

Autonomous optical navigation for interplanetary missions

S. Bhaskaran, J. E. Riedel, S. P. Synnott

Navigation and Flight Mechanics Section
Jet Propulsion Laboratory
California Institute of Technology

ABSTRACT

The automation of interplanetary spacecraft is becoming increasingly desirable to meet various mission requirements. A prototype of an autonomous spacecraft which will flyby an asteroid and comet is slated for flight in mid-1998 as part of NASA's New Millennium Program. This spacecraft will navigate by using optical data taken by the onboard camera to determine its orbit, and use this information to predict its future trajectory and make necessary course corrections. The basic navigation data type available from the camera are star- relative astrometric observations of solar system bodies which can be used to determine line-of-sight vectors to those objects. The directional sightings are obtained by determining the precise centers of the object and stars in the image. During interplanetary cruise, centerfinding is performed by using two pattern matching techniques inherited from the Galileo mission. Near-encounter images are processed with a separate algorithm employing image modelling and brightness centroiding. This paper describes the image processing algorithms, and the results of a ground-based test of the algorithms using real data.

Keywords: Optical navigation, autonomy, image processing

1.0 INTRODUCTION

Future plans being developed at NASA for deep-space exploration call for multiple, small, dedicated spacecraft operating autonomously or semi-autonomously over long periods of time to achieve specific science goals. In order to meet the size and cost constraints of these spacecraft, it is desirable to maximize the functions that the science instruments perform. In particular, the onboard camera used to image solar system objects during the cruise, approach, and orbiting phases of a deep-space mission can be used to navigate the spacecraft as well. Historically, such optical navigation techniques have been used only during the approach phase to planets or asteroids to supplement standard radio (Doppler and range tracking) navigation techniques. The images taken by the camera are transmitted to the ground, where they were processed and merged with radio navigation data to complete an orbit determination solution of the spacecraft's trajectory. Maneuvers needed to correct the spacecraft's trajectory are also computed on the ground, and then sent up to the spacecraft for execution. For future "sciencecraft" missions, however, constraints imposed by the number and types of missions being flown will preclude this personnel-intensive mode of operation, requiring that some or all of these functions be placed onboard the spacecraft itself. The optical data type is well suited to being acquired and processed autonomously, and forms the basis for a completely autonomous navigation system. Such a system is currently being developed for the New Millennium Program's Deep Space 1 (IX-1) flight. A description of a prototype of this complete system can be found elsewhere¹. In this paper, we will give a brief synopsis of the fundamentals of orbit determination using optical data, and then concentrate on describing the automation of the image processing subsystem which forms the core of the complete autonomous navigation system. In addition, we will present some results of a set of tests used to exercise the subsystem.

2.0 FUNDAMENTALS OF OPTICAL ORBIT DETERMINATION

Historically, navigation of deep-space satellites has involved the use of radio data types for determining trajectory of the spacecraft and then predicting its course in the future. These radio data types include Doppler data, which measures the line-of-sight velocity of the spacecraft relative to the tracking station receiving the signals, and ranging, which measures the line-of-sight position². During approach to target bodies, optical data taken with an onboard camera also measures the target-relative position of the spacecraft. This methodology of combining both radio and optical data has worked very well in the past for flybys of both planetary bodies (the Voyager missions) and for asteroids (the Galileo mission)^{3,4}. For purposes of developing autonomously navigating spacecraft, however, the software needed to process radio navigation data is very complex. In particular, the high precision of Doppler data (typically 0.1 to 1.0 rends) requires very precise modelling of the dynamic forces acting on the spacecraft, as well as corrections to account for propagation of the radio waves through the atmosphere, and other error sources. In addition, a radio system onboard the spacecraft cannot really be considered autonomous because an uplink from the ground is required. By its nature, optical data is not as precise (an instantaneous position fix is accurate to about several hundred km at typical interplanetary distances of 10^8 km), but has the distinct advantage of being self-contained onboard the spacecraft. A spacecraft's camera takes the pictures, and all processing is done onboard, including data editing, centerfinding, and filtering. Also, because of its lower precision, the dynamic models do not have to be as precise. The question then remains as to whether optical data alone is sufficient to navigate the entire mission. Using the processing techniques described in this paper, various analyses have shown that for many mission types, the optical data is sufficient. The basics of orbit determination using optical data will now be described.

2.1 Optical triangulation in interplanetary cruise

The fundamental concept of optical orbit determination is extremely simple. The spacecraft's camera shutters an image of a solar system body (which will almost always be an asteroid for reasons described later) against a star background. Assuming that the heliocentric location of the asteroid (referred to as a "beacon") is known, the location of the asteroid in the camera field-of-view (FOV) determines a line-of-sight (LOS) direction to that asteroid, so the spacecraft's position must lie along that LOS. Two such LOS fixes taken instantaneously place the spacecraft position at the intersection of the two LOS vectors (Fig. 1). If

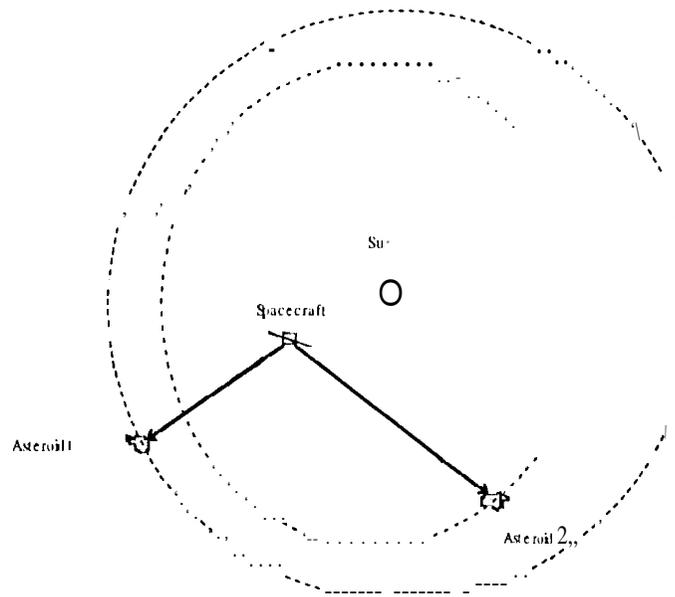


Fig. 1. Schematic of optical triangulation

two beacons are visible in the FOV, the spacecraft position in heliocentric space can be deterministically computed. In reality, however, the narrow camera FOV and the spacing of the asteroids precludes this occurrence except in rare instances. Thus, individual LOS images are taken over some interval of time, and the entire data is processed in a least-squares filter to determine the spacecraft's state (position and velocity). In addition, the LOS vectors need to be fairly widely spaced to provide optimal geometries for position estimation.

An implicit assumption in this process is that the inertial pointing direction of the camera boresight is known. This is determined by the stars in the FOV. Since the stars are effectively at an infinite distance away, their location can be thought of as fixed in the sky. Once the location of known stars are determined on the image, the pointing direction of the boresight can be computed using a least-squares process. Three parameters are estimated for the boresight direction - its right ascension, declination, and twist, which determines the angle the camera is rotated about the boresight. This implies that at least two known stars be visible in the FOV, but with a-priori information, even one star can fix the RA and DEC, with twist being relatively less well determined.

The accuracy of this procedure is dependent on several factors: the ability to determine the exact center of the stars and object in the image (a process termed "center finding"), the resolution of the camera, the distance from the spacecraft to the object, and knowledge of the object's heliocentric position. The first two of these issues will be addressed in a later section concerning image processing. Regarding the latter two, it can be seen that, with a given center-finding accuracy and camera resolution, the angular accuracy of the LOS fix is proportional to both the distance of the beacons from the spacecraft and to knowledge of the beacons' ephemeris (its position as a function of time). Thus, although the ephemerides of the planets are better known than that of asteroids, the proximity and number of asteroids (up to several thousand) makes them more viable candidates for use as beacons than planets, especially for missions such as DS-1 which spend all their time cruising in the inner solar system. The ephemerides of the larger and brighter of the asteroids are known to several tens of km; the smaller and dimmer ones are good to a few hundreds of km. This is accurate enough for most of the interplanetary cruise portions of the mission, but near encounter it is essential that many asteroids with accurate ephemerides are available. For DS-1 therefore, a ground-based campaign will be in effect to improve the ephemerides of the beacon asteroids used for the mission to accuracies in the lower tens of km.

2.2 Orbit determination filter

As mentioned above, the possibility of getting more than one object in the FOV simultaneously to obtain a deterministic fix on the position is unlikely. In any case, knowledge of the spacecraft's position alone is insufficient to determine the spacecraft's trajectory. Thus, some type of filter is needed to process the observations. The standard orbit determination procedure used is based on a linearization of the dynamical equations of motion. A nominal reference trajectory is generated from either a nominal mission plan or from a previous orbit determination solution by numerically integrating the equations of motion. Partial derivatives of the observables (in this case, pixel and line locations of the beacon in the image) with respect to the state (parameters to be estimated in the filter, which include position, velocity, and possibly other parameters which affect either the observables or the trajectory) are computed which depend on the nominal trajectory, and then accumulated into the information matrix. A residual vector is also generated which is the difference between the observed values of the asteroid pixel/line location and predicted values based on the nominal trajectory. After some time interval has passed, the accumulated information matrix is inverted and multiplied to the residual vector to obtain corrections to the state parameters. For the batch formulation of the filter just described, the state parameters are determined at some epoch time at the beginning of the data arc. This epoch state solution can then be mapped forward to any desired time using the equations of motion. The estimated, mapped trajectory then becomes the new nominal trajectory. For any given batch of data, the process is usually iterated several times to converge on the solution.

The dynamical equations which describe the spacecraft's motion include a term for the gravitational attraction of the central body (which is the sun for interplanetary cruise, or the planet or object being orbited for an orbiter) plus terms for various perturbing forces which affect the trajectory. Gravitational perturbing forces are introduced from other planets in the solar system, as well as unequal distributions of mass in the planet or asteroid for an orbiter. Non-gravitational perturbing forces include accelerations from solar radiation pressure, thruster firings from chemical engines on the spacecraft for either attitude maintenance or course corrections, and in the case for 1) S-1, the acceleration caused by the Solar Electric Propulsion system used for trajectory control. The equations of motion are numerically integrated using a Runge-Kutta integrator⁵.

3.0 IMAGE PROCESSING SUBSYSTEM

The heart of the autonomous optical navigation system is the image processing subsystem. This element reduces the raw image taken by the navigation camera to the fundamental observable needed for orbit determination -- the inertial LOS directions of the stars and beacon in the image. The procedure must be extremely robust in order to handle all types of errors which can corrupt the image, and perform the centerfinding tasks completely autonomously. Fortunately, a great amount of ground-based experience is available upon which to build and was used to full advantage in developing the system. Two technologies in particular which were originally developed for the Galileo mission have been adapted for use in autonomous image processing for IX-1. The first of these algorithms, termed the "autorover", provides the initial registration of the bright objects in the image and performs a pattern matching between these objects and an a priori pattern of the known stars and beacon to filter out unwanted bright spots. The output center locations from the autorover are crude, so the second algorithm, termed the Multiple Cross Correlator, or MCC, fine tunes the rough centers to determine the precise star-relative position of the beacon in the image. The output center location of the stars from this process is used to determine the inertial spacecraft pointing, and the center of the beacon becomes the fundamental observable for use in orbit determination. These aforementioned algorithms serve as the basis for optical navigation during cruise when the beacons are fairly distant and are therefore unresolved (i.e., their angular extent does not exceed a pixel in diameter). A third algorithm therefore performs centerfinding on an extended object (an object which exceeds a pixel in angular extent) for use during encounter or flyby of an asteroid.

3.1 The camera

Before going into the details of the image processing algorithms, a brief description of the camera system is needed. The camera to be used for navigation on IX-1 is a 670 mm focal length instrument with a 1024X 1024 pixel Charge-Coupled-Device (CCD). Each pixel is 9 microns square, resulting in an angular resolution of 13.4 prad/pixel, and a FOV of 0.788 deg. The full well of each pixel is approximately 80,000 electrons, with an electronic noise (i.e., exclusive of signal-induced shot noise) of about 40 electrons. The A/D conversion uses 12 bits. A dynamic range of at least 5 visual magnitudes is therefore possible, even if a signal-to-noise (SNR) ratio of about 10, required to achieve 0.1 pixel centroiding, is demanded.

The aperture of the camera is 100 mm. The exposures required to image navigation objects as faint as about 13th magnitude would be of the order of 1 second. However, the Spacecraft attitude rates of 4 or 5 pixels/see preclude such stable integrations, so another approach to producing images with sufficient SNR has been defined. The spacecraft attitude motion is confined to a deadband of about 40x40 pixels, and the shutter is left open for 100 seconds or more to allow the faint signal to build up a pattern of detected electrons on the CCD. Read noise is introduced only once at the end of each long exposure. The correlated image motions across the field form the basis of the centroiding techniques describe below.

3.2. Autorover

The Galileo spacecraft suffered a mission-threatening crisis with the failure of the high-gain antenna. A large effort was subsequently undertaken to rewrite the flight software to maximize the ability to compress as much data as possible into the severely limited downlink stream of the low-gain antenna. The nature of the Galileo mission is amenable to such a strategy in that short intense periods of science measurements are interspersed with month-long periods of cruise. However, the optical navigation frames, comprised of carefully pointed and timed images of the Galilean satellites and stars, are needed immediately for orbit estimation and control. In addition, the algorithms chosen for science data compression could seriously alter the appearance of the faint star images.

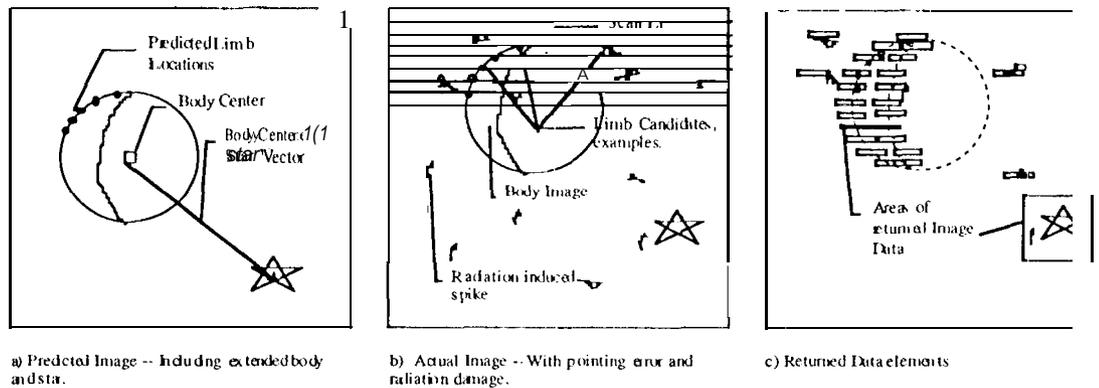


Fig. 2. Three stages of autonomous OPNAV data editing.

In response to this set of constraints, an autonomous optical navigation (OPNAV) editing capability was developed for use in the onboard computer. An algorithm was devised that could reliably and accurately locate the bright images of the Galilean satellites, and with that information, locate the star images in the field. Effective "compression" ratios of 100 to 400 could be achieved. Such compression factors would allow the down-link of the essential parts of an OPNAV image to occur in as little as ten minutes with no data loss or corruption in the returned regions. In contrast, several days would be required for the play-back of an entire image. In addition to locating the images of extended objects and stars, the algorithm had to successfully discriminate signal from radiation hits, which were expected to be common in the near-Jupiter environment.

The technique chosen was a pattern matching algorithm. The fundamental observable in the pattern match process is a threshold crossing. At some predetermined rate, the image is scanned (for the purpose of speeding the process, a predetermined number of lines is skipped after each scan), and areas of brightness are surveyed to see if they are above a minimum level and broader than a minimum width. Based on predictions of the appearance of the target satellite in the frame, a pattern of expected limbs, which meet the brightness and width criteria, are extracted in the form of image frame coordinates (column and row positions). These limb positions are compared with the growing set of threshold crossing positions (limb-candidates). These candidates (which may include both actual limb locations and radiation hits) are compared combinatorially to the predicted limb locations. The ensemble of limb candidates which best matches the predict set within predetermined error bounds defines the location of the object within the frame.

Fig. 2 shows the progression of the process from predicted limb locations, pattern matching, to extraction of image data. Fig. 3 shows an actual returned image from Galileo. The algorithm is able to successfully discriminate between radiation damage and actual limb locations while attempting to match the limb positions.

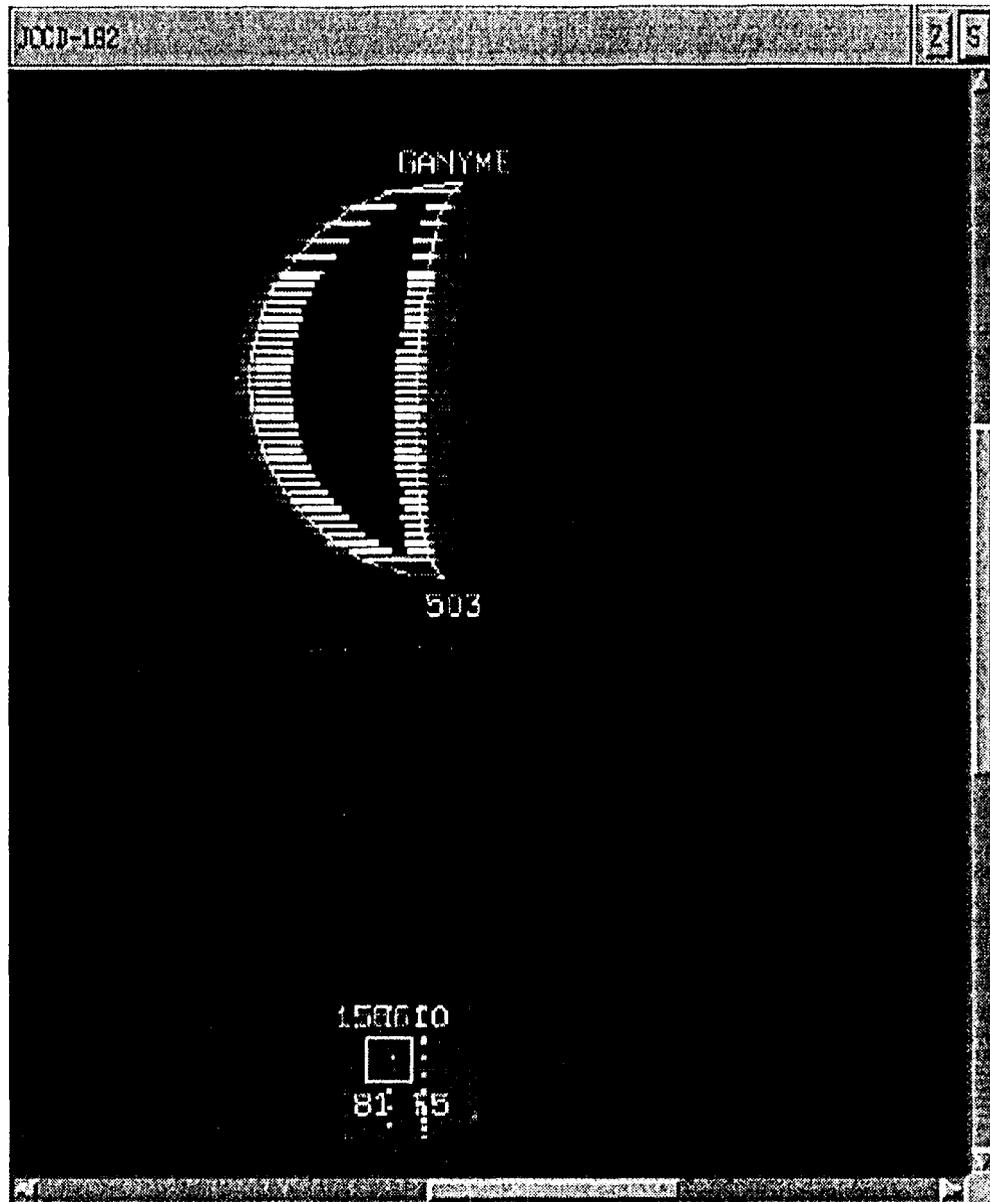


Fig. 3. Portion of OPNAV edited Galileo frame, with graphical identification overlays.

The mathematics of the algorithm proceeds schematically as follows. All possible unique differences between predicted limb locations are taken. For n predicted limb locations L_i , this forms an upper-triangular 3-dimensional matrix, with the upper-triangular portion two elements deep. Denote this matrix as:

$$PD(n \times n \times 2 \text{ upper triangular}) = \begin{bmatrix} L_i - L_j \\ \underline{1} \end{bmatrix} \text{ all } i, j \geq i \quad (1)$$

Similarly, for identified limb candidates, that is, those areas of sufficient brightness and width, a similar matrix can be computed; for m candidate limbs C_i :

$$CD(m \times m \times 2 \text{ upper triangular}) = \begin{bmatrix} C_1 & & \\ C_2 & & \\ \vdots & & \\ C_i & & \\ \vdots & & \\ C_j & & \\ \vdots & & \\ C_m & & \end{bmatrix} \text{ all } j \geq i \quad (2)$$

In general, $m > n$, due to having acquired as candidate limbs radiation hits or other false signals. It is also possible that not all of the limbs tallied in PD are represented in CD because of radiation damage or failure to meet the threshold test. In general, a successful "find" of the body is accomplished by matching a subset of a row of CD onto PD . The matching process then can be viewed as attempting to find the best match possible of the elements in a row of PD in the superset of elements of a row of CD . This matching is only approximate, with a match of row elements declared if they fall within a small position error, dx , of the predicted positions. The value of dx allows for small model errors and quantization effects due to the periodic scanning pattern. Some minimum number of element matches is required, and therefore a minimum number of elements in a row must match in order for a find to be declared. Because of the upper triangular nature of the difference matrices, if there are n limbs composing PD , and a minimum of k elements required for a find, then matches of PD against CD needn't proceed beyond row $n-k$ of PD .

For Galileo, typically 6 limb locations were used, with a minimum requirement of 4 to match for a successful find. For extended objects of about 1 (JO pixels and a scanning granularity of 2 or 3 lines), an error criteria dx of 2 or 3 was adequate. The onboard algorithm keeps a maximum of 16 candidates to reduce computation time. The results from the use of the auto-rover in the first encounter campaign have been excellent. Of the 33 scheduled OPNAV images, 6 were lost due to the expected obscuration of spacecraft booms, but none were lost due to the failure of the algorithm to locate the extended body, and that location was usually accurate to two or three pixels.

Few modifications were necessary to the auto-rover algorithm to make it suitable for use as part of the DS-1 AutoNav system. The principal difference in the systems is that for DS-1, the process of finding the limbs of an extended body was not the central problem. Instead, the registration of the frame[‡] is accomplished with a pattern of stars and target, all unresolved point-sources, but smeared into complex signal shapes (see Fig. 5). The threshold criteria is a minimum integrated brightness above background. An initial frame-wide search is performed to find such regions, or object candidates. This search is performed as a grid wise local integration of brightness, with two independent and offset grid patterns to catch those signals that might cross a grid boundary. As with the Galileo frames, there is a set of predict locations, and the combinatorial comparisons between difference matrices of these sets proceeds exactly as outlined above. In contrast to Galileo, several hundred pattern candidates are retainable, to be matched typically with a dozen or fewer predicted object locations. The pattern match error criteria of a half a grid-size width is used. The grid-size itself is a function of the expected smearing of the objects, and on DS-1 this is expected to be no greater than about 40 pixels. For the test with real data described below, the pattern smearing was somewhat greater. As an additional refinement, after the initial registration, local brightness centroids around the globally registered centers are determined, and the global-registration is tweaked to these average values. The net accuracy of the process is generally a small fraction of the initial searching grid size.

3.3 Multiple Cross Correlation (MCC)

Another Galileo derived navigation technology being adapted for New Millennium 1) S-1 Autonomous Navigation is the Multiple Cross Correlation process (MCC). The long exposures required to image faint asteroids and stars, and the relatively high spacecraft attitude rates combine to produce complex image patterns for the point source objects in each camera field. These image patterns are of course correlated. Furthermore, because the focal length is two orders of magnitude larger than the size of the CCD, the image smear can be assumed to consist of motions in the two directions normal to the LOS, with no effective "twisting" action around the LOS. Each image can thus be used as a template, or filter, with which to

[‡] registration is the process of placing the predicted object locations onto the observed images -- "Registering a predict OVERLAY", hence the name "ROVER"

locate every other. About each object-image center, determined by auto-rover, a region of image signal S_i is extracted to make a filter, F_i :

$$F_i = \frac{S_i - \frac{1}{m} \sum_{n=1}^m S_{i_n}}{\left(\sqrt{\sum_{n=1}^m S_{i_n}^2} \right)} \quad (3)$$

This represents a zero-mean normalized filter. The zero-mean attribute guarantees a zero response from the filter over a flat field, or a white-noise field. The normalization, when combined with a normalized data field, allows the filter to give a signal-level independent response during the subsequent convolution. Assuming the signal portions of the field, S_i , have been normalized, the convolutional product of filter i on data j , C_{ij} , a vector field whose locations correspond to locations of F on S , and whose values correspond to the convolutional product of same, is given by:

$$c_{ij} = F_i \otimes S_j = \sum_{k=1}^m \sum_{l=1}^n F_i^{kl} \cdot S_j^{kl} \quad (4)$$

For each element i and j of the cross-correlation, c_{ij} is evaluated for a maximum value and corresponding position, and parabolically interpolated to achieve sub-pixel accuracy in the center location. These centers form the data in a least squares estimate of position biases for each object. After several iterations of this solution process, the biases that are estimated represent shifts from the predicted positions to the best determined cross-correlated centers. The details of the solution process are given elsewhere⁶.

For Galileo, the accuracy of the process was approximately one-tenth of a pixel, or about a micro-radian. Several conditions made the Galileo images more difficult to process than the expected images from DS-1, making a 0.1 pixel (equivalent to 1 μ rad for DS-1) accuracy prediction for that system not unreasonable.

3.4 Extended body centroiding

The algorithms described above process images in which the beacon is unresolved. For flyby missions such as DS-1, however, the object's image will exceed a pixel in size somewhere between several days to several hours prior to encounter. At this stage, the cross-correlator will not work because the pattern formed by the extended object will be markedly different than that of the unresolved stars. Two separate approaches can be taken for these cases. The first assumes that an approximate value of the size of the object is available but the shape of the object is completely unknown, while the second assumes a known shape model. Regarding the first method, the assumption of known size is not unreasonable since ground-based photometry of the target object will exist which provides a rough estimate of its size. The onboard procedure then is to compute the center-of-brightness (COB) of the extended body in the image, and then apply a shift to correct for the offset between the COB and center-of-mass (COM). The computation of COB involves using a simple moment algorithm which weights the moment arm from a reference point (which, in this case is the upper left corner of the image) by the brightness value of the pixel. In order to minimize extraneous bright spots, the moments are computed only for points in the field within a 2 to 2.5 sigma uncertainty ellipse. This uncertainty ellipse is a projection of the combined dispersions of the spacecraft trajectory and object ephemeris onto the camera plane. A crude approximation to the COB to COM shift can then be computed as a function of the estimated radius of the body and its phase angle. Because this empirically determined formula assumes a spherical shape for the object, the further the object's true shape is from spherical, the larger the error in determining the shift. The formula, assuming uniform reflectance, is given as:

$$S = 3\pi \sin(\phi)(1 + \cos \phi) / \{16[\sin(\phi) + (\pi - \phi)\cos(\phi)]\} \quad (5)$$

where:

ϕ = the phase angle, and
S = the fractional shift value.

The accuracy of this procedure will vary widely depending on the actual shape of the object, and the errors incurred will have a systematic as well as a random component. For orbit determination analysis purposes, we have used an error value of 0.5 times the a-priori radius of the object.

A second approach can be used if a reasonable shape model of the object is available. This is an unlikely case for DS-1, however, so only a brief description of the algorithm will be given here; more detail on this topic was presented in a separate paper'. With the given shape model and a-priori information of the relative geometry between the spacecraft, sun, and target body, a predict of the scene in the camera is generated. Starting from the predicted COM location of the object in the computed image, the brightness is sampled along a radial scan towards the lit side until the limb is found. A sample of such scans are performed at some small angular increments to form an ensemble of limb locations. In addition, the COB of the image is also computed to obtain a predicted shift vector between the COM and COB. For the observed scene, the procedure is to first find the COB. The predicted COB to COM shift is then applied to the observed COB to get a rough approximation to the true COM. Radial scans are performed from this predicted COM location to once again obtain an ensemble of limb locations in this image. The observed limbs are then cross-correlated with the predicted set of limbs; the shift in the observations which produces the best match with the predicts becomes a new observable. The process is repeated for several sets of limbs spanning the sunward side of the object to produce a set of shift observables. Finally, the correction to the a-priori position of the spacecraft relative to the object is solved for in a least-squares sense, which minimizes the observables. These position observations can be incorporated into a complete orbit determination filter as described above. This procedure can produce centers with accuracies in the subpixel to several pixel range, depending on the accuracy of the shape model.

4.0 VALIDATION OF THE IMAGE PROCESSING SUBSYSTEM

The ability of the image processing subsystem to compute accurate centers autonomously under a wide range of conditions is critical to mission success since there are no other data types available to cross-check the results. For this reason, testing the algorithms using real data is an important step in validating the process. Ideally, images of asteroids taken from a spacecraft currently flying would be downlinked to the ground, processed through the autonomous navigation software, and then compared with results using standard radio navigation techniques. Unfortunately, no spacecraft currently flying can spare the resources to perform this test. The next best method therefore was to take test images from a ground-based telescope for processing. This section presents the results of preliminary ground validation of the image processing subsystem using observations of several asteroids over the course of three nights from a nearby observatory.

The equipment used to perform the validation was a 24 in telescope located at JPL's Table Mountain Observatory (TMO) facility. The telescope was equipped with a camera with a 512x512 CCD. The telescope focal length is 9503 mm. The FOV of the camera is about 1 mrad, with a per pixel FOV of 2.1 μ rad (note that this is considerably narrower than the camera which will be used for the DS-1 flight). Eight asteroids were observed over three nights. These asteroids were: 57 Mnemosyne, 61 Danac, 67 Asia, 73 Klytia, 114 Cassandra, 154 Bertha, 165 Loreley, and 168 Sibylla. The observations were taken during the nights of January 18, 20, and 21, 1996. Two types of observations were taken. The first was a standard exposure, with exposure lengths between 1 to 2 minutes, resulting in images of the stars and asteroid with a Gaussian smearing pattern (an example is shown in Fig. 4). In the second type of observation, called tailed observations, the images were smeared by physically altering the telescope pointing during the exposure. This was performed manually by tapping on the arrow keys which control the telescope pointing direction in a more or less random pattern. The resulting images were meant to mimic the attitude excursions on the spacecraft. Some sample observations of this type are shown in Fig. 5. The first type of observation is used for direct comparison of the centerfinding algorithms described in this paper with results using a standard center-finding technique employing a Gaussian pointspread function to model the star and

asteroid images⁸. Processing results from the trailed images were then compared to the results from the Gaussian spread images.

The output of the image processing are residuals obtained by subtracting the computed pixel and line locations of the asteroid and stars with its observed values. Since the location of the telescope and the coordinates of the stars are well known, and since the inertial pointing direction of the camera is computed using the star centroids, the residuals of the stars ideally should be very close to zero mean and be randomly distributed. The standard deviation of the star residuals then is a measure of the performance of the centroiding technique. The asteroid residuals, however, will be biased due to inaccuracies in its ephemeris, and the mean bias value of any given asteroid is a measure of its ephemeris error in the cross LOS direction. In reality, various errors such as those due to atmospheric distortions, star catalogue errors, and others will distort these results so they will not necessarily reflect the ideal conditions. However, since our purpose was not to obtain measurements of astrometric quality, no attempt was made to quantify or reduce these errors or sources. Instead, the residuals obtained from reducing the untrailed images employing standard centroiding techniques were used as a standard by which the trailed and untrailed residuals obtained from the MCC technique were compared.

The steps followed to obtain residuals using the MCC was as follows. First, predicted pixel and line locations of the asteroid and stars in the FOV are computed based on the a-priori pointing values. Using these predicts, the image is run through the autorover to locate the approximate centers of the desired objects, and to filter out signals which may be too weak or saturated. Then, a second filter is applied which deletes objects which are near the edges of the frame, and also deletes stars whose separation is smaller than the size of the MCC template. If at any time the object deleted is the asteroid itself, the entire image is removed from further processing. This methodology removed about 20% of the star observations.

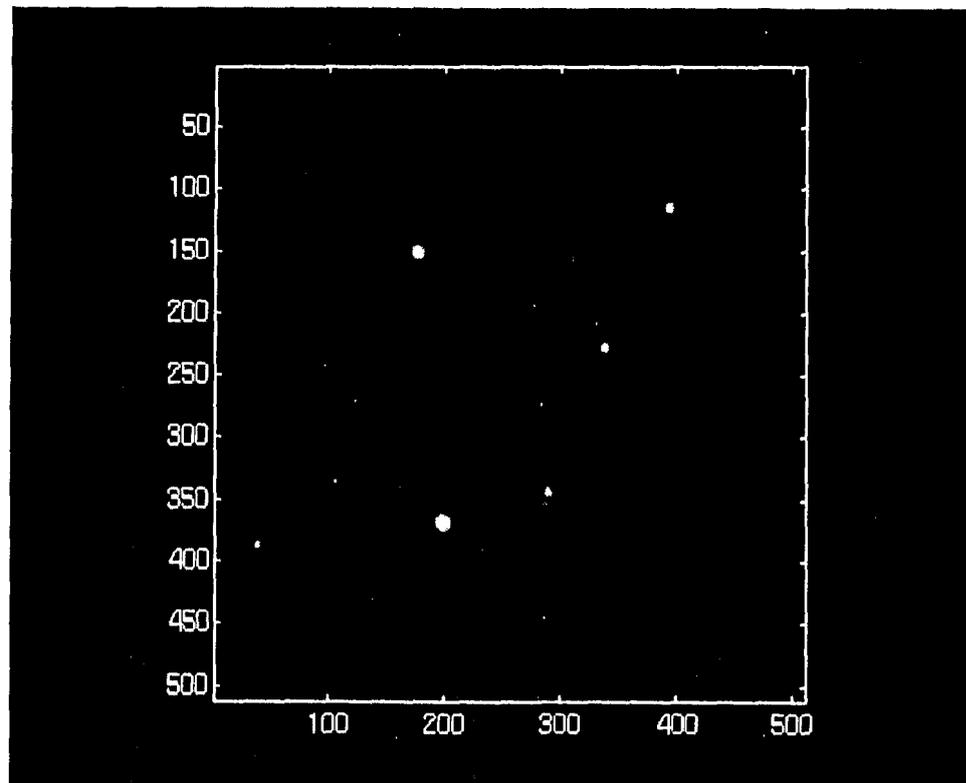


Fig. 4. An example of an untrailed image.

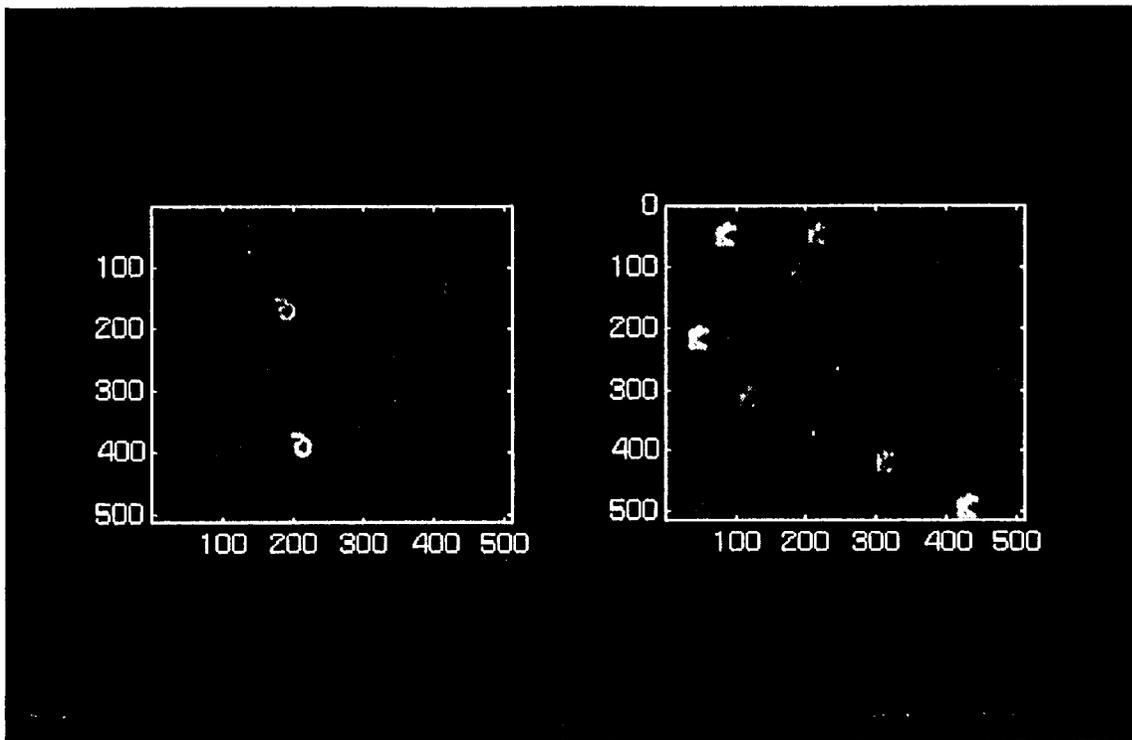


Fig. 5. Examples of trailed images.

After initial registration with the autorover and filtering out bad data points, a pointing solution is computed. This new pointing solution removes over 95% of the error in the initial pointing values and becomes the new nominal. Using the updated centers and pointing values, the image is sent to the MCC code. In the process of cross-correlating, the code also deletes objects which it cannot match. As a final filter, the values of the average response of the template with the data are checked, and any object which shows a low response is deleted. The precision centers output from the MCC method are then used to obtain the final pointing solution. Using this pointing, predicts for the pixel/line locations of the objects are recomputed and subtracted from the observed values to obtain residuals.

The results of this processing are shown in Tables 1-3 for each of the three nights (values for each night are shown separately because varying atmospheric conditions can alter the magnitude of the residuals). For each asteroid, the mean and standard deviation of the residuals for that night are computed. A comparison of the mean of the residuals on the untrailed images shows that the MCC technique matched the results of the standard method to better than 0.5 pixels, with the majority of values falling in the 0.1-0.2 pixel range. The only exception was for the asteroid Klytia on the first night which had a difference of nearly 0.6 pixels. Closer examination of these images revealed that one of the stars used for centroiding was either a binary star or had a dimmer, uncatalogued star very nearby. In any case, the result is that the MCC method had trouble obtaining a good cross correlation with this image, which biased the results. Manual removal of this star from processing improves the match to better than 0.1 pixels.

A comparison of the standard deviation of the untrailed image residuals show that the scatter using both methods are comparable. The noise level varies between about 0.04 to 0.3 pixels depending on various factors which affect the observations. The consistency of the values obtained verifies that the MCC technique is working properly and can obtain results which match those using conventional methods.

Table 1: Residuals for January 18, 1996

| | Standard Centroiding on Untrailed Images | | | MCC Centroiding on Untrailed Images | | | MCC Centroiding on Trailed Images | | |
|-----------|--|-------------------|--------------------|-------------------------------------|-------------------|--------------------|-----------------------------------|-------------------|--------------------|
| | # of obs | mean (pixel/line) | sigma (pixel/line) | # of obs | mean (pixel/line) | sigma (pixel/line) | # of obs | mean (pixel/line) | sigma (pixel/line) |
| Mnemosyne | - | | | | | | | | |
| Danae | 9 | 0.432 1.628 | 0.214 0.107 | 9 | 0.491 1.569 | 0.158 0.133 | 3 | 0.681 1.416 | 0.088 0.211 |
| Asia | | | | | | | | | |
| Klytia | 3 | -2.922 1.936 | 0.040 0.042 | 3 | -3.518 1.390 | 0.056 0.063 | 4 | -2.839 2.259 | 0.369 0.427 |
| Kassandra | - | | | | | | | | |
| Bertha | 3 | -2.440 -0.438 | 0.052 0.027 | 3 | -2.701 -0.407 | 0.040 0.036 | 4 | -2.777 -0.405 | 0.142 0.064 |
| Loreley | 4 | 6.483 -0.998 | 0.113 0.081 | 4 | 6.678 -0.593 | 0.154 0.042 | 4 | 5.759 -0.859 | 0.166 0.276 |
| Sibylla | 7 | -2.818 -0.367 | 0.047 0.086 | 7 | -2.848 -0.194 | 0.109 0.114 | 9 | -2.871 -0.142 | 0.229 0.164 |
| Stars | 92 | -0.059 0.061 | 0.787 0.611 | 92 | -0.084 -0.049 | 0.865 1.184 | 69 | -0.093 0.029 | 1.249 0.674 |

Table 2. Residuals for night of January 20, 1996

| | Standard Centroiding on Untrailed Images | | | MCC Centroiding on Untrailed Images | | | MCC Centroiding on Trailed Images | | |
|-----------|--|-------------------|--------------------|-------------------------------------|-------------------|--------------------|-----------------------------------|-------------------|--------------------|
| | # of obs | mean (pixel/line) | sigma (pixel/line) | # of obs | mean (pixel/line) | sigma (pixel/line) | # of obs | mean (pixel/line) | sigma (pixel/line) |
| Mnemosyne | - | | | | | | | | |
| Danae | | | | | | | | | |
| Asia | 5 | 7.583 -1.243 | 0.027 0.077 | 5 | 7.584 -1.293 | 0.047 0.316 | 4 | 7.552 -1.185 | 0.155 0.867 |
| Klytia | 4 | -3.589 3.034 | 0.118 0.034 | 4 | -3.518 3.247 | 0.115 0.076 | 4 | -3.431 3.339 | 0.235 0.056 |
| Kassandra | - | | | | | | | | |
| Bertha | 12 | -2.713 -0.377 | 0.239 0.094 | 12 | -2.812 -0.218 | 0.193 0.089 | 6 | -3.064 -0.196 | 0.441 0.273 |
| Loreley | | | | | | | | | |
| Sibylla | 3 | -3.329 0.331 | 0.043 0.022 | 3 | -3.210 -0.388 | 0.054 0.028 | 4 | -2.966 -0.334 | 0.109 0.093 |
| Stars | 88 | -0.021 -0.032 | 0.191 0.300 | 74 | -0.040 0.012 | 0.998 0.500 | 46 | 0.167 0.339 | 1.259 1.645 |

With this baseline established, the real test is in processing the realistic trailed observations. Examination of the residual means for these images shows that once again, they fall within 0.5 pixels of the standard results. The one exception was the pixel mean for Loreley on the first night, which had a 0.7 pixel difference. As of this writing, however, no explanation for this difference has been found. Regarding the residual standard deviations, the MCC technique applied to trailed observations has values which, with three exceptions, range from two to four times the values for the untrailed images. The result is not surprising and reflects the fact that the image wanders across the focal plane and therefore has less time to integrate on any one spot for a given exposure time. As a consequence, the signal is weaker and does not stand out as sharply over the noise which makes it more difficult for the cross-correlator.

The three exceptions noted were for Klytia on the first night, Asia on the second, and Danae on the third. Each had residual scatters which were an order of magnitude greater than those for the corresponding untrailed images and were therefore examined in more detail. For Klytia, the cause of the higher sigma was revealed to be a combination of a blemish of unknown origin which corrupted the image, and a very low

signal to noise ratio. The response from the correlator was just above the threshold to accept the observation, which implies that perhaps the threshold was set too low. For Asia and Danae, the problem was traced to the rather large signature of the trails which extended well beyond the limits of the templates. Increasing the size of the templates improves the results considerably but also slows the algorithm down. For actual flight, the template size will be set by the amount of deadband that the attitude control system can maintain, and the extent of the trails should not vary as much as the hand-generated trails used for this test.

Table 3. Residuals for night of January 21, 1996

| | Standard Centroiding on Untrailed Images | | | MCC Centroiding on Untrailed Images | | | MCC Centroiding on Trailed Images | | |
|-----------|--|-------------------|--------------------|-------------------------------------|-------------------|--------------------|-----------------------------------|-------------------|--------------------|
| | # of obs | mean (pixel/line) | sigma (pixel/line) | # of obs | mean (pixel/line) | sigma (pixel/line) | # of obs | mean (pixel/line) | sigma (pixel/line) |
| Mnemosyne | 2 | 4.745 -1.94(I) | 0.000 0.018 | 2 | 5.027 -1.880 | 0.054 0.060 | 1 | 5.144 -1.761 | |
| Danae | 5 | -0.472 2.070 | 0.039 0.027 | 5 | -0.480 2.008 | 0.045 0.068 | 4 | -0.695 1.694 | 0.275 0.320 |
| Asia | | | | | | | | | |
| Klytia | 4 | -3.952 3.590 | 0.019 0.011 | 4 | -3.915 3.572 | 0.029 0.036 | | | |
| Kassandra | 15 | -4.108 -1.226 | 0.129 0.357 | 15 | -4.174 -1.240 | 0.165 0.257 | 4 | -4.164 -0.942 | 0.266 0.256 |
| Bertha | 15 | -2.779 -0.586 | 0.087 0.090 | 15 | -2.898 -0.377 | 0.069 0.118 | 4 | -2.533 -0.896 | 0.505 0.144 |
| Loreley | | | | | | | | | |
| Sibylla | 3 | -2.541 -0.831 | 0.052 0.039 | 3 | -2.280 -0.593 | 0.042 0.148 | 4 | -2.543 -0.325 | 0.109 0.114 |
| Stars | 199 | .021 0030 | 0.445 0.384 | 180 | -0.089 0.083 | 0.500 0.440 | 67 | -0.110 0.385 | 1.375 1.436 |

5.0 CONCLUSIONS

The concept of building an entirely autonomous system to navigate spacecraft presents difficult challenges in algorithms and procedures. Successful results from earlier incarnations of the algorithms for the Galileo mission lend credibility to our thesis that it can be done. In addition, although not conclusive, the preliminary analysis of the test images from TMO adds further confidence that the system should perform in flight as expected. The checks already in place successfully weeded out most unprocessable data; the remaining discrepancies between the MCC and standard processing techniques have been explained, with some additional work necessary to ensure that these types of images are properly handled. The ultimate test, however, will be performed in the 1X- I flight to validate this technology for use in many future missions.

6.0 ACKNOWLEDGEMENTS

The authors would like to thank Bill Owen in the Optical Systems Analysis Group for his help in processing the TMO data. The work described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

7.0 REFERENCES

1. J. E. Riedel, S. Bhaskaran, S. P. Synnott, W. B. Bollman, and G. W. Null, "An autonomous optical navigation and control system for interplanetary exploration missions", IAA Paper IAA-I,-0506, IAA Conference, Laurel, Maryland, April 1996.

2. W. G. Melbourne, "Navigation between the planets", *Scientific American*, June, 1976.
3. J. E. Riedel, W. Owen, J. Stave, S. P. Synnott, and R. Vaughan, "Optical navigation during the Voyager Neptune encounter", Paper AIAA-90-2877, AIAA/AAS Astrodynamics Conference, Portland, Oregon, August, 1990.
4. P. G. Antreasian, F. T. Nicholson, P. H. Kallemeyn, S. Bhaskaran, R. J. Haw, and P. Halamek, "Galileo orbit determination for the Ida encounter", Paper AAS 94-132, AAS/AIAA Spaceflight Mechanics Meeting, Cocoa Beach, Florida, February, 1994.
5. W. Cheney and D. Kincaid, *Numerical Mathematics and Computing*, pp. 61-81, Brooks/Cole Publishing Co., Monterey, California, 1980.
6. R. M. Vaughan, J. E. Riedel, R. P. Davis, W. Owen, and S. P. Synnott, "Optical navigation for the Galileo Gaspra encounter", AIAA Paper 92-4522, AIAA/AAS Astrodynamics Conference, Hilton Head, South Carolina, August, 1992.
7. S. Bhaskaran, J. E. Riedel, and S. P. Synnott, "Demonstration of autonomous orbit determination around small bodies", Paper AAS 95-387, AAS/AIAA Astrodynamics Conference, Halifax, Nova Scotia, August, 1995.
8. S. P. Synnott, A. J. Donegan, J. E. Riedel, and J. A. Stuve, "Interplanetary optical navigation: Voyager Uranus encounter", Paper AIAA 86-2113, AIAA/AAS Astrodynamics Conference, Williamsburg, Virginia, August, 1986.