own electronics and therefore its own bias level, gain, and read noise.

Camera Model

The observables for optical navigation are the measured sample (column) and line (row) coordinates of an image. The expected values are determined by first calculating the apparent position of a target from the camera, corrected for light time and stellar aberration, and the midpoint of the exposure. That vector is then rotated into a coordinate system that allows for projection into the nominal focal plane of the camera. Distortions and misalignments in the camera optics are then accounted for, as are geometric distortions in the detector. Finally an offset from the optical axis, taken as the geometric center of the camera, is added, giving the sample and line coordinates of the image in units of pixels, such that $(s,l) = (1.0,1.0)$ corresponds to the center of the upper left pixel. Further details and equations describing the mathematical model are given by Owen[3] in these proceedings.

Camera Calibration

Three types of calibration activities are generally needed to be able to accurately analyze opnav pictures. First, one needs to know the distortions and misalignments in the detector and optics, as mentioned above. This is done be taking detailed calibration pictures of star fields which place star images across the entire FOV, and then using these images to solve for the various geometric parameters as components of the camera model. Second, pictures are taken of set of standard stars of various colors to determine the photometric calibration of the detector. These are used to plan exposure times for opnav observations, and are of use in scientific analysis of the comet images. Finally, the alignment between the camera and the spacecraft must be determined, by taking pictures of star fields and comparing the opnav-determined spacecraft attitude with that determined by ADCS.

All of these calibrations -- geometric, photometric, and alignment -- were done several times during the Deep Impact mission, and were determined to be extremely stable, such that no changes were needed during the EPOXI mission. We did test the camera-to-spacecraft alignment a couple times, but determined that it had not changed within the margin of error, and we were thus confident that we knew the direction of the camera boresight in spacecraft coordinates to better than one pixel, and the (less important) twist angle of the camera to better than one milliradian.

PICTURE PLANNING

In planning the opnav picture schedule, we needed to know the number and frequency of pictures that would provide us with an orbit determination solution for encounter that met project requirements. The ultimate application of the orbit determination solution was to orient the spacecraft during encounter for the best possible view of the comet near the time of closest approach. This was not as strict as for Deep Impact,[4,5] where the requirement was to position the spacecraft to deliver the impactor to hit the comet, and then also to get the best possible view near closest approach. Thus, we did not need nearly as aggressive an opnav campaign for EPOXI. On Deep impact, we eventually increased the picture frequency to an average one picture every three minutes for the final ~4 days, while for EPOXI, the final schedule was one picture per hour.

Covariance analyses were done, combining simulated optical and radio data of various amounts to arrive at the final decision regarding opnav frequency. We also took into account possible data losses from cosmic rays corrupting star or comet images, outages at the tracking station, and possible hardware glitches or pointing errors. The schedule we settled on was to take (1) one picture every 6 hours from 60 days to 50 days before encounter (E-60 d to E-50 d), (2) one picture every 2 hours from E-50 d to E-40 d, (3) none from E-40 d to E-34 d, to avoid spacecraft orientations that would heat up the HRI-IR detector during a critical time, (4) one picture per hour for 16 hours a day from E-34 d to E-8 d, and (5) one picture per hour continuously from E-8 d to E-26 hours. From this schedule, we lost only 16 pictures during a DSN problem around E-28 d, and 6 pictures during a planned break around our final TCM. We thus ended up with 712 opnav pictures, all of which were delivered to the ground uncorrupted. A few were found to be unusable because the image of the comet fell directly on top of a background star, but we still had about 690 usable pictures.

These opnav pictures were taken using the same Autonav software used for the Autonav imaging during the final hours of encounter. Several imaging parameters were determined and uploaded in advance. Two such parameters are the number and size of “snip” boxes, sub frames of the full FOV that are downloaded to minimize data volume. We requested a maximum of 6 snips, one for the comet and 5 for the brightest stars, each 250 by 250 pixels in size (see Figure 1). This would ensure at least 2-3 star boxes in each picture, since a few could be lost when the snipping algorithm looks for the brightest stars in a circle circumscribing the square FOV, and those outside of the actual FOV are dropped. In practice, the comet box usually includes several usable stars, more than enough to determine pointing; but, by selecting boxes around several stars across the FOV, we establish a larger baseline for determining the pointing.

The requested inertial pointing for each picture was determined by examining the star field behind the comet. Fortunately, the relative spacecraft and comet trajectories were such that there was always a large number of stars in the FOV, and thus we avoided the tedious process required for Deep Impact of fine-tuning the pointing every day to capture enough stars. We could simply point the camera at the comet and be guaranteed to have enough stars in the picture. We offset this pointing slightly, to make sure that the comet did not fall on the exact center of the FOV, since that is the boundary of the four quadrants of the detector, and falling exactly at that point would slightly degrade the results of the comet centroiding process (see below).

Exposure time calculations were also fairly simple, as compared to Deep Impact, where the exposure times were shorted over a dozen times over the final week to avoid saturating the comet. Using early estimates of the brightness of Hartley 2, we began the opnav campaign with exposure times of 10 seconds – long enough to detect the comet, but short enough to avoid saturating too many star images. As we obtained more pictures over the following weeks, it was determined that the comet was faint enough that 10 second exposures would not saturate it even up to the last opnav taken at E-26 hours.

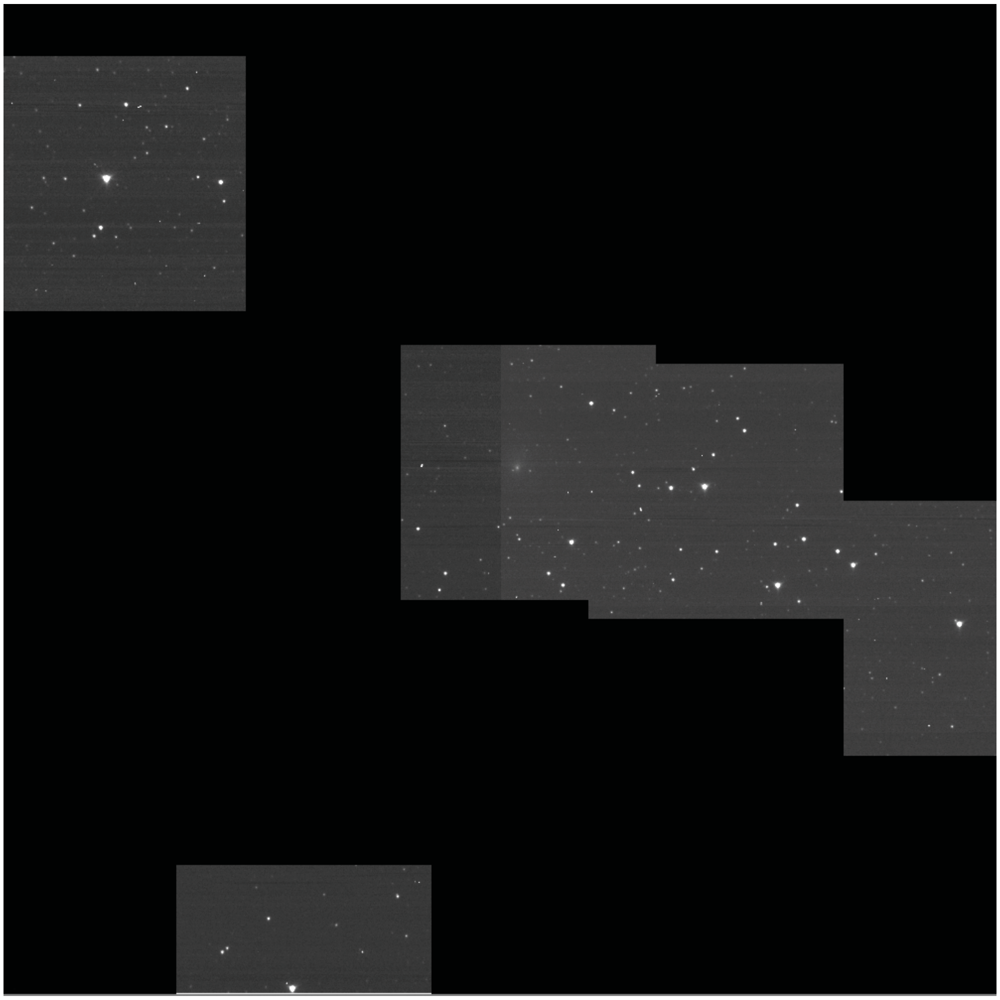

Figure 1. Opnav Picture showing 6 250x250 "snip" boxes

OPNAV DATA PROCESSING

After downlinking the pictures and converting them to a usable file format, the first step in the data analysis was to register the pictures, i.e., to determine an initial guess for the orientation of the field of view by comparing all star and comet images in the picture with the predictions of a star catalog and the input comet ephemeris. At first we did this manually, by displaying each picture, and then using a graphic overlay of predicted image locations based on the assumed pointing from ADCS, and then shifting the overlay until the predicted and observed images lined up.

We then moved to an automated procedure, which was especially useful during the last few days when we had to provide solutions on a tighter time scale. After comparing the results of manual and automatic registration for several sets of pictures early on, we were confident that the automatic process was more than sufficient. This automated procedure calculates the centroid of every local maximum that it finds in the picture, and compares this to the predicted locations of the comet and all stars that are expected to be in the FOV. The primary product of this procedure was an updated set of pointing angles for each picture, which was used to update the predicted image locations that were used as a starting point for the image centroiding described next.

Comet Centroiding

Determining the true center of the comet images was a difficult task. While it is easy to determine the center of brightness of the pixels which constituted the comet in each picture, it is not necessarily true that that point represents the 2-dimensional location in the picture of the center of the physical comet nucleus that we are tracking. That nucleus is surrounded by a coma that can vary in strength and shape depending on solar illumination, and the nucleus itself may not be near spherical (as we later found out), so even if we did resolve it clearly the center of mass would not be straight forward to determine. We also knew that we would not be able to resolve the comet nucleus until very late in the mission, after the last opnavs were taken. Figure 2 shows pictures of the comet: one picture of a typical opnav, at about 27 days before encounter when the scale was 271 km per pixel; and one picture of the last approach opnav, taken at about 19 hours before encounter, when the scale was still about 11 km per pixel, or about 5 times the size of the nucleus.

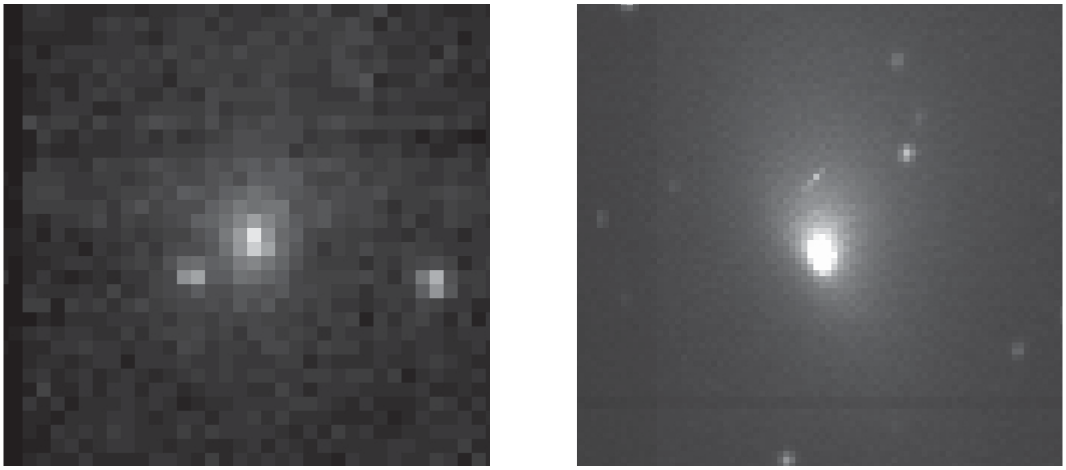

Figure 2. Opnav images, at E-27 days (left) and E-19.5 hours

We had to figure out a way to measure these brightness distributions to determine the most likely location of the center of the nucleus. To do so, we used six centroiding algorithms, and decided to use the one which gave the cleanest post-fit residuals after the final filtering process where we solved for the spacecraft and comet locations. Three of these were simple 2-D moment algorithms, determining the center of brightness in a 9 x 9 box centered on the predicted comet location after subtracting a constant background, limiting the pixels included in the calculation to those with at least 50 %, 75 %, and 99 % of the counts in the maximum pixel, to avoid both noisy background counts and off-nucleus coma. The latter of these methods simply equates to taking the center of the brightest pixel – a very simple but surprisingly acceptable way of determining the centroid in many cases. Four other methods were to use a Gaussian fit to the counts in a box of 3 x 3, 5 x 5, 7 x 7 , or 9 x 9 pixels centered on the brightest pixel. The final method was to use the center 3 x 3 box, fitting a parabolic function separately to the averaged counts in the sample and line direction, determining the sample and line coordinates of the resultant center.

We found out that the 5 x 5 Gaussian performed best. As shown in Figure 3, we typically obtained post-fit residual to the position of the comet of about 0.1-0.2 pixels in both the sample and line directions. The larger Gaussians gave worse results, being more affected by noise and biased by the coma, as was the case with the moment algorithms. The 99 % moment algorithm was inherently no more accurate than half a pixel, which made that method even less accurate. The 3 x 3 Gaussian and 3 x 3 marginal distributions provided results that were almost as good as the 5 x 5 Gaussian and would have been good enough to use if necessary.

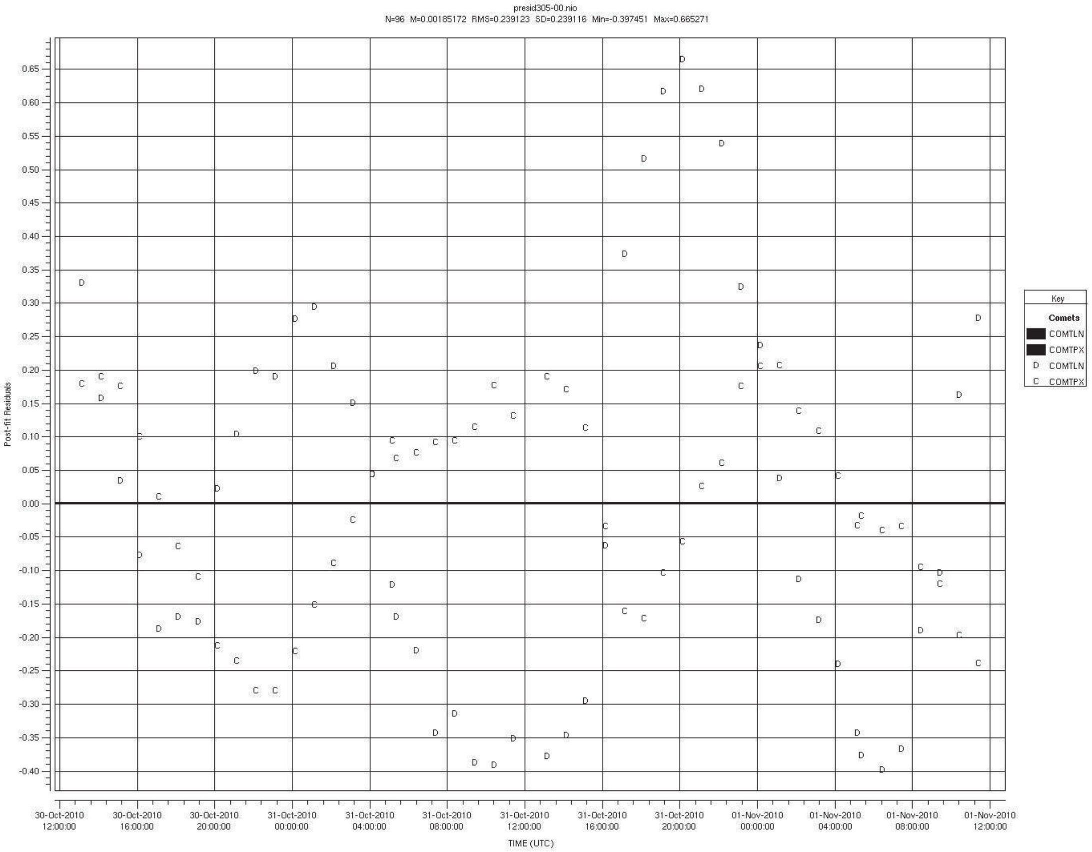

Figure 3. Residual plot for 48 opnav observations, from ~OCT-30 12:00 to ~NOV-30 12:00

Star Centroiding

Centroiding the star images was much more simple. In all cases, we used a Gaussian fit to an 11 x 11 pixel box centered on the predicted star location, after subtracting a constant background. We used the star at the center of each snip box, as well as any other stars that fell within the snips (including the comet snip) that were bright enough to detect without being saturated. This provided us with dozens of stars in most pictures. Star residuals, after the camera pointing had been updated, were typically less than 1/20th of a pixel rms.

ENCOUNTER SCIENCE PICTURES AND EPHEMERIS RECONSTRUCTION

The primary purpose of optical navigation for the EPOXI mission was to calculate the opnav inputs to the orbit determination process, to determine the spacecraft location up to E-26 hours. Autonav then took over and planned the final hours up to encounter.

We did however, perform opnav analysis with pictures past E-26 hours. Although there were no further dedicated opnav pictures, we obtained 142 science pictures, i.e., pictures taken of the comet for scientific analysis, and processed these as opnavs. Forty-nine of these pictures were taken from about half a day before to about half a day after closest approach, and 93 were taken during the departure period, from November 9 to 26.

The latter were analyzed in the same way as the other opnavs, using automatic registration for the stars, and then using a 5 x 5 Gaussian fit to centroid the comet images. The pictures closer to encounter, however, were processed differently, because the comet image was significantly resolved into an obvious bilobate nucleus, and thus no simple method of automated centroiding would yield accurate results. For these pictures, used manually registration based on an overlay of the star field, and then determined the center of the comet image by eye. This gave larger pointing uncertainties, of up to 6 pixels, especially closest to encounter when the star images were smeared because of the fast rate at which the spacecraft was turning to stay locked onto the comet; and larger uncertainties in the comet centroid, of 1-20 pixels, depending on the size of the comet image in the picture. However, these pictures were taken at such small distances from the comet that these large uncertainties in pixels mapped to very small uncertainties in kilometer space. This allowed these opnavs to significantly increase the accuracy of both the final reconstructed spacecraft trajectory, and of the Hartley 2 ephemeris.

SUMMARY

Over a period of 59 days, the MRI camera on board the Deep Impact spacecraft was commanded to take 734 opnav images. Of the 712 that were delivered to the ground, all were processed successfully. Of these, 692 were met the criteria to be kept in the final delivered dataset using the 5x5 Gaussian centroiding method. Typical post-fit residuals in the sample and line directions were on the order of 0.1-0.2 pixels. These data were combined with doppler and range radio observations to provide the final orbit determination solution that was used to seed the onboard Autonav. This solution had formal 1-σ uncertainties of 3 km in the B-plane and 1 second in time-of-flight. This was sufficient for Autonav to perform the final guiding, allowing for a successful encounter at 694 km from the Hartley 2 nucleus.

ACKNOWLEDGEMENTS

We with to thank Nick Mastrodemos for giving frequent advice and insight based on his detailed experience with Deep Impact optical navigation, from planning for and executing the Tempel 1 encounter.

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

REFERENCES

[1] D. L. Hampton *et al.*, “An Overview of the Instrument Suite for the Deep Impact Mission,” *Space Sci. Rev.* **117**, 43-93, 2005.

[2] Owen, W. M., Jr., Mastrodemos, N., Rush, B. P., Wang, T.-C. M., Gillam, S. D., and Bhaskaran, S., “Optical Navigation for Deep Impact,” *16th Annual AAS/AIAA Space Flight Mechanics Meeting*, American Astronautical Society, San Diego, CA, Jan. 2006, pp. 1231-1250.

[3] W. M. Owen, Jr., “Methods of Optical Navigation,” paper 11-215, *21st Annual AAS/AIAA Space Flight Mechanics Meeting*, New Orleans, 2011.

[4] Frauenholz, R. B., Bhat, R. S., Chesley, S. R., Mastrodemos, N., Owen, W. M., Jr., and Ryne, M. S., “Deep Impact Navigation System Performance,” *Journal of Spacecraft and Rockets*, Vol. 45, No 1., Jan.-Feb. 2008, pp. 39-56.

[5] Ryne, M., Jefferson, D.,Craig, D., Higa, E., Lewis, G., and Menon, P., “Ground-Based Orbit Determination for Deep Impact,” *16th Annual AAS/AIAA Space Flight Mechanics Meeting*, American Astronautical Society, San Diego, CA, Jan. 2006, pp. 1179-1202.