(Preprint) AAS 11-486

NAVIGATION OF THE EPOXI SPACECRAFT TO COMET HARTLEY 2

Shyam Bhaskaran, Matt Abrahamson, Steven Chesley, Min-Kun Chung, Allen Halsell, Robert Haw, Cliff Helfrich, David Jefferson, Brian Kennedy, Tim McElrath , William Owen, Brian Rush, Jonathon Smith, Tseng-Chan Wang, Chen-Wan Yen,*

On November 4, 2010, the EPOXI spacecraft flew by the comet Hartley 2, marking the fourth time that a NASA spacecraft successfully captured high resolution images of a cometary nucleus. EPOXI is the extended mission of the Deep Impact mission, which delivered an impactor on comet Tempel-1 on July 4, 2005. EPOXI officially started in September 2007 and eventually took over 3 years of flight time and had 3 Earth gravity assists to achieve the proper encounter conditions. In the process, the mission was redesigned to accommodate a new comet as the target and changes in the trajectory to achieve better imaging conditions at encounter. Challenges in navigation of the spacecraft included precision targeting of several Earth flybys and the comet encounter, uncertainties in determining the ephemeris of the comet relative to the spacecraft, and the high accuracy trajectory knowledge needed to image the comet during the encounter. This paper presents an overview of the navigation process used for the mission.

INTRODUCTION

On July 4, 2005, the Deep Impact (DI) spacecraft successfully impacted the comet Tempel 1 using a 400 kg impactor while the mother ship flew by the comet at a distance of 500 km and captured images of the impact. Its primary mission completed, the mother ship, with a suite of instruments including a high and medium resolution camera and an infrared spectrometer, and over 70 m/s of propellant remaining, was placed on a trajectory which would eventually result in a flyby of the comet Boethin in late 2008. In mid-2007, NASA announced that a new mission for DI was funded under the Discovery Mission of Opportunity Program. The mission combines two separate proposals into one mission. The first, called Extrasolar Planet Observation and Characterization (EPOCh) was to use the high resolution camera to image distant stars and observe the transit of planets orbiting those stars. The second, called the Deep Impact eXtended Investigation of comets (DIXI) was to flyby Boethin and take high resolution images and infrared spectra of the nucleus. The two missions were fused together and renamed EPOXI (a merging of EPOCh and DIXI).

The mission officially began on September 26, 2007 when commands from the Earth brought the spacecraft out of hibernation. Unfortunately, during this time attempts to image Boethin using ground-based telescopes proved unsuccessful, leading to the conclusion that Boethin may have broken up. The decision was made in November 2007 to redirect the spacecraft to a new comet,

*Navigation and Mission Design Section, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, MS 301-820, 4800 Oak Grove Dr., Pasadena, CA 91109.

Hartley 2. The rendezvous with Hartley 2 would be a longer mission involving three Earth gravity assist flybys and an encounter in the fall of 2010. On November 4, 2010, EPOXI successfully flew by Hartley 2 at a distance of approximately 700 km and relative velocity of 12.3 km/s, fulfilling the mission's science goals.

This paper provides an overview of the navigation techniques used by the EPOXI mission to fly by Hartley 2. We describe the process to design the trajectory that satisfied the mission constraints for the flyby, and how the reference design trajectory was navigated using ground-based and onboard navigation methods. We pay particular attention to the specific challenges faced when navigating to comets, which often have poorly determined orbits due to the large non-gravitational forces acting on them. Finally, we close out with a brief description of the closed-loop onboard autonomous navigation (AutoNav) that was used to track the nucleus during closest approach.

MISSION DESIGN

The first step in navigating a deep space mission is to design the reference trajectory. The reference trajectory must satisfy the geometric requirements of the science instruments in observing the comet during the flyby while still meeting mission constraints. The process of finding an alternate flyby target began late in 2006; prior to this in August 2005, a 97 m/s maneuver was performed to nominally direct the spacecraft towards an Earth flyby and a possible encounter with Boethin. The main constraint of the extended mission was the remaining propellant available for diverting the spacecraft to the next comet, which was about 71 m/s. Starting from the DI orbital state on January 1, 2006, test runs to target to various comets were made with an efficient conic trajectory optimizer called MIDAS.[1] Tests were made to many short-period comets possessing a perihelion distance of less than about 1.6 AU and perihelion dates later than January 2008 but before 2011. Since the remaining amount of propellant was small, the possibility depended intimately on the Earth gravity assist(s) already in the plan made to go to Boethin.

Among the potential targets found, comet Hartley 2 was selected as the destination by the Principal Investigator (PI) Mike A'Hearn on Oct 19, 2007. Factors such as ΔV requirement, flight times, viewing conditions and nature of the comet and reliability of ephemeris knowledge led to the decision. The characteristics of the designed trajectory was such that the first Earth Gravity Assist (EGA-1) in December 2007 would alter the spacecraft's orbital period from 1.5 to 1 year, and result in an orbit with small eccentricity (about 0.11), and low inclination (about 0.2 deg). This made possible the subsequent Earth Gravity Assists (EGA-2 and EGA-3) on December 2008 and December 2009. The last flyby was to be used to bend the spacecraft orbit to encounter Hartley 2 on Oct 11 2010, with the encounter taking place very near its perihelion at a solar distance of 1.037 AU.

However, it was noted that the resulting Sun-comet-spacecraft phase angle of 70 deg on approach to Hartley 2 was not ideal for a science measurement. In particular, an approach phase angle near 90 deg was preferred for the Infrared (IR) spectrometer to remain cooler on approach. So, in the spring of 2008, the trajectory was further modified to attain a more favorable approach phase angle. This modification was made by shifting the EGA-3 date to June 2010, half a year later than the original. Since the orbit is nearly 1 AU circular, a 360 deg transfer between EGA-2 to EGA-3, or adding another half a revolution to make EGA 3 occur in June 2010, can both be easily accommodated. This shift resulted in a phase angle of 86 deg at encounter. On June 19, 2008, a relatively large maneuver was performed to insert the spacecraft into this alternate trajectory, which became the new baseline, and EPOXI encountered Hartley 2 on November 4, 2010 at a relative velocity of 12.3 km/s (Figure 1). The trajectory contained two very distant Earth flybys (DFB-1 and 2) in addition to the two remaining closer Earth flybys (EGA-2 and 3 as noted). The new reference trajectory also required one additional, small maneuver a month before the last Earth flyby.

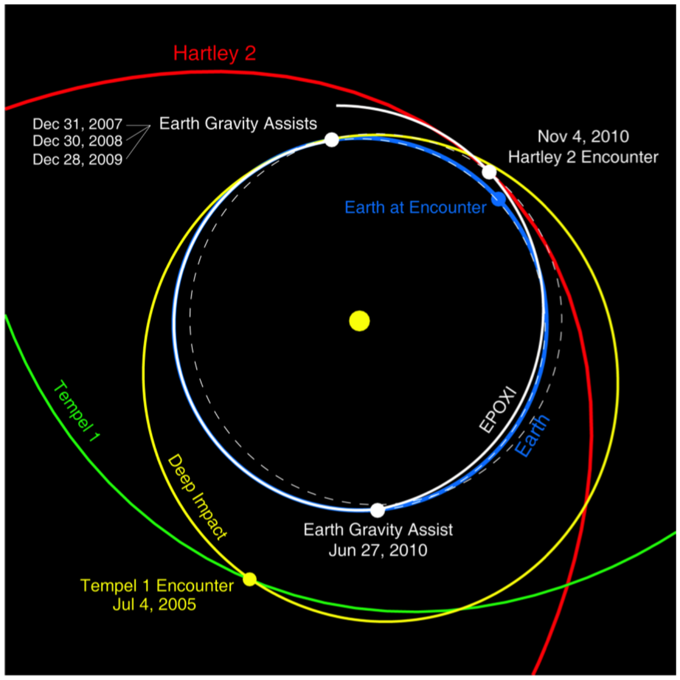

Figure 1. EPOXI Reference Trajectory to Comet Hartley 2

NAVIGATION

Once the reference trajectory is created, the purpose of the navigation is to fly as close as possible to that reference. Navigation involves two steps: the first, Orbit Determination (OD), is the process of determining the current and past state of the spacecraft, and then predicting its future course. The second step is to perform periodic trajectory correction maneuvers (TCMs) to adjust the trajectory as necessary. These processes will be described in the next two sections.

Orbit Determination

Deep space OD involves fitting a high fidelity dynamic model of the spacecraft's trajectory to various types of tracking data. The dynamic model of the orbit is a set of second order differential equations which describe the various forces acting on the spacecraft in a cartesian coordinate system. The forces include the gravitational attraction of the Sun, all the planets and the Moon, solar radiation pressure, the effects of non-spherical gravity field harmonics (when within the sphere of influence of a planetary body), spacecraft thrusting events, and relativistic effects. For high speed cometary flybys, the gravitational attraction of the comet on the spacecraft is negligible and not typically included. The differential equations are numerically integrated to provide a complete description of the position and velocity of the spacecraft at any time. The origin of the coordinate system depends on the mission phase. Since EPOXI spent most of its time in deep space far from any planets, the solar system barycenter was used for the origin. As the spacecraft approached and flew by the Earth, the center of the Earth was used as the origin. In addition, during the Earth flybys, a 70×70 spherical harmonic expansion was used to model the Earth's gravity field.

Thrusting events deserve special mention because they are typically the most problematic force to model. For planned TCMs, the thrust is modeled as a finite burn whose force is described by a second-order polynomial. The mass flow rate is also modeled as a second order polynomial. These models, however, are only approximations as the actual force variation is more complex, leading to a source of error in the OD process. The other type of thrusting event is an angular momentum desaturation. The spacecraft is equipped with three orthogonal reaction wheels which are used to reorient the spacecraft and maintain its attitude. From time to time, the wheels must be spun down to avoid exceeding their design limits; this is accomplished by firing thrusters to change the momentum state of the spacecraft. Typical sizes of these events are on the order of a few mm/s, so they are modeled in the dynamics as impulses which change the instantaneous velocity of the spacecraft but not its position. Thus, at each of these events, the numerical integrator is stopped, the velocity change is added, and the integrator restarted. For propagating the trajectory into the future, the spacecraft Attitude Control Team provides predicted values of the desaturation burns based on predicted spacecraft momentum states. The events happen on average once or twice a week during normal cruise activities.

Tracking Data

Given a set of initial values for the spacecraft state and other parameters describing the dynamics of the spacecraft, the equations of motion are integrated to get a nominal trajectory. To get the actual trajectory, tracking data are acquired and compared against predicted values based on the modeled trajectory. The difference between the two (the residual) is rectified by adjusting the state and other parameters in a least-squares process to minimize the residuals. Throughout the large portion of the cruise and Earth flybys, standard Earth-based radiometric tracking data are used. These include two-way coherent Doppler, which compares the received frequency at the tracking station against the transmitted frequency and thus provides a measure of the line-of-sight velocity of the spacecraft relative to a tracking station, and two-way range, which computes the time it takes for a transmitted signal to return to the spacecraft an thus provides a measure of the spacecraft's distance relative to a tracking station. These measurements are augmented by Delta Differenced One-way Range (DDOR), which measures the delay between the spacecraft signal acquired simultaneously from two stations, and differences that with the delay in signal from a nearby (in an angular sense) quasar taken at roughly the same time. The differencing process removes much of the common error sources inherent in radio signals propagating through the Earth's atmosphere; the final observable is a measure of the angular position of the spacecraft projected into the plane-of-sky. Thus, DDOR complements Doppler and range in that it provides information about a direction perpendicular to the line-of-sight information from Doppler and range. All three are scalar data types, which must be used to solve for the multi-dimensional spacecraft state.

The above data are sufficient to accurately determine the heliocentric state of the spacecraft. For encounters with solar system bodies however, the important factor at encounter is the knowledge of the spacecraft state relative to the target body. In the case of many of the planets, the heliocentric ephemeris of the body is also well known so the ground-based radiometric data also provides enough information for the encounter. For the large outer planets where the ephemeris may not be accurate, the acceleration to the spacecraft caused by the gravity of the planet provides a strong signature in the Doppler data which can be used to infer the relative state. For small bodies however, and comets in particular, the ephemerides of the body is not known well, and there is little or no gravity signature, so a target relative data type is needed. This is provided by optical data, which uses an onboard camera to take pictures of the target body against a star background. The measured position of the star images in the camera frame provides the inertial pointing direction of the camera boresight, and the location of the image of the object within the picture provides an angular measurement of the target position relative to the spacecraft. Details of the use of optical data for interplanetary missions are given in Reference 2. Optical data collection typically begins about 1-2 months before the encounter when the target object becomes bright enough to see in the camera. For EPOXI, the optical data were acquired by the onboard Medium Resolution Imager (MRI) camera, which has a field-of-view (FOV) of 0.6 deg, spread over a 1024 square pixel array. More information about how optical navigation was used on EPOXI is provided in Reference 3.

Filter

Given the above scalar data sets, the objective is to fit the combination of them to solve for the multidimensional state of the spacecraft. The fit is done through a batch-sequential least-squares process which adjusts the parameters of the dynamic model and models of the observation data to minimize the sum-of-squares of the residuals.[4] Since the fit is a statistical process, an error estimate falls out of the least-squares fit which can be used to determine the accuracy of the solution at any given time, and mapped and rotated into any coordinate frame. For deep space applications, it is convenient to map the error covariance into B-plane coordinates. The B-plane is centered on the target body and perpendicular to the incoming asymptote of the trajectory; the location of the asymptote as it pierces the B-plane are values of interest, along with the time when it cross the plane, or distance along the asymptote. These components are referred to as B•T (the projection in the plane in the horizontal, or T axis), B•R (the projection in the plane in the vertical, or R axis), and the time to cross the plane (the linearized time-of-flight, LTOF). This system is especially useful for small body flybys since the target body is massless for all practical purposes and thus the B-plane coordinates represent the actual flyby conditions.

In addition to parameters that are estimated, there are also some parameters that are "considered"; these are parameters which aren't adjusted, but their error contributes to the filtered error covariance. They are used to provide more realistic error estimates which can be too optimistic otherwise. In the standard OD process, estimated parameters include the spacecraft state (position and velocity), solar radiation pressure coefficients, maneuver ΔVs, random accelerations, and biases in the range data. Consider parameters include Earth orientation parameters, media calibration errors, and the Earth-moon barycenter location.

Maneuvers

Given the knowledge of the current state of the spacecraft and its predicted path, TCMs can be performed to correct the trajectory as needed. There are two types of TCMs: "Deterministic" maneuvers are those planned in the reference trajectory and need to be performed in order to achieve the desired flyby. "Statistical" maneuvers have zero magnitude in the reference but are placed at strategic locations to clean up errors in the orbit determination and execution errors from previous maneuvers. Table 1 lists all the deterministic and statistical maneuvers used by the EPOXI mission following the decision to encounter Hartley 2 instead of Boethin. The deterministic maneuvers are designed in the reference (as described above), usually with some type of trajectory optimizer to place the TCMs in an optimal location to minimize the total amount of ΔV needed. For the statistical TCMs however, the maneuvers can be computed more simply by linear methods since the deviation of the actual trajectory from the reference is small. Typically, three orthogonal components of the target location (BR, BT, and LTOF) are controlled by the three Cartesian components of the maneuver, but critical plane targeting, where only the components in the B-plane but not LTOF are controlled, are also used.

Table 1. Maneuver Schedule

Maneuver	Date	Designed ΔV (m/s)	Achieved ΔV m/s	Comment
TCM 9	November 1, 2007	14.492	14.633	
TCM 10	December 11, 2007	-		Cancelled
TCM 11	January 16, 2008	-	-	Cancelled
TCM 12	June 19, 2008	31.543	31.563	
TCM 13	December 11, 2008	0.594	0.594	
TCM 14	February 19, 2009	0.828	0.829	
TCM 15	March 18, 2009	-	-	Cancelled
TCM 16	December 8, 2009	0.494	0.496	
TCM 17	January 21, 2010	-	-	Cancelled
TCM 18	May 28, 2010	0.109	0.109	
TCM 19	July 19, 2010	0.845	0.846	
TCM 20	September 29, 2010	1.533	1.534	
TCM 21	October 27, 2010	1.587	1.590	
TCM 22	November 2, 2010	1.360	1.361	

Ephemeris

The target body trajectory is a key constraint at every stage of a flyby mission, from the earliest design phase through the encounter. Few comets and asteroids have observations from spacecraft in deep space, and those are usually optical navigation measurements associated with a mission encounter, and so the pre-encounter ephemerides are dominated by ground-based optical measurements (angles-only, Right Ascension and Declination) and occasionally some near-Earth observations from spacecraft such as WISE or Hubble Space Telescope (HST). Late ephemeris updates are often able to take advantage of optical navigation measurements obtained by the flyby spacecraft, leading to dramatically reduced uncertainties in the spacecraft plane of sky. Also, in an unusual twist, the EPOXI mission was able to take advantage of radar astrometry (delay and Doppler) from Arecibo Observatory.[5]

A key distinction between comets and asteroids is the outgassing of volatile compounds from the surface of comets. This activity is mostly dominated by the sublimation of water, which imparts a difficult-to-model nongravitational acceleration on the body. Also, the gases carry entrained dust with them creating a large cloud, or coma, that obscures the exact location of the cometary nucleus. The combination of a poor dynamical model and a poor measurement model for comets makes them far less tractable for ephemeris prediction than asteroids. And because the potential errors associated with these factors are also poorly known even the estimated ephemeris uncertainties are poorly constrained.

Early ephemeris products, e.g., from the design phase of the mission, can have prediction intervals of up to several years and the associated uncertainties are typically too great to actually implement the flyby. However, simulated future observations can be folded into the covariance analysis to give mission planners a coarse estimate of what the ephemeris uncertainties should be in the weeks leading up to encounter. Conversely, ephemerides derived shortly before encounter have relatively long data sets and short prediction intervals and so tend to be much more stable and reliable than the early estimates. However, in the face of significant mismodeling of the nongravitational accelerations it is important to rely most heavily on the recent observations. Therefore the data arc length may need to be shortened repeatedly, and sometimes dramatically, in order to obtain decent fits to the most recent and most accurate measurements. The details of how all ephemeris solutions were obtained for the EPOXI mission are out of the scope of this paper; here we will summarize the operational results and its consequence in navigating the spacecraft.

NAVIGATION OPERATIONS BY MISSION PHASE

EGA-1

Following the Deep Impact encounter in 2005, the surviving mother ship was placed into "hibernation" mode, where it shut down most functions and was placed into a spinning sun-coning attitude to maintain the solar panels on the sun. On September 25, 2007, the EPOXI mission began when the spacecraft was taken out of hibernation and put into 3-axis mode, with the solar panels on the sun, the high gain antenna pointed to the Earth, and attitude maintained with the reaction wheels. Navigation data also began, with the Navigation Team receiving Doppler, range, and periodic DDOR data. The first major event was a planned TCM 9 on November 1, 2007, which would target an Earth flyby on December 31, 2007. The official plan at this point was to still go to comet Boethin. However, since it was becoming obvious that Boethin couldn't be found, the Navigation Team had been planning the Hartley 2 trajectory. On October 19, a project meeting was held with the PI and it was decided to use TCM 9 to to go to Hartley 2.

The standard navigation procedure for designing a TCM is to first obtain the current OD knowledge by fitting all the tracking data in a specified "data arc." Data arcs usually span about 2–3 months, with the end occurring 7–10 days prior to the maneuver. The gap between the end of the data arc and the maneuver is used to design the ideal maneuver, adjust it based on spacecraft constraints, test the maneuver on a high fidelity simulation testbed, create command products to uplink, and finally uplink the products during a tracking pass. For TCM 9, the data arc went from the start of the mission in late September to October 23, 2007, and included over 7000 Doppler points, over 3200 range points, and 6 DDOR points. The least-squares filter used the parameters listed above, and the epoch state solution was mapped and rotated into the Earth flyby B-plane coordinates and compared to the desired flyby B-plane provided by the reference trajectory. Also as is typical, variations on the solution were tried by using different data arc lengths and combinations of the three data types to evaluate consistency among the solutions. An example of these solutions is shown in Figure 2(a) for a variety of cases, with the baseline solution shown in blue. The ellipses shown are the 1σ error covariances of the solutions also mapped to the B-plane. Since the variations showed relatively good consistency with respect to their uncertainties, the baseline solution was picked as the starting point for the design of TCM 9.

TCM 9 was designed to correct back to the original Hartley 2 reference trajectory by targeting the desired location in the Earth flyby B-plane coordinates. The computed value for TCM 9 was 14.49 m/s, and its effect in the B-plane is shown in Figure 2(b). Reconstruction of the burn after the maneuver showed that the error in the execution amounted to a roughly 1% overburn. A cleanup maneuver, TCM 10, was planned for December 11 to correct any OD and maneuver execution errors from TCM 9. Plots of the OD at the time of the data cutoff for the TCM 10 design indicated, however, that the estimated trajectory was within 1σ of the targeted location, so TCM 10 was canceled.

EGA-1 occurred as planned on December 31, 2007. The desired time of perigee and radius at perigee were 19:29:05 UTC and 21,944.3 km. The post-flyby reconstruction of the trajectory indicated the actual location was at 19:29:20 UTC and 21,943.6 km, an error of 15 seconds and 700 m. Because this was within the envelope of expected errors for the flyby, another planned statistical maneuver, TCM 11 on January 16, 2008, was canceled.

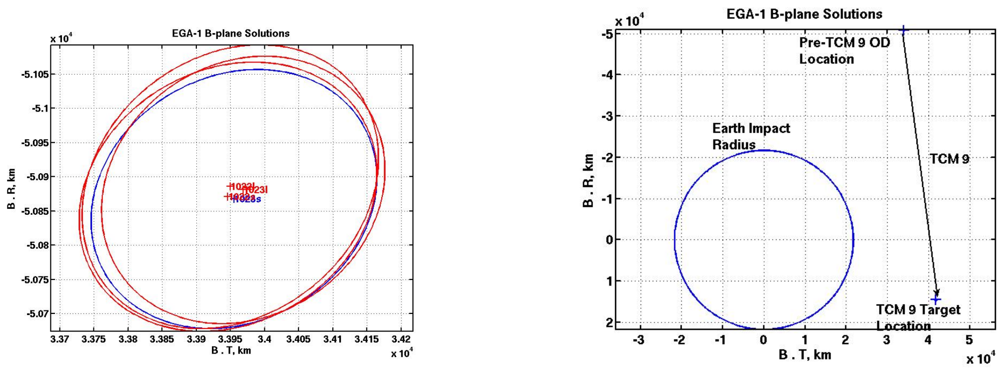

(a) Pre-TCM 9 OD Solutions

(b) TCM 9 Maneuver and Target

Figure 2. TCM 9 OD in Earth B-plane

TCM 12 and EGA-2

As mentioned above, a search for an alternate trajectory with more favorable encounter geometries was initiated in the spring of 2008. The resulting trajectory, although better for science, did use more fuel, with a deterministic maneuver of over 31 m/s needed for TCM 12 in June, as opposed to 12 m/s. The OD arc to design TCM 12 spanned the period April 18, 2008 to June 8, 2008, and included Doppler, range, and DDOR points. The final design of TCM 12 based on this OD solution had a magnitude of 31.543 m/s. Post-maneuver reconstruction indicated the actual maneuver executed was 31.563 m/s, for a 0.062% overburn.

Following TCM 12, navigation operations became fairly quiet during the summer of 2008 as the EPOCh astronomical observations took place. Operations resumed in the fall of 2008 as the team prepared for TCM 13, twenty days prior to EGA-2 on December 29, 2008. Because the spacecraft had gone for considerable time without navigation updates, even the small error in TCM 12 had to be cleaned up prior to the Earth flyby. The data arc for TCM 13 started on September 2, 2008, and had a cutoff on December 2, 2008, and once again, used all three radiometric data types. In order to get back on the reference trajectory, TCM 13 was designed with a nominal magnitude of 59.4 cm/s and targeted to the reference trajectory's flyby point in the EGA-2 Earth B-plane. Figure 3(a) shows the baseline OD solution used for the design and Figure 3(b)the effect of the maneuver in the B-plane. The reconstruction of the maneuver indicated the achieved ΔV to only a 0.005 % underburn, very close to the design. The resulting flyby condition had a perigee on December 29, 2008 21:39:57 UTC at a radius of 49,835 km, less than 1 sec and 7 km from the desired location.

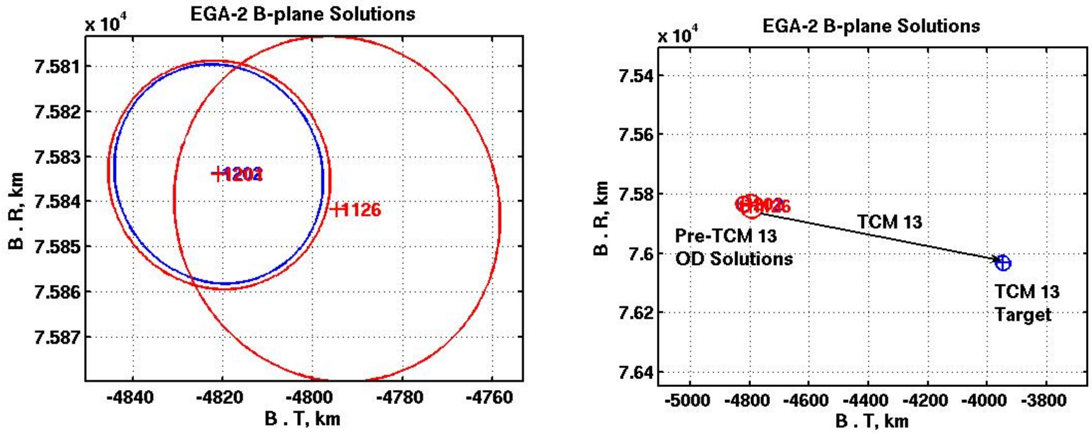

(a) Pre-TCM 13 OD Solutions

(b) TCM 13 Maneuver and Target

Figure 3. TCM 13 OD in Earth B-plane

DFB-1 and DFB-2

The standard navigation procedure following the Earth flyby would be to insert a statistical maneuver a few weeks after the flyby to correct any errors from the gravity assist, and thus TCM 14 was scheduled in mid-January 2009. We also had a statistical TCM 15 scheduled in mid-March because of the sensitivity of the low energy trajectory to maneuvers in this location. In practice, we decided to combine the two maneuvers in order to minimize the load on the operations teams. Analysis had indicated that by delaying TCM 14 by a month, we could remove TCM 15 with little or no penalty. Unlike most statistical maneuver designs, TCM 14 was re-optimized to find the minimum ΔV solution to achieve the comet flyby. One complicating factor, however, was that the direction of the ΔV vector in the optimal solution violated spacecraft attitude thermal constraints. Two options were considered. The first was to split TCM 14 into two maneuvers a day apart, each satisfying the thermal constraint and with the vector sum of the two achieving the same result. The second was to compute a value for TCM 14 that forced it to satisfy the thermal constraint and then insert a deterministic maneuver at the TCM 16 time (in December 2009). The optimization for the latter case resulted in finding a solution which had a negligible deterministic ΔV for TCM 16, and the penalty for adding the constraint was only 6 cm/s, so this option was eventually used.

The design of TCM 14 used a data arc from December 12, 2008 through February 7, 2009, which spanned the Earth flyby. The design magnitude of TCM 14 was 82.8 cm/s, and its reconstructed value was 82.9 cm/s, a 0.15% overburn.

The first of the two distant Earth flybys, DFB-1, occurred on June 29, 2009. Because of the large range at perigee (over 1.3 million km), it was mostly a non-event from a navigation standpoint. By the time of TCM 16 on December 8, 2009, however, enough error had accumulated that a correction was necessary. The strategy used to compute TCM 16 was similar to TCM 14; the trajectory optimum was found by adjusting both TCM 16 and TCM 18 (planned for May 28, 2010) simultaneously, with the target for TCM 16 being the cartesian position of the spacecraft obtained from the reference trajectory at the time of TCM 18, and the target for TCM 18 being the comet encounter condition. Starting from the OD solution using a data arc from October 15 through November 29, 2009, the solutions for TCM 16 and 18 were found to be 49.36 and 0.683 cm/s, respectively. TCM 16 executed properly, with the result being 49.58 cm/s, or a 0.44% overburn. DFB-2 occurred uneventfully on December 28, 2009.

EGA-3

The final Earth flyby was critical because the gravity assist changed the parameters of the spacecraft orbit that would allow it to encounter Hartley 2 in November, and any errors in the flyby conditions would be expensive in terms of ΔV to clean up later. It was also important at this point to get the best available ephemeris for Hartley 2 as any changes to the ephemeris could be corrected by using the Earth's gravity to adjust EPOXI's path rather than the remaining fuel. Thus, a new Hartley 2 ephemeris solution using all ground-based data until April 2010 was delivered prior to the design of TCM 18. The solution was labeled sb-103p-51, and the difference between this and the one used up to this point was about 2200 km in position at the nominal encounter time.

The OD arc for designing TCM 18 was from March 29 through May 17, 2010. The target point used to compute the maneuver was the 700 km flyby location in the Hartley 2 B-plane at closest approach. Visually, however, it was easier to see the discrepancy between the desired and actual trajectories in the upcoming Earth flyby B-plane. A plot of the OD solutions leading up to TCM 18 is shown in Figure 4. Because of the buildup of execution and OD errors from TCM 16, plus the change in the comet ephemeris, TCM 18 increased from the original design to the final value of 10.86 cm/s; its effect in the B-plane is shown in Figure 4. The actual achieved value was 10.92 cm/s, an overburn of 0.58%. The resultant flyby had a perigee on June 27, 2010 22:03:48.7 UTC at a radius of 36,875.0 km, a difference of 1.2 sec and 16 km from the designed reference.

One interesting aspect of this flyby was the fact that we closely monitored the data to see if the mysterious "flyby anomaly", seen on several spacecraft (Galileo, NEAR, Rosetta) that performed Earth gravity assists, manifested itself.[6] The nature of the anomaly is such that these spacecraft experienced a small increase in its Earth relative velocity at some unknown time near perigee which is not predicted in our current dynamic models. No physical explanation has been found for this behavior; however, Reference 6 provides an empirical formula which seems to predict the magnitude of the velocity change when applied to previous flybys. Application of the formula to the EPOXI EGA-3 case results in a predicted anomalous ΔV of over 5 mm/s. However, careful fitting of the pre- and post-encounter Doppler data showed no such anomaly down to the noise level of the data (around 0.1 mm/s). The previous EPOXI gravity assists similarly did not indicate any unaccounted ΔV, although the predicted values were less than a mm/s. We conclude that the formula's apparent ability to predict the magnitude of the anomaly for the other missions was coincidence, or that another scale factor, perhaps related to the altitude, needs to be incorporated into the formula. The latter is especially pertinent since EPOXI's flybys were all over 15,000 km altitude, whereas the previous flybys where the anomaly was seen were under 2500 km.

Although Hartley 2 was being regularly observed from the Earth through the EGA-3 time period, no ephemeris updates were received prior to the design of TCM 19 (scheduled for July 19, 2010), intended as cleanup for the Earth flyby. And, as small as the Earth flyby error was, it had a magnifying effect such that the trajectory error was noticeable when projected to the time of the encounter. Using an OD arc from May 29 through July 12, 2010, the trajectory mapped to the Hartley 2 B-plane (shown in Figure 5) was off by over 5000 km with respect to the desired 700 km radius flyby location. Thus, TCM 19 was designed with a magnitude of 84.5 cm/s to correct the error. The maneuver executed properly, with the reconstruction showing the achieved magnitude to be 84.6 cm/s.

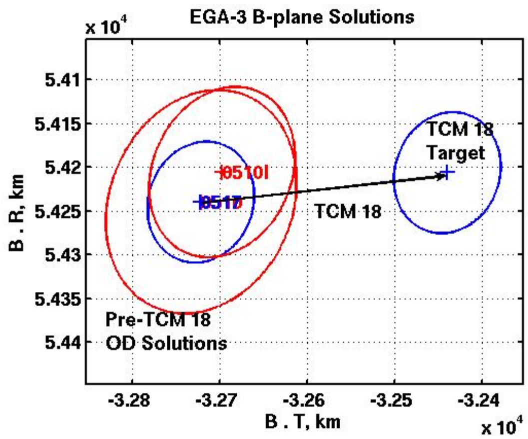

Figure 4. TCM 18 OD Solutions and Maneuver in Earth B-plane

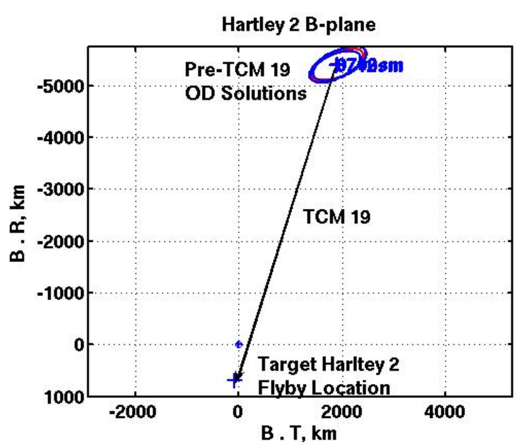

Figure 5. TCM 19 OD Solutions and Maneuver in Hartley 2 B-plane

Approach

The approach phase for small body missions begins with the first optical navigation (Opnav) images taken of the comet several months prior to encounter. At this point, the spacecraft's heliocentric trajectory is known to a few tens of km or better inside the data arc, and the uncertainty mapped to encounter is on the order of a few hundred km. The uncertainty of where the comet is based on ground astrometric observations, however, is much larger – on the order of a thousand km or more. The addition of the spacecraft-based comet observations provides the third leg of the triad to determine the spacecraft's comet-relative position. Initially, this data is used primarily to improve the comet's ephemeris and thus, using a radiometric-only spacecraft trajectory as the observing platform, the Opnav observations are combined with the ground-based astrometry to improve the comet ephemeris. Around the last week or so prior to encounter, the uncertainty in comet location is commensurate in size with that of the spacecraft and so a simultaneous estimate of the spacecraft and comet trajectory is performed. This provides the best comet-relative information until AutoNav takes over an hour before encounter to get a purely Opnav-based comet-relative solution for tracking the comet through closest approach.

Specific details of the Opnav process on EPOXI (picture frequency and scheduling, processing the images, details of the camera system, etc.) are provided in Reference 3; here we give an overview of the results. The first Opnav was taken on September 5, 2010. With only a few days of additional observations, it was apparent that the ground-based ephemeris for Hartley 2 being used at that point (sb-103p-56) had large errors. Figure 6 is a plot of one of the image frames taken during this time which shows the unresolved comet as a bright blob. Overlaid on this picture is the brightness centroid of the blob (the green plus sign, indicating the comet location) and the predicted location of the comet (shown as a blue circle). The discrepancy between the comet's predicted and observed location is roughly 1.3 pixels, which at this distance translated to about 616 km as projected into the camera FOV. Combined opnav and ground-based astrometry solutions computed during this time indicated that the projected flyby was clustered in an area around B•R of -3100 km, and B•T of 2500 km, almost 4000 km away from the target location (see Figure 7). This was disturbing because the error was well over 3σ of the formal covariance of the last ground-based ephemeris solution, sb-103p-56 (shown in green in Figure 7(b)), indicating serious deficiencies in the dynamic modeling of Hartley 2. This led to considerable work over the next few months to understand the source of the problem, which turned out mostly to be due to the volatile nature of Hartley 2's outgassing and the difficulty in modeling this behavior. The first opportunity to redirect the spacecraft towards the updated comet location was TCM 20 on September 29, 2010. The data arc for the OD leading up to the maneuver went from July 20 through September 23, 2010. All the optical data taken from the spacecraft up to that point were included in the combined ground- and space-based comet ephemeris solution (sb-103p-122) used for the TCM design. The projected flyby location and uncertainty of the spacecraft in the Hartley 2 encounter B-plane relative to this ephemeris is shown in blue in Figure 7. Also shown are two other ellipses (in red) representing two examples of alternate Hartley 2 models. Many variations were tried, using different comet dynamic models, combinations of observing locations, data arc lengths, etc., and in the end, the choice of which solution to use was based as much on intuition than science, weighing various factors such as the goodness and consistency of the data types, knowledge of which observatories provide better data reductions, and experience from previous comet flyby missions.

The nominal flyby target location was 700 km from the center of the comet. The specific location in the B-plane was at the coordinates B•R of 695.8 km, and B•T of -77.4 km. This location was

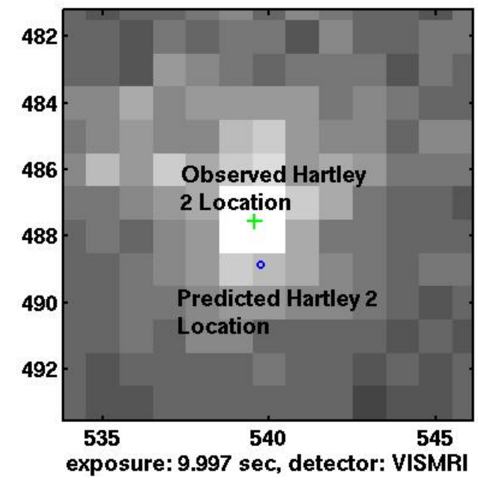

Figure 6. Opnav Frame Showing Unresolved Hartley 2

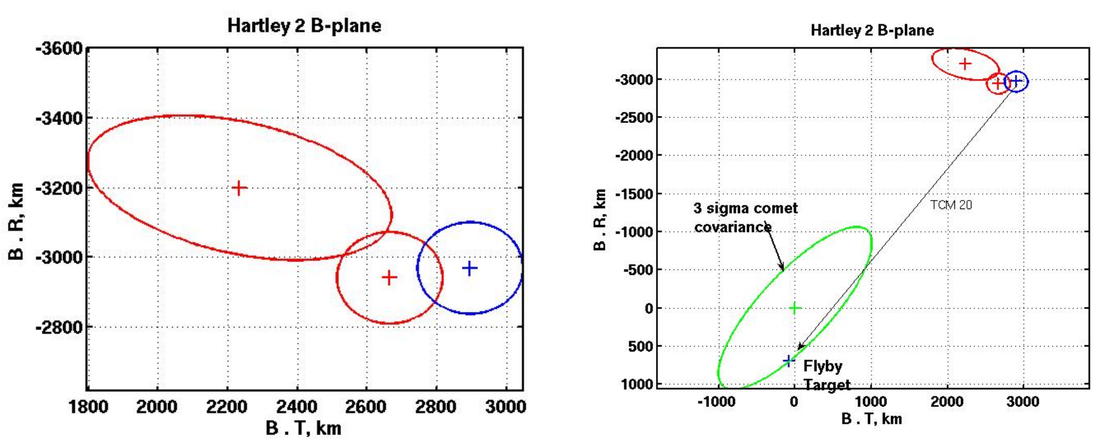

(a) Pre-TCM 20 OD Solutions

(b) TCM 20 Maneuver and Target

Figure 7. TCM 20 OD in Hartley 2 B-plane

chosen by the Science Team to optimize the observations from the various instruments through the flyby and results in a solar phase angle of 77 deg at closest approach. In practice, it would be impossible to hit this spot exactly and so what was needed was a range of allowable flyby locations that was acceptable and realistically achievable by the Navigation Team. Thus, negotiations between the scientists, navigators, and project management resulted in a wedge region, defined by a range between 650 and 750 km, and a closest approach phase angle between 75 and 85 deg, that constitutes an acceptable region for the flyby. This region is shaded in green in Figure 8. Furthermore, it was decided that prior to the last targeting maneuver (TCM 22), if the predicted flyby location was in the green region, no TCM 22 would be performed. The yellow region in the figure is an area where the science was not optimal, but could be acceptable if necessary, and thus, the decision to perform TCM 22 would rest with the PI and Project Manager. If the flyby location were in the red region, TCM 22 would automatically be necessary and needed no decision.

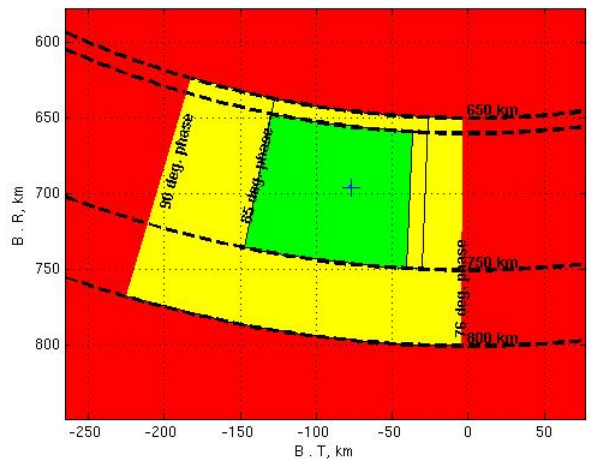

Figure 8. Wedge Region for Acceptable Flyby

The designed magnitude of TCM 20 was 1.533 m/s to correct the large targeting error (Figure 7(b)), and the achieved value was 1.534 m/s (0.07% overburn). One complicating factor in the design of this, and the subsequent two maneuvers prior to encounter, was that the direction of the burn had constraints relative to its angle with the sun. The reason for this was that the project Science Team strongly desired data from the IR Spectrometer, which had thermal constraints that precluded certain spacecraft attitudes. The details of how the maneuvers were designed within the constraints is described in Reference 7. The next maneuver opportunity was TCM 21 on October 27, 2010, about 34 days prior to encounter. During this time, the spacecraft, as well as observatories on the ground, was continually imaging the comet, and ephemeris solutions were being provided regularly to monitor the location of the comet. The volatility of the Hartley 2 ephemeris was very apparent as shortly after TCM 20, the ephemeris was starting to wander. The data cutoff for the design of TCM 21 was on October 24, 2010, and the OD arc used for the spacecraft began on September 5. As they were for TCM 20, a number of Hartley 2 solutions were examined; examples of two of them, sb-103p-122 (in green) and sb-103p-183 (in blue) are shown in Figure 9 where it is clear that the spacecraft is several hundred km away from the target. The discrepancy in the two trajectories seen in this plot is entirely due to the different comet ephemerides; spacecraft solutions at this time were very consistent with each other. The final decision was to use sb-103p-183. TCM

21 was designed to retarget the spacecraft to the nominal flyby and had a magnitude of 1.587 m/s. The reconstructed value was 1.598 m/s, or a 0.15% overburn.

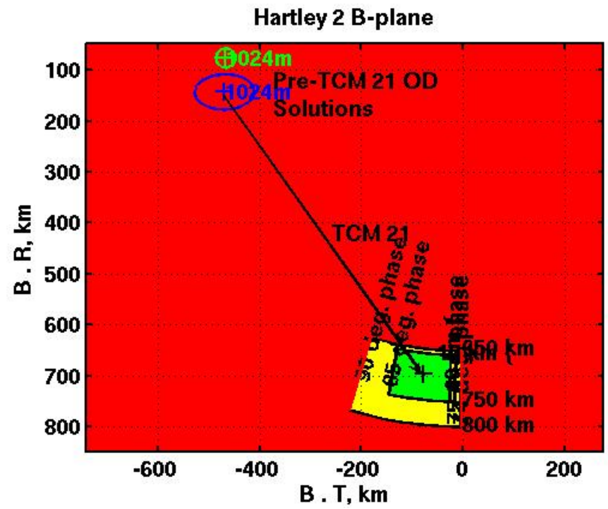

Figure 9. TCM 21 OD Solutions and Maneuver

The final targeting opportunity was TCM 22 on November 2, 2010, about 2 days prior to encounter. With the spacecraft Opnavs providing increasingly higher accuracy observations of the comet and the comet dynamic modeling issues becoming more settled, the error from the previous maneuver had become much smaller. Nevertheless, the current best spacecraft solution relative to the chosen Hartley 2 ephemeris at this time (sb-103p-228) showed it to be right at the edge of the desired region at roughly 750 km range and 76 deg phase angle (Figure 10(a)). The PI and PM agreed at this point to execute TCM 22, which had a design value of 1.360 m/s, and would once again target the nominal aim point. The maneuver executed properly, with an achieved value of 1.361 m/s (.08% overburn).

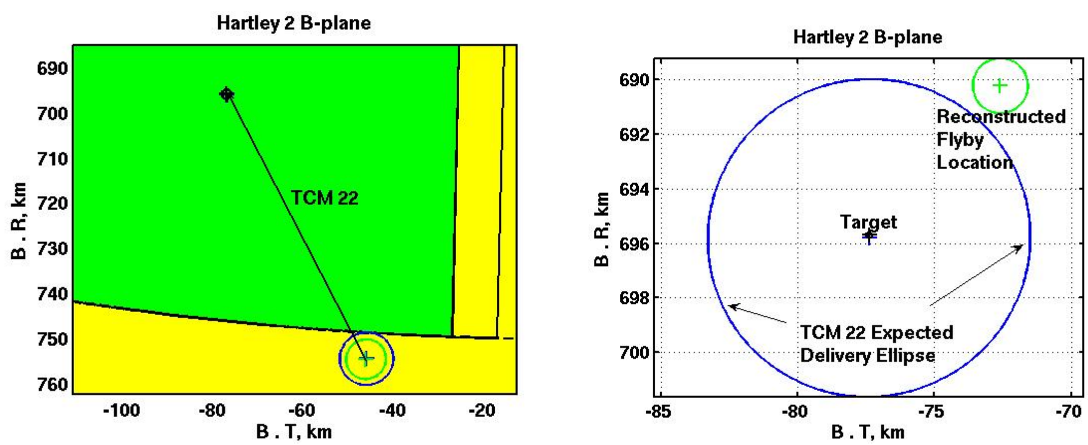

(a) Pre-TCM 22 OD and Effect of TCM 22

(b) Final Reconstruction of Flyby

Figure 10. TCM 22 and Final Flyby Location

Encounter

Following TCM 22, no maneuvers were planned. However, the accuracy of the delivery of TCM 22 (shown by the blue uncertainty ellipse in Figure 10(b)) was not sufficient to keep the instruments pointed at the nucleus through closest approach. In particular, it was a mission requirement to image the nucleus continuously through closest approach using the MRI camera. With a FOV of 0.6 deg, this meant that the knowledge of the spacecraft position relative to the comet at a radius of 700 km had to be better than $\pm$ 3.6 km. This is especially problematic in the downtrack direction since that value translates to a knowledge of less than 0.3 seconds in the arrival time, well below the expected knowledge from ground-based navigation. Thus, AutoNav was needed to maintain visual lock on the nucleus through the flyby.

Details of the DI AutoNav system are provided in Reference 8. Briefly, AutoNav takes images of the nucleus, processes the images to get a brightness centroid, and then performs OD solutions to continuously update the spacecraft's comet relative state, all onboard. Thus, the latest and best information is used at all times without the delay caused by the round-trip light time and time needed for ground processing. The system has been validated on several comet missions including DI.[9]

For the Hartley 2 flyby, AutoNav was initiated at E-50 minutes. The starting point was the best ground-based spacecraft and comet information; this was provided by a ground OD solution using a data arc from September 5 through November 3, 2010. Then, using sb-103p-228 as the reference, Opnav data from October 30 through November 3 was combined with the radiometric data to simultaneously solve for the spacecraft and comet. The resulting solution is shown in Figure 10(b), where the OD solution is shown right on top of the target. At E-50 minutes, AutoNav took images at roughly 15 second intervals and at E-42 minutes, these were combined into the first onboard solution. Subsequently, after every image was taken, AutoNav performed an OD update. AutoNav terminated at 60 minutes past encounter. Images returned shortly after encounter indicated the process was successful, and all 307 images taken had the comet in the MRI FOV. A post-flyby reconstruction of the encounter, which included the frames taken by AutoNav, indicated that the spacecraft was delivered to a B•R and B•T location of 690.222 and -72.600 km, respectively, with a formal 1σ uncertainty of 334 m, circular. This estimate is about 6 km distant from the intended target location (see Figure 10(b)). Figure 11 shows the image of Hartley 2 taken at closest approach.

CONCLUSIONS

Navigation to comets can be a challenging task due to the unpredictable nature of comets and the need to combine multiple data sets and make meaningful conclusions from the process. This paper provided an overview of how this mission was flown and the results.

ACKNOWLEDGEMENTS

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

REFERENCES

[1] C. G. Sauer, "MIDAS: Mission Design and Analysis Software for the Optimization of Ballistic Interplanetary Trajectories," *The Journal of the Astronautical Sciences*, Vol. 37, No. 3, 1989, pp. 251–259.

[2] W. M. Owen, "Methods of Optical Navigation," *AAS/AIAA Space Flight Mechanics Conference*, February 2011. Paper AAS 11-215.

Figure 11. Encounter Image of Hartley 2

[3] B. Rush, W. M. Owen, S. Bhaskaran, and S. P. Synnott, "Optical Navigation for the EPOXI Mission," *AAS/AIAA Space Flight Mechanics Conference*, February 2011. Paper AAS 11-217.

[4] B. D. Tapley, B. E. Schutz, and G. H. Born, *Statistical Orbit Determination*. San Diego, CA: Elsevier Academic Press, 1st ed., 2004.

[5] J. K. Harmon, M. C. Nolan, E. S. Howell, and J. D. Giorgini, "Comet 103/P Hartley," *IAU Circular No. 9179*, November 2010. Ed. Daniel W. E. Green.

[6] J. D. Anderson, J. K. Campbell, J. E. Ekelund, J. Ellis, and J. F. Jordan, "Anomalous Orbital-Energy Changes Observed during Spacecraft Flybys of Earth," *Physical Review Letters*, Vol. 100, No. 9, 2008, pp. 1–4.

[7] M. K. Chung, S. Bhaskaran, S. R. Chesley, C. A. Halsell, C. E. Helfrich, D. C. Jefferson, T. P. McElrath, B. P. Rush, T.-C. Wang, and C. w. Yen, "EPOXI Trajectory and Maneuver Analysis," *AAS/AIAA Astrodynamics Specialist Conference*, August 2011. Paper AAS 11-482.

[8] D. Kubitschek, N. Mastrodemos, R. Werner, B. Kennedy, S. Synnott, G. Null, S. Bhaskaran, J. Riedel, and A. Vaughan, "Deep Impact Autonomous Navigation: The Trials of Targeting the Unknown," *AAS Guidance and Control Conference*, February 2006. Paper AAS 06-081.

[9] S. Bhaskaran, J. Riedel, and S. Synnott, "Autonomous Target Tracking of Small Bodies During Flybys," *AAS/AIAA Astrodynamics Specialist Conference*, February 2004. Paper AAS 04-236.