

Autonomous Navigation for Deep Space Missions

Shyam Bhaskaran¹

Jet Propulsion Laboratory, Pasadena, CA, 91109, USA

Autonomous navigation (AutoNav) for deep space missions is a unique capability that was developed at JPL and used successfully for every camera-equipped comet encounter flown by NASA (Borrelly, Wild 2, Tempel 1, and Hartley 2), as well as an asteroid flyby (Annefrank). AutoNav is the first on-board software to perform autonomous interplanetary navigation (image processing, trajectory determination, maneuver computation), and the first and only system to date to autonomously track comet and asteroid nuclei as well as target and intercept a comet nucleus. In this paper, the functions used by AutoNav and how they were used in previous missions are described. Scenarios for future mission concepts which could benefit greatly from the AutoNav system are also provided.

I. Introduction

For over 40 years, nations have sent spacecraft to visit other natural bodies in the Solar System. The mission profiles have included high speed flybys, rendezvous and orbit, and landings. One of the challenging aspects of flying deep space missions is navigation, that is, knowing where the spacecraft is at any given time (orbit determination), and controlling its path to achieve desired mission objectives (maneuver control). For the vast majority of deep space missions, this process is performed on the ground using standard radiometric tracking data (two-way Doppler, two-way range, Delta-DOR), and, when required, onboard optical data (images taken by the spacecraft camera of nearby target bodies). This combination has worked well and resulted in remarkably accurate navigation over the years.

Despite its success, one of the inherent drawbacks in ground-based navigation is the delay caused by the round-trip light-time – the time it takes for a two-way radio signal to be transmitted from an Earth ground station, received by the spacecraft and retransmitted back to Earth, and finally received by the Earth. For Earth orbiting missions, this delay is negligible, but for interplanetary missions, this can be several tens of minutes to many hours long. Thus, it is always necessary that the delay be accounted for in performing navigation functions. In addition to this physical limitation, there is also the delay needed to process the data once it gets to the ground which includes the time needed to perform the orbit determination and maneuver calculation computations, analyze the results, convene meetings and make decisions on what to do, implement the decision in terms of spacecraft commands, and finally transmit the commands and execute them on the spacecraft. The end-to-end process can be very time consuming and typical mission scenarios call for many days to over a week to perform these functions for non-time critical maneuvers or sequence commanding, down to many hours for very critical events. Even best case scenarios require over eight hours for the process. As a result, the latest and best information about the spacecraft's location is not used to execute the current event; the net loss could be for mission parameters such as fuel needed for the mission, or for science as the instruments aren't pointed exactly where they should due to imprecise knowledge of where the spacecraft is. Often, it's a loss for both, such as not having the best spacecraft control of atmospheric entry parameters to achieve the most accurate landing on a planetary body such as Mars.

One obvious way to overcome this limitation is to perform some, or all, of the navigation functions onboard the spacecraft. This not only eliminates the light-time delay but also circumvents the human-related delays for performing the navigation functions as described above. Thus, the turnaround time can be theoretically reduced to minutes, or even seconds, for reacting to late-breaking navigation information. A system to perform the navigation function autonomously onboard the spacecraft can enable certain classes of missions and greatly enhance science return on others. At JPL, we have developed such a system, called AutoNav, and applied it to several missions. A

¹ Supervisor, Outer Planet Navigation Group, Mission Design and Navigation Section, MS 264-820, 4800 Oak Grove Dr., Pasadena, CA 91191.

particularly useful application of AutoNav has been in the small body (asteroids or comets) flyby scenario, where the ephemeris of the target body is a major source of error. Using standard ground navigation techniques, the sequence of pointings to track the object through closest approach must be designed ahead of time, and the knowledge of the target-relative state of the spacecraft is not accurate enough to point the narrow angle cameras exactly at the target. Thus, mosaics were used to cover the uncertainty ellipse of the position estimate by rastering the camera across it, resulting in a substantial portion of the images being of empty sky. Onboard updates of trajectory knowledge enable rapid turnaround for closed-loop control, and in all the flybys that used AutoNav, most or all of the images taken had the target body in the camera FOV. If the objective is to impact a small body (as in the case of the Deep Impact mission), the mission simply could not have been accomplished without AutoNav due to the same lack of accurate target-relative navigation information as in the flyby case.

In this paper, we will describe the basic functions and algorithms used by AutoNav, its use to date, and present cases where it can be very useful in the future.

II. AutoNav Overview

Although in principle AutoNav can be used with any data type, in practice it is most straightforward to use data that are self-contained onboard the spacecraft. Radiometric data requires either a link with stations on the Earth (and thus having to deal with issues such as scheduling the station, and modeling time varying atmospheric delay effects on the radio signal), or other spacecraft (thus having at minimum some knowledge of where the other spacecraft is). Self-contained data, on the other hand, only requires controlling instruments on the spacecraft itself and includes such data as optical images, lidar, or radar altimetry. The current AutoNav system relies exclusively on optical images as the data type, so the remainder of the paper will focus on the use of this data.

The principle behind an optical-based navigation system is that the images taken of a natural body, either from a distance, or close-up, provide one or more line-of-sight (LOS) vectors to that body, or to locations on that body. So, for example, to navigate a spacecraft in an interplanetary cruise scenario, images can be taken of visible asteroids or the planets. By using a centerfinding technique, the brightness center of the unresolved (the extent of the target body is less than a CCD pixel) target object provides the inertial LOS vector. The inertial pointing direction of the camera is obtained from the stars in the frame. A succession of these images provides multiple LOS vectors which are then input into a non-linear least-squares filter to estimate the spacecraft's position and velocity. The accuracy of this method depends on many factors, including the camera system parameters, the distance to the target objects, the frequency of the images, and the knowledge of the ephemerides of the targets. An example of the accuracy of this technique as it was used on the Deep Space 1 (DS1) mission will be given later.

For the flyby or impact scenario (which forms the bulk of AutoNav use to date), the images taken are solely of the target body, starting minutes to hours before closest approach. In this scenario, any error in the heliocentric trajectory of the target body and spacecraft are folded into the spacecraft's trajectory error relative to the target alone. Once again, centerfinding on the body in a series of images provides a sequence of LOS observations. Here, however, the target will become resolved (its extent will be greater than a detector element (pixel)) at some point in the sequence and so the method for centerfinding will be slightly different. The initial images provide a direct measure of the closest approach location in the crosstrack (perpendicular to the incoming asymptote) axes. As the spacecraft gets nearer to closest approach, the changing geometry provides very accurate information about the spacecraft's location in the downtrack, or time-of-flight, axes. Examples of this use for the DS1, STARDUST, and Deep Impact (DI) missions will be provided later.

For the orbiting scenario, the images are of the extended body being orbited, and the body either fills a substantial portion of the camera frame or extends considerably beyond it such that each image only sees a portion of the body. In this case, the objective is to locate landmarks on the body that have been previously identified and whose body-centered position is known to some accuracy. Analogous to the deep space cruise scenario described above, the coordinates of each landmark in the frame provide a LOS vector to that landmark. If three landmarks are available in a frame, then the 3-D position of the spacecraft and the three components of the camera boresight pointing can be solved for deterministically. More than three landmarks in a frame allows a least-squares fit of the kinematic position and boresight pointing vector at each image time. Combining a series of such measurements into a dynamic filter allows estimation of the complete trajectory of the spacecraft. This technique can be especially

powerful in the proximity of small primitive bodies (asteroids and comets) where the gravitational pull is weak and therefore provides little signature in radiometric data; the geometric information provided by landmark tracking is essential to determining the orbit. The main restriction is that an *a priori* shape model with known landmark locations must be constructed prior to invoking AutoNav. Such techniques are available and documented¹.

All the above techniques can obviously be performed on the ground, with the penalty of turnaround time to get the solution. At its core, AutoNav simply takes the processes typically done by ground navigation teams and places these onboard a spacecraft. Thus, AutoNav must perform the following functions without human intervention: (1) process raw images into a form that can be used by the orbit determination filter, (2) perform orbit determination using the processed data by applying appropriate edits and weights and feeding the data into a least-squares filter, and (3) computing and executing maneuvers at appropriate times to guide the spacecraft to its destination. The product of these steps is the current, past, and predicted ephemeris of the spacecraft that can be used by other spacecraft subsystems, such as the Attitude Control System (ACS) or science instrument subsystems. Each of these AutoNav functions will now be described in detail.

A. Image Processing

The first step in AutoNav is to process images and convert them into an observation data point that can be input to an orbit determination filter. The images are obtained from an onboard camera whose various physical parameters (e.g., focal length, aperture) are determined by a combination of science and navigation requirements. In modern camera systems, the light is focused onto a Charge-Coupled-Device (CCD) mounted on the focal plane. The size of pixels in the CCD array, combined with the camera focal length, determine the resolution of the image and hence, the metric accuracy obtainable by the camera data. The number of pixels determines the total field-of-view. A filter wheel is often mounted in front of the lens to get color images, but for navigation, the clear filter is always used to maximize the light throughput and get the highest signal with the shortest integration time in order to reduce smear caused by spacecraft motion, as well as to reduce the exposure to cosmic rays and camera noise build-up. The dynamic range of the camera is determined by the digitization level, which in the past has varied from 8 to 12 bits per pixel for missions that used AutoNav.

Regardless of the type of image being processed (unresolved object centroid location, resolved object centroid location, or landmark location), the optical observation data used by the OD filter are the (x,y) coordinates of the location in the detector, here referred to as the sample and line coordinates, in the camera frame. The basic steps used to determine the coordinates is the same as it is for ground-based optical navigation processing and depends on the image type².

1. Unresolved Objects

For unresolved objects (which include the stars and solar system bodies when they are distant), the signal from the light emanating from the object is effectively convolved with the pointspread function of the camera optics to create an image which extends over several pixels. The pointspread function can be approximated by a Gaussian function; by measuring the pointspread of bright, known stars, the parameters of the Gaussian function can be estimated. This information can then be used to locate the position of other stars and target object by overlaying the Gaussian function on the image and adjusting its position in a least squares sense to best match the image. For ground processing, the initial locating can be done by eye; onboard, the *a priori* location of the stars and object can be used to determine an approximate location, and then a brightness centroid in the local vicinity can be performed to get a rough location. The Gaussian fit is then applied to fine tune the center location; the resultant accuracy is typically 0.05 to 0.1 pixels.

It may turn out, however, that the exposure duration needed to for stars and/or Solar System objects to have a sufficient signal is longer than the ability of the spacecraft to maintain a stable attitude. In this case (as it was for DS1), the images are smeared such that the signal no longer appears as a Gaussian. A more sophisticated centerfinding method was developed for this purpose, originally to support the Galileo flyby of asteroid Gaspra³. In this method, a multiple cross-correlation method was used to accurately determine the pixel/line coordinates of the target relative to the stars.

2. Resolved Objects

During flybys or the early phases on approach, the target object at some point becomes resolved. If the shape and orientation of the body is known, then a predicted image, given an initial guess of the location of the body relative to the spacecraft and the shading, can be generated and compared against the actual image. By correlating the predicted and observed image, the sample/line coordinate of the body can be determined as the point where the correlation is maximized, and this can be done to sub-pixel accuracy.

In typical flybys or approaches for primitive bodies, however, the shape and orientation are not known. In this case, a brightness centroiding method can be used to determine a less accurate center location. The process must be robust such that the centroiding does not pick up extraneous bright spots, such as stray light or cosmic rays, so we have developed methods to mitigate these failure modes. One method in particular is dubbed the “blobber”. The blobber searches the image for all contiguous bright spots that are between a minimum and maximum brightness threshold⁴. Each “blob” is then listed and ordered in terms of size. Since the approximate size of the object is known, any blob which falls outside the range of possible sizes is deleted. The largest remaining blob is picked as the target, and a simple brightness moment algorithm is applied to the circumscribing box around the blob to get the centroid. This process has been found to be very robust in numerous simulations and in actual flight. The accuracy of the centroid in determining the center of the target body can range from a sub-pixel level when the object is small in the FOV, to several tens or even hundreds of pixels as the object gets large. The reduced accuracy is still sufficient, however, to get accurate determination of flyby parameters.

3. Extended Object Landmark Tracking

During extended operations in the proximity of a target, landmark tracking techniques can be used in an AutoNav system to get very accurate orbit information. The first step is to create a detailed shape and albedo model, which must be done on the ground. A technique developed recently (called OBIRON – OnBoard Image Registration and Optical Navigation) uses stereo-photoclinometry to generate the shape and albedo model, and tested on the asteroid Itokawa during the Hayabusa mission¹. Once this shape model is created, for onboard use either part or all of the model parameters can be stored. As images are taken during the spacecraft’s path around the body, the scene in the image is compared against a computed scene based on the model; OBIRON will then locate and correlate landmarks on the surface. Each identified landmark then becomes a LOS vector, which as described above, can be input into an OD filter. This technique has yet to be demonstrated onboard, but has been used successfully on the Dawn mission for ground navigation. Prototype versions of AutoNav that include OBIRON have also been extensively tested in ground-based simulations.

B. Orbit Determination

Orbit determination is the process of combining a set of measurements to estimate the spacecraft’s trajectory (position and velocity) in a least-squares filter, as well as any ancillary parameters that affect the observations or the trajectory. Examples of the former are spacecraft attitude or camera bias errors; the latter includes such things as solar pressure scale factors, or miscellaneous un-modeled accelerations. The OD process in standard ground-based navigation is well known and documented⁵, but there are some factors to consider when applying it to an onboard system. The important ones will be described here.

4. Data Editing

Large outliers in the measurement data set, whether random or systematic, can corrupt an OD solution if not removed prior to filtering. Thus, instead of using data immediately as they become available, we first store a series of measurements. After a predefined set of measurements have been taken, the data are statistically analyzed to remove outliers. If, after some length of time, enough data has not been accumulated for whatever reason, the filter is not invoked, with the philosophy being that it is better to stick with a previous, known good solution that might be stale as opposed to risking getting bad data into the solution. The parameters chosen for the data editing are mission and scenario specific, and are usually determined by running numerous Monte Carlo simulations on the scenario to see what parameter set yields the highest probability of success. Because the optical data set is very sparse (data appears at the rate of one every several minutes as opposed to 1-10 Hz for radiometric data), the choice of editing

parameters is crucial and a good portion of ground testing is devoted to this task. Once the data passes the statistical tests, they are passed to the orbit determination filter.

5. Force Modeling

For ground-based navigation, the force models that describe the trajectory of the spacecraft are chosen to be as high fidelity as possible since radiometric tracking data are of very high accuracy. Also, with modern computer systems, the speed of the numerical integration (both trajectory and state transition matrix) and partial derivative generation is not usually a constraint. Onboard, however, speed does become an issue since spacecraft processors are not generally as fast as the latest available on the ground, and rapid turnaround is important for certain scenarios, such as high-speed flybys. We thus limit the force modeling for AutoNav to include central and third body point mass gravitational accelerations, simple flat plate or sphere model for solar radiation pressure, and instantaneous ΔV s for chemical propulsion maneuvers, and linear polynomials for low-thrust maneuvers obtained from ion propulsion systems⁶. To get the latest information about thrusting events in the past, thruster activity is converted to a “non-grav history file”, which accumulates and records the ΔV as a function of time; this file is read into the trajectory propagation as discrete events which change the instantaneous velocity at the recorded time.

A faster AutoNav version was also created specifically for use on comet and asteroid flybys, and used on DS1 and STARDUST. This version, dubbed the Reduced State Encounter Navigation (RSEN), simplified the trajectory by modeling it as straight line motion past the body. At the time AutoNav was initiated (tens of minutes prior to encounter), the trajectory was initialized with the best ground-based position and velocity, and the velocity was not updated through the flyby. For these effectively massless bodies, this model has sufficient accuracy, as was proven in flight⁷.

6. Filtering

Filtering is the processing of obtaining a least-squares fit to the trajectory using the available data. For many applications in the Aerospace field, a Kalman filter formulation is used, where, as each data point becomes available, it is used to update the filter. This works well for situations where very high frequency data are available and rapid filter updates are needed, such as for attitude control or missile guidance, but for most deep space applications, this speed is not necessary. Furthermore, as was described above, optical-based AutoNav has a sparse data set and needs to be pre-edited. For this reason, we chose a batch filtering approach, where data are stored, edited, and then input to the filter. The filter uses variations on the batch formulation which solves for the state at some epoch time. After iterating the filter to convergence, the epoch state solution is integrated forward to get the complete trajectory from epoch to some future time. Details on the filtering method used by AutoNav are provided in Ref. 6.

C. Maneuver Planning and Execution

Once an OD solution is performed and the current best estimate of the trajectory available, it must be evaluated to see if the predicted path satisfies the mission requirements. If it doesn't, then maneuvers are needed to re-target the course. The current version of AutoNav can handle two types of maneuvers, low-thrust (such as for an ion propulsion system), and near instantaneous (for standard chemical propulsion systems). In both cases, the optimization of the maneuvers needed to minimize fuel consumption, flight time, or other parameters, is done on the ground. Thus, the times and (in the case of low-thrust) overall duration of the burn is pre-planned. AutoNav will then perform adjustments to the nominal burn in a linearized fashion; the partial derivative of the burn parameters with respect to the target parameters are evaluated on the pre-defined reference trajectory. This Jacobian matrix is inverted to get the values of the burn, with several iterations being necessary to converge the solution. It is possible that the solution will not converge if the deviation from the reference is large, and in this case, the ground must intervene. It is envisioned that future versions of AutoNav will have the ability to do its own re-optimization of the trajectory if the course has strayed far from the reference.

For low-thrust maneuvers, the ΔV imparted can last anywhere from many hours to several months, depending on the mission phase and time to go relative to the target condition. The approach used was to break up the ground optimized continuous thrust profile into linear segments, with each segment consisting of the Right Ascension (RA)

and Declination (DEC) of the thrust vector, and the magnitude of the thrust. Using the linearized method described above, the RA/DEC of multiple segments were adjusted from the nominal, based on the current best OD solution⁸.

Impulsive burns are much simpler. Each maneuver can be solved for independently, with the three parameters of the burn able to control two or three target parameters. For flybys of primitive bodies, the target parameters are the location in the “B-plane” – a plane perpendicular to the incoming trajectory and centered on the body – and the time of closest approach. If the encounter time is not critical (such as for an impact), then just the two parameters in the plane can be targeted⁹. For proximity operations scenarios around a body, the target parameters can be waypoints along a reference trajectory or the landing location on the body¹⁰.

III. Mission Results

AutoNav has been used successfully on five missions: Deep Space 1, STARDUST, Deep Impact, EPOXI (the follow-on mission to Deep Impact), and STARDUST NEXt (the follow-on mission to STARDUST). A brief description of these missions and how AutoNav was used on them will be given in this section. Since fast flybys were the main use of AutoNav in these missions, Table 1 provides a concise tabulation of the various parameters of the flyby, including the approach phase angle, flyby velocity, and flyby distance.

Table 1. Summary of flybys which have used AutoNav to date.

Mission/Target	Flyby Radius (km)	Flyby Velocity (km/s)	Approach Phase (deg)
DS1/Borrelly	2171	16.6	65
STARDUST/Annefrank	3076	7.2	150
STARDUST/Wild 2	237	6.1	72
DI/Tempel 1	500	10.2	62
EPOXI/Hartley 2	694	12.3	86
STARDUST NEXt/Tempel 1	182	10.9	82

D. Deep Space 1

DS1 was the first mission in the NASA’s New Millennium Program. This program was intended to fly unproven, high-risk technologies and validate them in flight so that subsequent science-focused missions could then use these technologies. Twelve such technologies were flown on DS1; among them were the first deep space ion propulsion system, a low mass combined visible/infrared/ultraviolet imaging system (named MICAS – Miniature Imaging Camera and Spectrometer), and the first onboard autonomous navigation system (AutoNav). The spacecraft was launched in October 1998, and the original mission plan included flybys of the asteroid Braille and the comets Wilson-Harrington and Borrelly. During flight, however, the failure of the only onboard star tracker following the Braille encounter resulted in a long period where onboard software was modified to fly the spacecraft without the star tracker, resulting in the loss of the Wilson-Harrington encounter. DS1 flew by Braille in July 1999 and Borrelly in September 2001.

For AutoNav, the original plan was to use it to navigate the cruise and flyby portions of the mission. OD for the cruise portion would be handled by using bright asteroids in the asteroid belt as beacons as described in Section II. The choice of which asteroids to choose was done on the ground and uploaded to the spacecraft. At roughly weekly intervals, a 4-6 hour block of time would be used to cycle through a list of pre-planned beacons and image each of them. The centroiding of the stars and asteroid provided the raw data for the OD filter, which then computed the spacecraft’s trajectory. Weekly segments of ion thruster pointing direction vectors for the future were then computed by the maneuver planner. Pre-flight analysis of this process indicated that the accuracies obtainable from the system were sufficient to guide the spacecraft to the vicinity of the next flyby target body⁶.

These plans had to be abandoned after the initial images from MICAS revealed several unforeseen camera problems that affected how it could be used by AutoNav. The three biggest issues were that the pictures were heavily corrupted by stray light, the sensitivity of the camera was not at the expected level, and the camera distortions were unusual and proved to be difficult to model¹¹. This resulted in many months of reprogramming AutoNav to handle these problems. As a result, the validation of onboard OD did not occur until the summer of 1999, a few months prior to the encounter with Braille. From June through end of July 1999, until preparations for the encounter were started, AutoNav was performing OD and maneuvers and the results sent to the ground. Since standard radiometric ground-based navigation was being performed concurrently, we could compare the results of AutoNav with the ground “truth”. The accuracy of the solutions as compared to the ground is plotted in Fig. 1, which shows the OD was good to roughly 500-2000 km, and 0.2-0.8 m/s. The formal uncertainties of the solution are also plotted, indicating the error was at the level of the uncertainty, or as good as can be expected. However, due to the camera anomalies, this wasn’t as good as the pre-flight planning had indicated, so some of the targeting maneuvers were done with ground intervention in order to achieve the Braille flyby.

The Braille flyby also was problematic. Due to an AutoNav software error, the spacecraft experienced a safing event 18 hours before encounter. After recovery, maneuvers were planned from the ground to retarget the flyby. At about 30 minutes prior to closest approach, the RSEN version of the code was initiated to begin closed-loop tracking of Braille. RSEN was, at that point, updating the OD after every image, but using images from a not-well calibrated alternate APS (Active Pixel Sensor) detector. In the first image taken of the asteroid, the asteroid itself was too dim

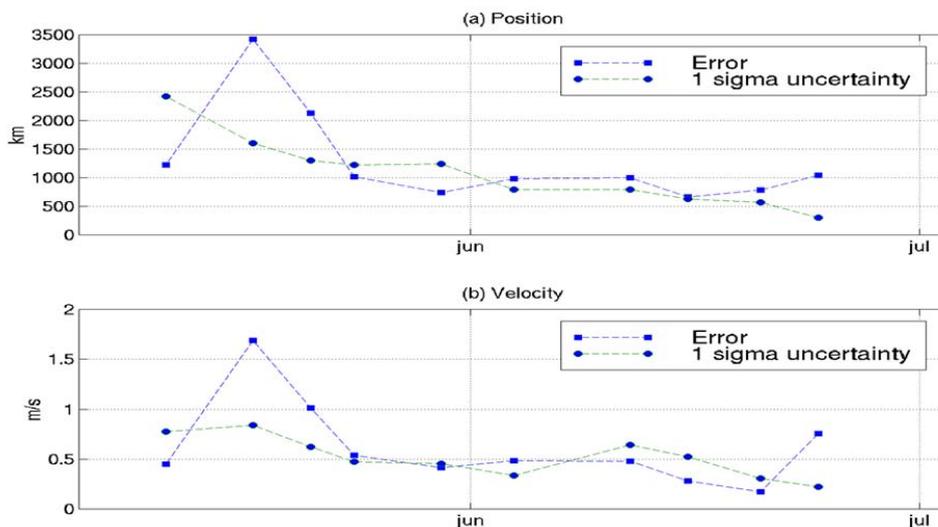


Figure 1. Orbit determination results from DS1 cruise¹¹.

to detect, but a random cosmic ray was and resulted in RSEN updating the spacecraft state from erroneous data. This effectively took the asteroid out of the camera field-of-view; no further updates occurred since no asteroid was subsequently detected, and the closest approach images were lost. A major outcome of this failure was the restructuring of the code so that rigorous error checking was performed before ingesting the data. It was also subsequently determined that the APS detector was non-functional for anything but the brightest signals

Following the flyby, AutoNav was once again put in control of the low-thrust “cruise” to the next encounter. However, in August 1999, the sole onboard star tracker failed. Since the spacecraft relied on the star tracker to provide attitude knowledge, the mission had to be put on hiatus until a workaround using MICAS as an alternate star tracker could be found. The mission was restarted late in summer of 2000, but the Wilson-Harrington encounter was lost and AutoNav could no longer be used for cruise.

By the summer of 2001, the revised RSEN AutoNav code was ready for the Borrelly flyby planned for September. Among the revisions was a change to accumulate data in batches and pre-edit data before use by the filter, robust error checking, and the use of the “blobber” to weed out random cosmic rays and identify the actual

target. The flyby itself occurred on September 22, 2001. RSEN was initiated at Encounter (E) – 32 minutes and initialized with ground-based navigation information from E – 12 hours. The uncertainty in the spacecraft’s comet-relative position was roughly 15 km the cross-track directions, and 200 km (or equivalently, 12 seconds) in the downtrack directions. Unlike the Braille encounter, AutoNav performed as expected, capturing the nucleus in 45 out of the 52 total images shuttered. The closest image was taken at E – 2 min, 45 sec, at a distance 3556 km and resolution on the surface of 46 m/pixel (these were the highest resolution images of a comet taken to date). After this time, the comet drifted out of the camera FOV due to known limitations in the spacecraft’s ability to turn the spacecraft at a fast enough rate. Following the flyby, both the mainline and RSEN AutoNav were validated as part of DS1’s technology demonstration.

E. STARDUST

STARDUST was the fourth mission in NASA’s Discovery Program. The primary goal of the mission was to fly through the tail of comet Wild 2 and return samples of the coma back to Earth, with a second goal to also capture high resolution images of the comet’s nucleus during the flyby. The spacecraft was launched on February 7, 1999, encountered Wild 2 on January 4, 2004, and successfully returned a coma sample to the Earth on January 6, 2006. En route to Wild 2, the spacecraft also flew by asteroid Annefrank, which was used as an engineering demonstration to practice the Wild 2 flyby.

As for DS1, the RSEN flavor of AutoNav was used for closed loop tracking of the nucleus during the flyby. The development of AutoNav for STARDUST closely matched that of DS1, and the lessons learned from the Braille and Borrelly flyby’s were incorporated into STARDUST. Algorithmically, the code was nearly identical to the final DS1 version, with the primary difference being that STARDUST used a unique camera/scan mirror combination for imaging. The camera itself was fixed along one of the spacecraft axes, but a mirror was placed in front of the camera lens that could rotate about 180 deg about the boresight axis. This allowed the camera FOV to sweep along a plane from the front to the rear of the spacecraft and thus, the spacecraft attitude to remain fixed while the mirror handled the motion of the comet, provided the comet was in the sweep plane⁷. The modifications in the code from the DS1 version were mainly for accommodating this camera geometry.

The opportunistic encounter with Annefrank occurred in November 2002 which allowed a practice run for AutoNav before the main event. This was in some ways more challenging than Wild 2, primarily because the approach phase angle to the asteroid was 150 deg, resulting in poor visibility on approach. In fact, the asteroid was never seen in ground processed approach Optical Navigation images, and AutoNav was initiated at E – 20 minutes with no ground updates to the target-relative ephemeris. Nevertheless, AutoNav did successfully track Annefrank through the flyby, providing extra confidence that Wild 2 would also be successful.

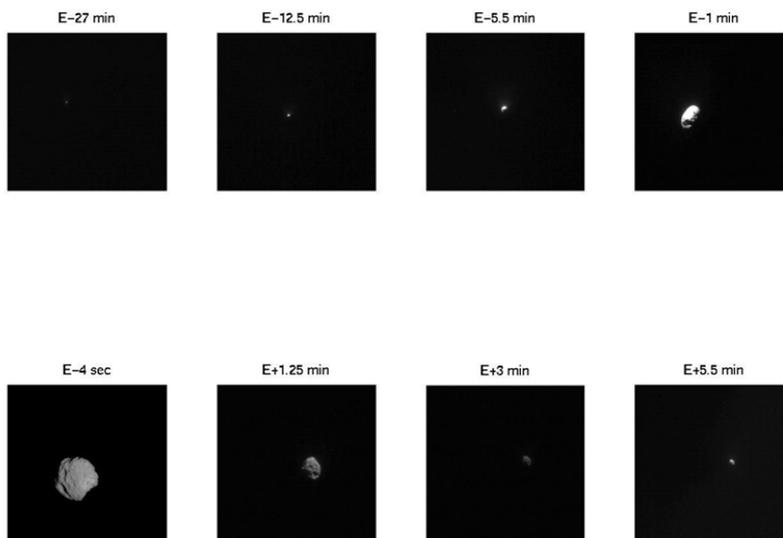


Figure 2. Sequence of STARDUST AutoNav images of Wild 2 surrounding encounter.

For Wild 2, AutoNav was initiated at E – 30 min using the best ground-based target relative trajectory solution at E-48 hours. The cross-track uncertainty was around 5 km, and the downtrack was 2000 km (or equivalently about 50 seconds). The large downtrack uncertainty was a consequence of poor ephemeris knowledge of the comet from Earth-based astrometry, due to the large distance to Earth at the time of encounter (> 2 AU), and the low elevation, resulting in poor astrometry. After initiation, AutoNav took images at 30 second intervals; the first OD update after accumulation of the data was at E-10 min. In

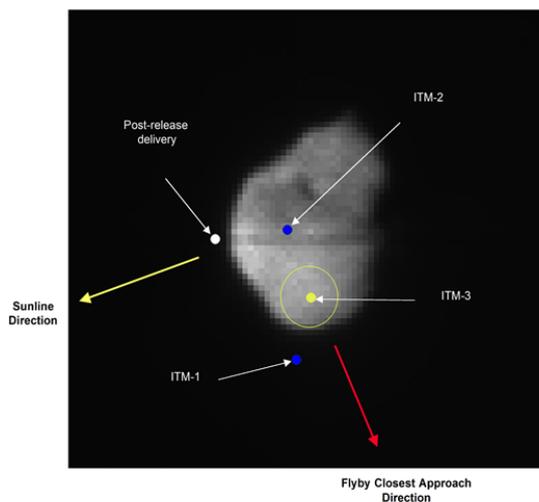


Figure 3. Impact points after each maneuver on DI Impactor spacecraft⁹.

high speed impact of a comet with one spacecraft while observing the event with another. The flyby/impactor combination was launched on January 12, 2005 for a 6 month cruise to impact comet Tempel 1 on July 4, 2005. Several maneuvers in the weeks leading up to encounter targeted the dual spacecraft for the impact. Twenty-four hours before encounter, the impactor separated from the flyby spacecraft; the latter then performed a large deflection burn to target a 500 km altitude flyby and slow the spacecraft down so it would arrive at closest approach 10 minutes prior to the impactor. Both spacecraft performed their functions successfully and images from the impactor as it approached Tempel 1 and the flash resulting from the impact were sent to the ground for analysis.



Figure 4. Impact as viewed from DI Flyby spacecraft.

on DI were impulsive chemical burns, rather than the low-thrust maneuvers on DS1. The impulsive ΔV logic was simpler than that for low-thrust, so this modification was relatively simple. DI also had a few additional capabilities needed. These included code to perform: (1) “scene analysis” – selecting the actual impact site to be on a lit area of the nucleus, (2) time-of-impact and time-of-imaging updates – these synchronized the timing of the imaging sequence starting on the flyby based on the Impactor’s expected impact time, and (3) autonomous coma cutoff – a modification to process the comet images for ensuring that the brightness centroid is on the nucleus and not

order for the flyby plane to be aligned with the mirror sweep plane, the OD from AutoNav was used at E-4 min to perform a spacecraft roll to make this alignment. Following the roll, the images were taken at a faster rate, every 20 seconds, to better capture the close approach images. AutoNav was terminated 8 minutes past the nominal time of closest approach.

Examination of the post-encounter telemetry and images indicated AutoNav had successfully tracked Wild 2, with 114 total images taken and all capturing the nucleus. The closest image was shuttered at E – 4 seconds at a distance of 239 km and a resolution of 14 m/pixel. Fig. 2 shows a sample of the sequence of images surrounding closest approach.

F. Deep Impact

DI, the eighth mission in NASA’s Discovery Program, had the most challenging use of AutoNav to date. The goal of the mission was to perform a

DI used standard ground radiometric and optical navigation techniques for the cruise and approach phases of the mission. AutoNav was used only in the hours prior to the impact event and had dual roles. On the Impactor, AutoNav needed to use its own onboard calculated maneuvers to guide the spacecraft to a sunlit area on the comet nucleus, and bias the impact site towards the side where the flyby spacecraft had its closest approach. On the Flyby spacecraft, no maneuvers were computed but AutoNav needed to maintain camera lock on the nucleus and predict where the impactor would hit, as well as determine the time of closest approach and initiate a high-rate imaging sequence. Since maneuvers were needed on the Impactor, the full AutoNav version was used, and then to maintain consistency, the same version was used on the Flyby.

The bulk of the computational elements of the DS1 AutoNav code was used without modifications on DI. One component that changed was the maneuver computations

corrupted by coma¹². Because AutoNav was critical to the success of the mission, extensive ground testing was performed to verify that AutoNav would work as planned. Since the lack of knowledge of what the comet nucleus would look like was a large error source, much of the testing involved simulating all types of comet nuclei to see how AutoNav would respond. One result of the simulations discovered after the spacecraft had launched was that subtle interactions between the attitude estimated by the Attitude Control System and AutoNav estimates of the comet-relative velocity could cause the targeting maneuver to behave somewhat erratically. Nevertheless, overall, the simulations indicated that, barring an unforeseen failure in the flight system, the probability of success for the impactor was greater than 99%¹².

Based on these results, the scenario for AutoNav on the Impactor was baselined. AutoNav was initiated at E – 2 hours, with images taken at a rate of 4 images per minute. After 10 minutes of accumulating data, the first OD was performed; subsequent OD solutions were computed every minute. Three targeting maneuvers were executed at E – 90 min, E – 35 min, and E – 12.5 min; their magnitudes were 1.27, 2.26, and 2.29 m/s, respectively. Post-flight reconstruction showed that, as was predicted, the first maneuver actually moved the spacecraft further away from the target than prior to the maneuver, but subsequent maneuvers successfully targeted to the nucleus, as can be seen in Fig. 3⁹. The Flyby spacecraft also performed as planned, imaging the flash as the Impactor hit the nucleus (Fig. 4). The image was taken at a distance of 887 km at E – 72 seconds. AutoNav was terminated at E - 50 seconds so that the spacecraft could go into “shield mode” to protect the instruments from comet dust impacts through closest approach.

G. EPOXI

EPOXI was the follow-on mission to DI. It was selected as a NASA Discovery Mission of Opportunity flight in the summer of 2007. With the announcement of the selection, the DI flyby spacecraft, which had been placed in hibernation following the Tempel 1 encounter, was reactivated for a planned flyby of comet Hartley 2 in November 2010. Once again, ground-based navigation was used for the cruise and approach, and AutoNav was used just to track the nucleus on the flyby.

Originally, it was assumed that running AutoNav for EPOXI would be relatively simple since there was no longer an Impactor spacecraft, the Flyby had performed successfully on the Tempel 1 encounter, and no code

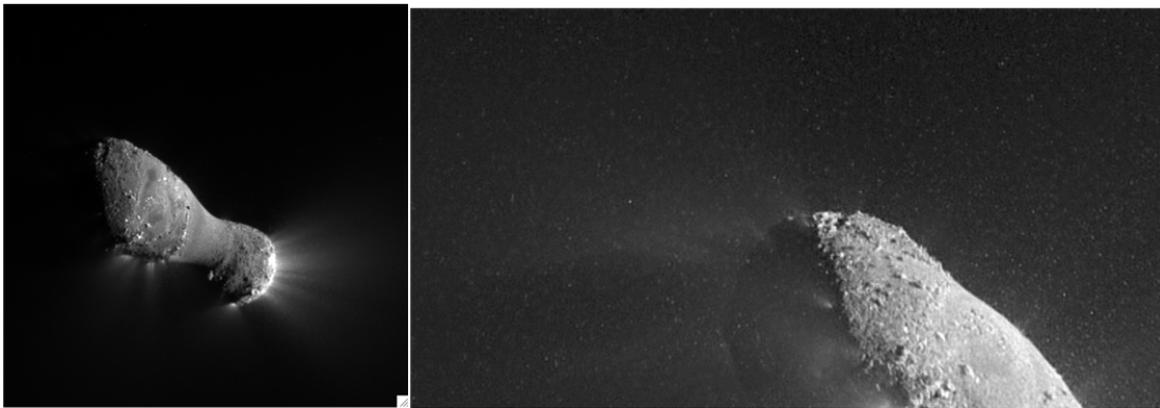


Figure 5. Hartley 2 close approach image taken by EPOXI. Left panel shows the nucleus taken by the medium resolution imager, right panel shows an enhanced view from the high resolution imager that reveals the particles ejected from the nucleus.

modifications were needed, only parameter updates. However, the detailed analysis of the changes needed in the parameters to accommodate the different flyby geometry revealed some unforeseen complications. In particular, for the DI case, AutoNav was not engaged through the point of closest approach since the spacecraft went into shield mode. For Hartley 2, the minimum flyby altitude was roughly 700 km; any closer and the spacecraft ACS could not keep up with the turn rate. At this distance, analysis indicated that comet dust impacts were not of high concern, so the instruments could keep taking data throughout the encounter. Simulations of AutoNav performance in this time

period showed that the same ACS/AutoNav velocity estimate interactions that were problematic for the Impactor on DI also caused problems when the spacecraft went through a 180 turn through encounter. Furthermore, the interactions with the spacecraft Fault Protection software also had to be carefully choreographed so that unintended behavior did not occur. Details of these issues and how they were resolved can be found in Ref. 13.

The thorough analysis paid off in the end and AutoNav successfully tracked Hartley 2. One final surprise was that in the images nearest periapse, Hartley 2 was offset by almost a third of the camera frame from the center. The telemetry from AutoNav and the ACS, however, showed that everything had worked smoothly and the offset errors should have amounted to far less. It was finally discovered that the biasing on the DI Flyby to account for the Impactor hitting on the lit hemisphere was still resident in the spacecraft's memory and had never been cleared. Fortunately, this bias was relatively small compared to the flyby distance; had it been larger, the close approach images might have been lost. In the end, the bias turned out to be fortuitous for science; the discovery of large chunks of ice ejected from the comet was found due to the offpoint¹⁴. Fig. 5 shows the closest image and an enlarged image of the ice discovery.

H. STARDUST NExT

NExT was the follow-on mission to the STARDUST mission. It was also a Discovery Mission of Opportunity, selected at the same time as EPOXI. NExT's purpose was to revisit comet Tempel 1 and image the crater that DI created since the flash caused by the Impactor obscured a clear view of the impact site from the DI Flyby spacecraft. Another mission goal was to image more of the surface of Tempel 1. This would mark the first time in history that the same comet was visited more than once, and viewed before and after a perihelion passage. The mission used ground navigation for cruise and approach, and the same AutoNav code used on Annefrank and Wild 2 for Tempel 1.

Because, unlike on EPOXI, the system was used in nearly identical fashion to the earlier encounters, analysis for the Tempel 1 encounter was straightforward. This flyby was considerably faster (10.9 vs 6.1 km/s), and was targeted to be a little closer (180 vs 250 km altitude), and so the parameter settings were adjusted to optimize the probability of success. The two primary concerns were that the low flyby was at the limit of the scan mirror's fastest rate, and the image rate after the spacecraft alignment roll turn was more rapid in order to get high resolution stereo images of the impact site. The former was a problem in that, if the rate exceeded the mirror's capability, the images would be smeared and corrupt the AutoNav solution. The latter concern had to do with the ground navigation ability to predict and control the time of closest approach; if the error in the predicted periapse time was larger than +/- 2 minutes, it would be beyond the high rate imaging sequence and much of the stereo information would be lost. As it turned out, the timing was only off by 15 seconds, but the final radius was at 181 km and thus, one of the near images was slightly clipped as the mirror did not quite keep up. Nevertheless, AutoNav performed as expected and all science imaging requirements were met. Fig. 6 shows the closest approach image, shuttered at E + 3 sec at a distance of 184 km.

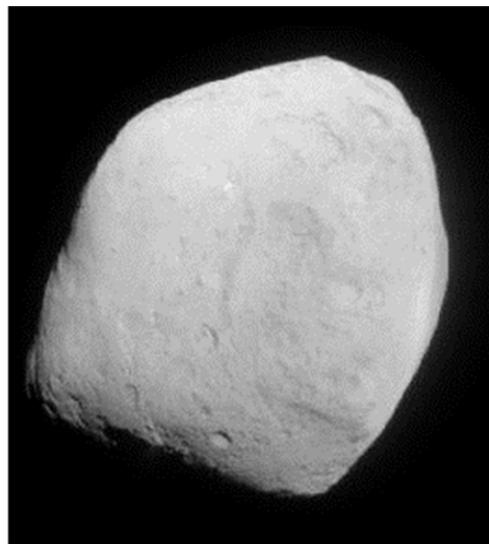


Figure 6. Tempel 1 close approach image taken by STARDUST NExT.

IV. Future Mission Uses for AutoNav

The original AutoNav code developed on DS1 has proved to be quite versatile, working successfully on five missions and four different spacecraft built by three different contractors. The code, however, is now well over a decade old, and the limitations imposed by the inexpensive, rapid development, needed to be fixed for future uses. Over the past several years, this has been undertaken with the next generation AutoNav in prototype use¹⁵. This version builds on the experience gained from AutoNav's use in the missions described above, and incorporates many enhancements. The major ones are: (1) a new interface that is written in the Virtual Machine Language

(VML), which allows for greater flexibility and ease of use, and (2) the addition of attitude guidance and control merged with the current translational navigation capabilities, and (3) the landmark tracking capability using a flight version of OBIRON. This new version is now called AutoGNC to reflect the enhanced capability, and enables AutoNav use for scenarios where the translational and rotational motion are tightly coupled, such as landings or atmospheric flight. The code, however, is modular and flexible so that only the computational capabilities that are required for any specific mission need to be used. The following subsections describe possible missions where the use of AutoNav can be valuable.

I. Asteroid Deflection

The possibility of a Near Earth Asteroid impacting the Earth and causing large scale destruction is of low likelihood, but very high consequence should it happen. One realistic option for mitigating this threat is to deflect the asteroid well before its Earth impact. The DI mission demonstrated that the technology is already available using AutoNav. The modifications needed to hit a much smaller target (potentially hazardous asteroids are typically in the 100s of meters or less in diameter, as compared to the 7 km diameter of Tempel 1), and at potentially higher velocities (perhaps as high as 20-25 km/s) represents an evolutionary, not revolutionary improvement. A point solution analysis for a specific example has been performed¹⁶. A full system trade study, simulating the range of possible asteroid sizes and shapes, would help to define the spacecraft and AutoNav performance necessary to execute the deflection reliably.

J. Small Body/Lunar Pinpoint Landing

The NEAR and Hayabusa asteroid landings demonstrated that such missions are quite feasible using ground-in-the-loop navigation at the tens of meters accuracy^{17,18}. For future landings on asteroids or comets, it may be necessary to achieve accuracies of less than 5 m or better, either because of the lack of safe landing spots at larger scales, or to target very specific regions for science. Furthermore, it may also be necessary to tightly control the velocity at touchdown for spacecraft safety. This combination of requirements is very difficult, if not impossible, to do open loop from the ground due to the light time and other lags between the navigation knowledge update and control. AutoNav could be ideally suited for this type of mission, and Monte Carlo simulations have been performed that demonstrated this, achieving position control to within 3 m and horizontal velocity control better than 2 cm/s¹⁰. Simulations have also been performed for precision landings on the Moon, which show that landings to within 20 m are possible¹⁹.

K. Aerobraking

Aerobraking has been used on several Mars missions to modify the orbit of the spacecraft from its initial capture orbit to the final desired science orbit. The process is very operations intensive, and towards the end game, the frequency of atmospheric passes and maneuvers becomes very rapid, on the order of several hours. This requires large staffing to maintain round-the-clock coverage. Onboard AutoNav, using acceleration data from an onboard Inertial Measurement Unit, could reduce the staffing needed by maintaining the knowledge onboard and using it to control the drag corridor. AutoNav could also be used as a monitor to ensure spacecraft safety and perform an abort in case the orbit gets too low in the atmosphere.

L. Outer Planet Satellite Tour

The complex dynamics of the large satellite systems of Jupiter and Saturn allows for unique mission possibilities that can be enabled by AutoNav. By exploiting some of these dynamics, the delta-v can be reduced substantially for capture, or more satellites can be visited in a shorter amount of time²⁰. These options do require, however, short turnaround times between successive satellite flybys. These types of missions would be difficult to do from the ground, making them another candidate for onboard automation.

V. Conclusions

AutoNav is a technology that has matured to the point where its capability has been demonstrated in real-world applications repeatedly over the last decade. This paper described the basic workings of the system, and the examples of where it has been used. Some examples of future uses were also described, but there are many other situations that can be envisioned as well, and it is quite likely that descendants or variations on the system described will expand along with the range of mission possibilities.

Acknowledgments

The development of AutoNav would not have been possible without the contributions of a number of people. I am indebted to the following people who made substantial contributions to the concept, design, testing, and operations of AutoNav and which resulted in its successes. These include: Matt Abrahamson, Shailen Desai, Don Han, Brian Kennedy, Dan Kubitschek, Nick Mastrodemos, Bill Owen, Ed Riedel, Steve Synnott, Mike Wang, and Bob Werner.

The research described in this publication was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Copyright 2012 California Institute of Technology. Government sponsorship is acknowledged.

References

- ¹Gaskell, R. H., et al, "Characterizing and navigating small bodies with imaging data," *AIAA Meteoritics and Planetary Science*, Vol. 43, 2008, pp. 1049-1061.
- ²Owen, W. M., "Methods of Optical Navigation," *AIAA/AAS Spaceflight Mechanics Conference*, AAS 11-215, New Orleans, LA, February 2011.
- ³Vaughan, R. M., Riedel, J. E., Davis, R. P., Owen, W. M., and Synnott, S. P., "Optical Navigation for the Galileo Gaspra Encounter", *AIAA/AAS Astrodynamics Conference*, AIAA 92-4522, Hilton Head, S. C., August 1992.
- ⁴Russ, J. C., *The Image Processing Handbook*, CRC Press, 1999, Chapter 7.
- ⁵Tapley, B. D., Schutz, B. E., Born, G. H., *Statistical Orbit Determination*, Elsevier Academic Press, San Diego, CA, 2004.
- ⁶Bhaskaran, S., et al, "Orbit Determination Performance Evaluation of the Deep Space 1 Autonomous Navigation System", AAS 98-193, *AAS/AIAA Spaceflight Mechanics Meeting*, Monterey, CA, February 1998.
- ⁷Bhaskaran, S., Riedel, J. E., and Synnott, S. P., "Autonomous Target Tracking of Small Bodies During Flybys", AAS-04-236, *AAS/AIAA Astrodynamics Specialist Conference*, Maui, Hawaii, February 2004.
- ⁸Desai, S. D., et al, "The DS1 Autonomous Navigation System: Autonomous Control of Low-Thrust Propulsion Systems", AIAA 97-38819, *AIAA Guidance, Navigation and Control Conference*, New Orleans, LA, August 1997.
- ⁹Kubitschek, D. G., et al, "Deep Impact Autonomous Navigation: The Trials of Targeting the Unknown", AAS 06-081, *AAS Guidance and Control Conference*, Breckenridge, CO, February 2006.
- ¹⁰Bhaskaran, S., et al, "Small Body Landings Using Autonomous Onboard Optical Navigation", *Journal of the Astronautical Sciences*, Vol. 58, No. 3, 2012.
- ¹¹Bhaskaran, S., Riedel, J. E., Synnott, S. P., Wang, T. C., "The Deep Space 1 Autonomous Navigation System: A Post-flight Analysis", AIAA 2000-3935, *AIAA/AAS Astrodynamics Specialist Conference*, Denver, CO, August 2000.
- ¹²Mastrodemos, N., et al, "Autonomous Navigation for Deep Impact", AAS 06-177, *AAS/AIAA Spaceflight Mechanics Conference*, Tampa, FL, January 2006.
- ¹³Abrahamson, M., Kennedy, B. M., Bhaskaran, S. B., "AutoNav Design and Performance for the EPOXI Hartley 2 Flyby, *SpaceOps 2012 Conference*, Stockholm, Sweden, June 2012.
- ¹⁴Hearn, M. F., et al, "EPOXI at Comet Hartley 2", *Science*, Vol. 332, No. 6036, 17 June 2011, pp. 1396-1400.
- ¹⁵Riedel, J. E., et al, "Configuring the Deep Impact AutoNav System for Lunar, Comet and Mars Landing, AIAA 2008-6940, *AIAA/AAS Astrodynamics Specialist Conference*, Honolulu, Hawaii, August 2008.
- ¹⁶Bhaskaran, S., et al, "Navigation Challenges of a Kinetic Energy Asteroid Deflection Spacecraft", AIAA 2008, *AIAA/AAS Astrodynamics Specialist Conference*, Honolulu, Hawaii, August 2008.
- ¹⁷Antreasian, P. G., et al, "The Design and Navigation of the NEAR/Shoemaker Landing on Eros", AAS 01-372, *AAS/AIAA Astrodynamics Specialist Conference*, Quebec City, Canada, July 2001.
- ¹⁸Yano, H., et al, "Touchdown of the Hayabusa Spacecraft at the Muses Sea on Itokawa", *Science*, No. 5778, 2006, pp. 1350-1353.
- ¹⁹Riedel, J. E., et al, "Optical Navigation Plan and Strategy for the Lunar Lander Altair, OpNav for Lunar and other Crewed and Robotic Exploration Applications", AIAA-2010-7719, *AIAA GN&C Conference*, Toronto Canada, August 2010.
- ²⁰Lynam, A. E., Kloster, K. W., and Longuski, J. M., "Multiple-satellite-aided Capture Trajectories at Jupiter using the Laplace Resonance", *Celestial Mechanics and Dynamical Astronomy*, Vol. 109, No. 1, 2011.