

OPTICAL NAVIGATION FOR THE STARDUST WILD 2 ENCOUNTER

Shyam Bhaskaran⁽¹⁾, Nick Mastrodemos⁽²⁾, Joseph E. Riedel⁽³⁾, Stephen P. Synnott⁽⁴⁾

⁽¹⁾Jet Propulsion Laboratory, California Institute of Technology MS 264-280, 4800 Oak Grove Dr., Pasadena, CA, USA, Shyam.Bhaskaran@jpl.nasa.gov:

⁽²⁾Jet Propulsion Laboratory, California Institute of Technology, MS 301-150, 4800 Oak Grove Dr., Pasadena, CA, USA, Nick.Mastrodemos@jpl.nasa.gov:

⁽³⁾Jet Propulsion Laboratory, California Institute of Technology, MS 301-150, 4800 Oak Grove Dr., Pasadena, CA, USA, Joseph.E.Riedel@jpl.nasa.gov:

⁽⁴⁾Jet Propulsion Laboratory, California Institute of Technology, MS 301-150, 4800 Oak Grove Dr., Pasadena, CA, USA, Stephen.P.Synnott@jpl.nasa.gov:

ABSTRACT

On January 2, 2004, the STARDUST spacecraft flew by the short period comet P/Wild 2 at a distance of 237 km. The primary goal of the flyby was to collect samples of the coma and return them to Earth on January, 2006. An additional goal was to shutter images of the nucleus during the flyby. In order to meet these goals, the spacecraft had to be guided to a flyby about 250 km from the center of the nucleus with an accuracy of better than 50 km. This was accomplished by the use of standard radio navigation techniques, augmented by optical images of the comet using the onboard camera. This paper describes the optical navigation techniques used on approach to target the desired encounter conditions. The optical navigation phase of the mission began about a month prior to encounter when the first images of the comet were seen. Because of the dimness of the comet, plus unexpected problems with the camera due to contamination of the optics, individual frames could not be properly processed to get good navigation data. Techniques such as multiple frame co-addition were employed to boost the signal. The data was used to design a series of maneuvers to guide the spacecraft to its flyby target. At about 7 days prior to encounter, efforts to clean the camera paid off, and better signal to noise ratios were achieved. The final images for targeting were taken roughly 14 hours prior to encounter. These images indicated the previous targeting maneuver done at 2 days prior to encounter achieved the desired accuracy. The final flyby reconstruction showed the accuracy to have been better than 10 km.

1. INTRODUCTION

On January 2, 2004, the STARDUST spacecraft flew by the short period comet P/Wild 2 at a distance of 237 km, and a velocity of 6.1 km/s, becoming the first spacecraft to capture samples of the dust environment around a comet. In addition, 72 images of the comet's nucleus were also taken, revealing surface details with resolutions less than 30 m, and providing substantial

information on the comet's size, shape, and surface morphology. The primary means of navigating the spacecraft to this flyby were standard radio navigation techniques, augmented by the use of an onboard navigation camera (the Navcam) to image the comet on approach. The latter was essential as the ephemeris of the comet was highly uncertain, and could only be improved to the desired accuracies using images taken by the Navcam. Despite this, the challenges of achieving this flyby were considerable, owing to the problems with the camera optics, the uncertain nature of determining the center of brightness of a dim, irregularly shaped object buried within a diffuse coma, and the geometry of the flyby which made it difficult to determine the downtrack position of the spacecraft relative to the comet. This paper describes the techniques used by the optical navigators to overcome these challenges (the radio navigation methods used on STARDUST are described elsewhere [1]).

2. MISSION OVERVIEW

STARDUST is the fourth of NASA's Discovery class mission – missions which are guided by the “better, faster, cheaper” philosophy and are designed to be Principal Investigator driven and cost constrained. The primary goal of the mission is to return samples of the comet dust to Earth, with a secondary goal to also capture high resolution images of the comet's nucleus during the flyby. The STARDUST spacecraft was launched on February 7, 1999 on a Delta 2 launch vehicle. During its nearly 5 year trajectory to get to comet Wild 2, it swung by the Earth once to achieve enough energy for reaching the comet, and also flew by the asteroid Annefrank. The latter encounter, which occurred on November 2, 2003, was performed as a dry run for the comet flyby, testing spacecraft activities and training the flight team for the comet encounter. The flyby was successful, and the distant images of Annefrank also provided a modest science return as a bonus. The spacecraft then performed a successful encounter with Wild 2 on January 2, 2004, capturing

dust samples in its aerogel collector grid, which was then stowed and sealed in the sample return capsule. STARDUST is now on its final 2 year orbit which will return to Earth on January 15, 2006. Several hours before entry, the sample capsule will be ejected from the main spacecraft bus and will parachute down to the Air Force Utah Test and Training Range, where the samples will be retrieved and sent to NASA's Johnson Space Center for analysis.

3. OPTICAL NAVIGATION

Standard navigation data types used on STARDUST included Doppler and range, which measure the line-of-sight velocity and position, respectively, of the spacecraft relative to the tracking station. These data are sufficient for guiding the spacecraft through all of the mission phases except the flyby of the comet. This is because the largest navigational error source for encounters with small bodies is the ephemeris of the body itself. Although ground-based observation campaigns are used to improve the ephemeris of the target body, the typical large distances involved when observing the bodies result in ephemeris estimates whose accuracies are on the order of hundreds to even thousands of km. The only way to improve on this is to use an onboard camera to determine the spacecraft's position relative to the target. The optical data is used in conjunction with the radio data types to achieve the desired flyby conditions, which for STARDUST, was at a comet distance of 250 km at the time of 19:20:00 on January 2, 2004.

The methodology used in optical navigation (opnav) is to take images of the target body against a star background. Because the stars are far enough away that their parallax is negligible, they can be used to determine the inertial pointing direction of the camera. An accurate algorithm is used to determine the location of the target's brightness centroid relative to the stars, a process known as centerfinding. This provides a measure of the angular displacement of the spacecraft relative to the comet. This has the effect of determining very accurately the spacecraft's flyby location in the plane perpendicular to the incoming asymptote (the "B-plane", described in the Appendix). However, relatively little information is gained in the downtrack, or time-of-flight direction; this must be determined using the spacecraft's radio determined position along with the a priori ephemeris of the comet.

To obtain an orbit determination solution, the radio data is combined with the optical data, and a least-squares fit is computed for the orbit. Typical parameters adjusted in the fit are the initial state (position and velocity) of the spacecraft, solar radiation pressure coefficients, scale factors on thrusting events, and time varying non-gravitational accelerations which soak up unmodelled error sources which affect the trajectory. When optical

data are used, the ephemeris of the target body is also adjusted. When the target body is a comet, optical data taken far from encounter (weeks to months) generally have the largest effect on the comet ephemeris; as the spacecraft gets closer, the optical data becomes strong enough relative to the radio data such that the spacecraft's trajectory is also adjusted by the optical data.

3.1 The Camera

The most important hardware component for optical navigation is the camera used for taking images. On STARDUST, the camera has a focal length of 201 mm and focuses light to a 1024x1024 pixel array Charge Coupled Device. Each pixel element's field-of-view (FOV) is 59 microradians, and the total FOV is 61 millirad, or about 3.5 deg. The location of a star or comet in the FOV is measured in pixels (horizontal distance from the left), and lines (vertical distance from the top). The hardware is capable of 12 bit digitization, resulting in Data Number (DN) greyscale values between 0 and 4095, with 0 being black and 4095 being white. The hardware also has an option for square root compression of the DN values to 8 bit which could be used when faster throughput or increased storage capability is necessary; this results in DN values between 0 and 255. The camera itself was fixed to the spacecraft, however, the light path also included a scan mirror which could rotate about a single axis over a range of about 200 deg. Thus, the camera boresight could cover a 3.5 deg by 200 deg range without having to reorient the spacecraft. When looking at the 0 to 19 deg range, the light path also included a periscope which is necessary to see around the front shields used to protect the spacecraft from comet dust during the flyby.

3.2 Optical Navigation Challenges

Opnav analysts on STARDUST faced several challenges, some due to the camera hardware and some due to the comet itself. On the hardware side, two specific problems posed difficulties. The first was that some type of contamination was coating the optics, with the probable location somewhere in the camera housing. Removing the contamination could only be accomplished by heating the camera by a combination of turning on small heaters located near the camera, and reorienting the spacecraft such that the camera radiator was in direct sunlight. However, after some period of time, the contaminants would return, requiring the heating process to be repeated. Thus, prior to important opnav events, the camera would need to be heated to dissipate the contamination. Depending on the level of contamination, the exposure durations to raise the signal to sufficient levels sometimes required long exposures, resulting in smearing of the images.

Another major problem with the camera setup was that many images were corrupted by stray light. The problem was very geometry dependent; the largest stray light was at mirror angles greater than 100 deg, but also at very low mirror angles when looking through the periscope. The problem was mitigated by reorienting the spacecraft such that the direction to observe would force the mirror to be in a benign location for stray light. The solution worked well, but at the cost of increased operational complexity to sequence the spacecraft attitude adjustments. Additional problems in the camera included random noise spikes which increased with increased camera operating temperatures, and cosmic ray streaks which could spoof the centerfinding algorithms if they were located on or near the star or comet.

Centerfinding for comets poses several challenges due to the fact that the comet has an unknown, irregular shape, illuminated from the side. For the Wild 2 approach, the phase angle (the angle between the sun and approaching spacecraft as viewed from the comet) was about 72 deg. Determining a geometric center of such an image is nearly impossible, and so the centerfinding will always be biased to some extent. Furthermore, the coma surrounding the nucleus adds an additional brightness source which is asymmetric, and must be accounted for in the centerfinding.

3.3 Image Processing

Image processing includes the algorithms and techniques needed to compute astrometric quality centers of the stars and target object in the camera FOV. The simplest algorithm, the brightness centroid, determined using a moment algorithm, is not effective due to the problems listed above. However, a modified moment algorithm worked very effectively for the particular case of centroiding the comet in the last 10 days or so prior to encounter. For this, only pixel locations above a certain threshold were allowed in the moment computation. The threshold was set to be 0.75 times the peak signal of the comet. This has the effect of removing much of the signal of the coma which might otherwise bias the moment. The centroid thus computed has a higher probability of being on the comet.

For stars, the above method is adequate, but better techniques are available. The one used here was to compute a least-squares fit of a Gaussian surface to the signal of the stars. Although the pointspread function of the star signals was not truly a Gaussian shape, the fit did quite well in determining the center location. This method is better at handling varying signal strengths and noise sources, but could not be used for the comet.

Although both the above methods worked well when near the comet, and when contamination was at a

minimum, they suffered early on when the comet was very dim and long exposures were needed. In this situation, the signals were smeared and noisy, making centroiding fairly difficult. These cases required a more sophisticated approach, called the multiple cross-correlation (MCC) method [2]. This technique, first developed for the Galileo mission, uses the pattern of an object in the FOV as a template to then cross-correlate to the other objects. The location of the peak correlation of the other objects defines a shift vector from its predicted location. This process is repeated for each of the objects, and the least-squares fit is then computed for the ensemble of shifts to find the ones which maximize the cross-correlations. This technique has proved to be more robust than simple moment algorithms in the presence of weak and corrupted signals to compute high precision object centers.

Two further aspects to image processing were the techniques used to boost the signal for dim images. The first was to co-add images together to improve signal-to-noise, especially for the comet. This proved absolutely necessary as the required brightness necessary to process single images did not occur until a week prior to encounter due to a combination of a relatively dim comet and re-contamination of the camera. Additionally, to improve the signal to noise, images often had to be median filtered using a 3x3 median filter box. Although median filtering suppresses the peak signal of an object, it was very effective in removing the high amplitude point-to-point random noise which affected many images during the approach to the comet. Fig. 1 shows an example of the raw comet signal taken on November 17, and Fig. 2 shows the same signal after median filtering. Note the central peak signal of the comet stands out much better in the filtered image, even though its absolute value is lower.




Fig. 1. Mesh plot of raw comet signal




Fig. 2. Median filtered mesh plot of comet signal

The data weight assigned to the optical data used in orbit determination solutions was highly dependent on the signal to noise. The image frames which had large spikes and low signals, or had to be co-added to detect the signal, were assigned weights of 2-3 pixels. For stronger signals when the spacecraft neared the comet, data weights were dropped to about 1 pixel or less.

4. COMET EPHEMERIS DEVELOPMENT

One important task performed concurrently with the navigation of the spacecraft was the computation of the ephemeris of Wild 2. The geometry of the Wild 2 ephemeris was such that it was in conjunction as viewed from the Earth from about May 2003 until a couple of weeks prior to the encounter. Furthermore, the comet passed through perihelion in October 2003, a period where outgassing activity, a major contributor to the uncertainty in the computation of its orbit, was at its peak. With the planned Earth observation campaign to determine Wild 2's orbit ending in May, the formal uncertainty in the comet's position at the time of encounter was nearly 3000 km (1 sigma). And, the orientation of the maximum dimension of the uncertainty ellipse was along the direction of the spacecraft's incoming asymptote, the direction least observed by the opnav data. Since the magnitude of the uncertainty needed to be improved in order for the onboard sequencing to properly image the comet, it was decided that additional ground telescope data was needed to try and improve the comet's orbit estimate. Several telescopes were commissioned to observe the comet in the last weeks of December 2003; these observations were combined with the spacecraft images to obtain a better estimate of Wild 2's orbit than either of the data individually. With the combined data, the ephemeris accuracy was computed to be about 300 km (1 sigma).

5. RESULTS

Optical navigation for the encounter effectively began on November 13, 2003 when the first images of the comet were taken. At this point, the spacecraft was about 26 million km from the comet. Three images were taken, each with an exposure duration of 3 seconds. Although no discernible signal from the comet could be seen in any individual image, co-addition of the three revealed a distinct signal above the noise. The offset of this signal from the predicted location of the comet was about 1 pixel, indicating an error in the comet's ephemeris of about 1500 km which was within the expected uncertainties of the ground-based ephemeris to date. Four days later, 20 images were taken, 15 at 5 second exposures and 5 at 15 second exposures. The two sets were co-added separately and clearly indicated the comet's location. Applying the MCC method, a center location of the comet was determined to an accuracy of about 0.25 pixels. With the precise astrometric measure of the comet's location, the discrepancy between the predicted and actual location of the comet grew to over 3000 km.

Two more image sets were obtained on November 20 and November 24, and once again processed using median filtering, co-addition, and the MCC. The data set was combined with the radiometric solution of the spacecraft's heliocentric trajectory to predict the encounter condition of the spacecraft with the comet in the B-plane targeting system. This indicated that, without any course corrections, the flyby would occur at a B•R value of -1037, a B•T value of 2964 km, and a flyby time of 19:02:24 on January 2 (the B-plane location is plotted in Fig. 3, labelled "Pre-TCM10"). The uncertainty ellipse for the flyby location was computed to have a semimajor axis of 453 km and a semiminor axis of 256 km, and the time uncertainty was 175 seconds (all 1 sigma). This information was used to design the first Trajectory Correction Maneuver (TCM 10), for December 3, 2003, to target the B-plane to a location 150 km from the nucleus, and the time to 19:20:00.

TCM 10 executed as it was designed. Unfortunately, two planned spacecraft activities occurring on December 4 and December 11 that would cause additional perturbations to the trajectory were not accounted for in the design of the maneuver, and therefore, the net effect was to have a spacecraft still not targeted to the proper flyby location. This would have to be corrected in the next TCM which was planned for December 23, ten days prior to encounter.

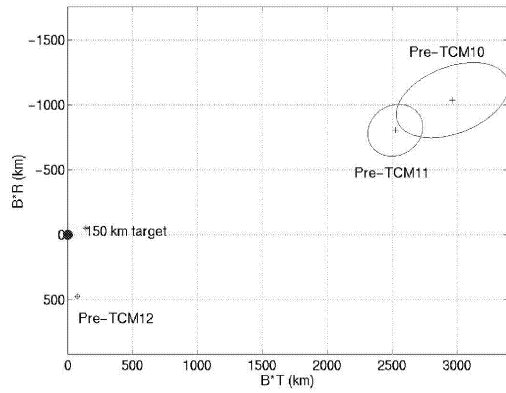


Fig. 3. Spacecraft flyby locations up to TCM 12

Following the TCM, image sets were taken on the 4th, 8th, and then at daily intervals starting on the 10th of December. These string of images indicated that the camera was clearly being re-contaminated; a measure of the peak signal of bright stars in the image showed that the signal had dropped in amplitude by a factor of almost 5. At this rate, the comet would become practically invisible until less than a day before encounter, much too late to plan targeting maneuvers. The problem was not unexpected, and resulted in the planning of another heating cycle to try and clear up the contamination for one final time before the encounter. The heating would take place during the TCM 11 maneuver, scheduled for December 23.

The location of the spacecraft in the B-plane determined by the final orbit determination solution prior to the design of TCM 11 was at a B•R value of -806 km, B•T of 2523.6 km (see Fig. 3, labelled “Pre-TCM11”), and a flyby time of 19:23:32. The uncertainty ellipse for the estimate had a semimajor and semiminor axis of 220 by 191 km, and the time uncertainty was 127 sec. The final opnav data taken for this solution was at a range of 10.8 million km. The shrinkage of the uncertainty ellipse at this point was not as good as expected because many of the images taken to this point were unusable for the reasons described above. Even the usable ones had to be deweighted in the orbit solutions to values of several pixels, as opposed to the predicted values of 0.5 to 1 pixel. On December 23, TCM 11 executed properly, retargeting the spacecraft to a location about 150 km from the nucleus. Although the planned flyby was at 250 km, the target for this maneuver was chosen such that future maneuvers would have the lowest probability of needing to rotate the spacecraft by 180 deg to perform it (all thrusters are located on one side of the spacecraft, and retargeting to further away from the comet does not need to have the thrusters reoriented).

On December 24, the first opnavs were taken following the heating. The images showed a substantial

improvement in throughput (a factor of 2), and combined with the nearer distance to Wild 2 (about 5.1 million km), resulted in the comet having peak signal values of about 25-30 DN above the background noise, without median filtering. Even though the camera state was not nearly as good as it could have been, for the first time, images would not have to be co-added to do good centerfinding. With the addition of several more days worth of data, it became apparent, however, that the previous solutions were biased by several hundred km in the vertical, or B•R direction. The cause of this is uncertain, but most likely due to the relatively poor conditions under which centerfinding had been done up to this point. By December 31, the solution had stabilized to a location nearly 500 km away from the designed target point. This error would have to be removed at TCM 12, planned on December 31 to retarget the spacecraft to a location 250 km from the nucleus. The solution used to design this maneuver had B-plane values of 474 and 76 km for B•R and B•T respectively, with uncertainty ellipse dimensions of 16 x 14 km (labelled “Pre-TCM12” in Fig. 3). The estimated time of encounter was at 19:24:05, with an uncertainty of 43 sec.

The goal of TCM 12 was to target a spot at the center of an annulus which defined a safe and acceptable flyby location (Fig. 4). Since the primary requirement for the mission was to collect a certain number of dust samples, the flyby could not occur too far out. With the latest dust models, this indicated that the furthest distance for optimal sampling was at a comet distance of 300 km, so this defined the outer ring of the annulus. The inner ring was at a distance of 200 km, chosen to minimize risk to the spacecraft from impacts from too large a particle which might compromise spacecraft safety. The upper and lower borders were chosen such that the spacecraft would not have to roll around the incoming asymptote direction by more than 13 deg; any larger would mean that the fixed high gain antenna used to communicate with the Earth would lose contact.

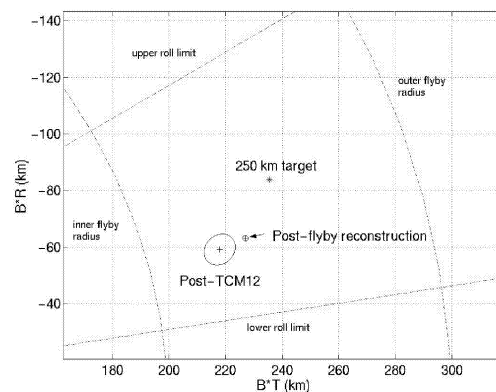


Fig. 4. Spacecraft flyby locations following TCM 12

Orbit solutions following TCM 12 indicated that, unlike at TCM 11, the accuracy of the pre-maneuver solutions was quite good and had resulted in TCM 12 placing the spacecraft well within the annulus. A planned final maneuver, TCM 13, was then cancelled, based on data taken 27 hours prior to encounter. The flyby location determined using this data had a B•R and B•T of -59 and 218 km, with an uncertainty ellipse of 6×5 km (Fig. 4). This solution was used to initialize the autonomous tracking algorithm which tracked the comet during the flyby [3]. The final opnav images were taken 14 hours prior to encounter at a distance of about $300,000$ km, and the solutions from this data showed the estimated flyby location holding steady. The post-encounter reconstruction of the trajectory using the 72 images taken in the minutes surrounding the encounter showed the actual location of the flyby to have been at a B•R and B•T of -63 and 227 km (Fig. 4), at a time of 19:22:36 on January 2, 2004. Fig. 5 shows the image shuttered at closest approach at a distance of 237 km.




Fig. 5. Closest approach image of Wild 2

6. CONCLUSIONS

The challenge of hitting a corridor roughly 100 km wide for a spacecraft over 2 AU away from the Earth was met using a combination of radiometric and optical data. Because the comet's ephemeris as determined from the ground had accuracies in the thousands of km, it was imperative that optical data be used to determine the spacecraft's position relative to the comet. Despite the complexities of determining centers of an irregularly shaped object surrounded by a hazy coma, the results showed that the methodology used to process the opnav images delivered the spacecraft to within 15 km of its planned target. This allowed the spacecraft to capture the required amount of dust, as well as provide the highest resolution images of a comet to date, an important science result in itself.

7. ACKNOWLEDGEMENTS

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

8. REFERENCES

1. Carranza, E., et al, "Orbit Determination of Stardust from the Annefrank Asteroid Fly-by through the Wild 2 Comet Encounter", Paper AAS 04-221, AAS/AIAA Spaceflight Mechanics Conference, Maui, Hawaii, February 8-12, 2004.
2. Vaughan, R., et al., "Optical navigation for the Galileo Gaspra encounter", AIAA Paper 92-4522, AIAA/AAS Astrodynamics Conference, Hilton Head, South Carolina, August, 1992.
3. Bhaskaran, et al, "Autonomous Target Tracking of Small Bodies During Flybys", Paper AAS 04-236, AAS/AIAA Spaceflight Mechanics Conference, Maui, Hawaii, February 8-12, 2004.

9. APPENDIX

Targeting at JPL is performed in the so-called B-plane coordinate system. The B-plane, shown in Fig. 6, is a plane passing through the center of the target body and perpendicular to the incoming asymptote, S, of the hyperbolic flyby trajectory. Coordinates in the plane are given in the R and T directions, with T being parallel to the Earth Mean Equatorial plane of 2000; to complete the right-hand coordinate system, R is positive downwards. The angle theta determines the rotation of the semi-major axis of the error ellipse in the B-plane relative to the T-axis and is measured positive right-handed about S. The horizontal coordinate in the B-plane is referred to as B•T and the vertical is B•R.




Fig. 6. The B-plane