NAVIGATING STARDUST-NEXT: THE ROAD TO TEMPEL 1

Aron Wolf, Paul Thompson, David C. Jefferson, Shadan Ardalan, Timothy McElrath, Matthew Abrahamson, Shyam Bhaskaran, C. Allen Halsell, Ramachand Bhat, Stephen Gillam, William Owen, J. Ed Riedel, Stephen Synnott, Tseng-Chan "Mike" Wang, Steven Chesley (Jet Propulsion Laboratory, California Institute of Technology), Kevin Gilliland, Greg McAllister (Lockheed Martin Astronautics)

The Stardust-NExT (New Exploration of Tempel) mission, a follow-on to the Stardust prime mission, successfully completed a flyby of comet Tempel-1 on 2/14/11. However there were many challenges along the way in navigating this mission to its successful conclusion, most significantly low propellant margin and detection of the comet in imagery later than anticipated. These challenges and their ramifications forced the navigation team and the project to respond with flexibility and ingenuity. As a result, the resulting flyby at an altitude of 178 km was nearly flawless, accomplishing all its science objectives.

Introduction

The Stardust prime mission collected a coma sample from comet Wild 2 during a flyby in 2004 and returned the sample to Earth in January, 2006 by jettisoning a sample return capsule which landed successfully on the Utah Test and Training Range. The spacecraft bus continued in a 1.5-year heliocentric orbit, with a planned return to Earth in January, 2009. Subsequently, it was found that the 2009 Earth flyby could be used to retarget the Stardust spacecraft to fly by comet Tempel 1 in February, 2011. This mission concept was proposed as Stardust-NExT (**N**ew **Ex**ploration of **T**empel) as a Discovery Mission of Opportunity and selected by NASA in July, 2007, by which time the S/C had been in space for over 8 years since its launch on Feb. 7, 1999.

The proposed baseline mission objectives were:
1. Document the style and amount of sublimational erosion and other surface changes occurring between successive perihelion passages of a comet.
2. Extend the geologic mapping of the nucleus of Tempel 1 to elucidate the extent and nature of layering and help constrain models of the formation and structure of comet nuclei .
3. Extend the study of smooth flow deposits, active areas, and known exposures of water ice.
4. If possible, determine the size and depth of the crater formed by Deep Impact (DI) and map any evidence of crater ejecta to provide constraints on models of crater formation and to derive further information on the structural properties of the nucleus of Tempel 1. (The DI impact produced so much ejecta that DI did not succeed in imaging the crater.)

The performance floor objective (required by NASA as part of the proposal) was to return at least one stereo image pair at a resolution of 20 m/pixel or better with a stereo separation angle between 10 and 30 deg., and to image at least 25% of the hemisphere seen by Deep Impact at 80 m / pixel or better.

These objectives were to be accomplished with high-resolution imaging using the Stardust navigation camera (NAVCAM); and measurements of the composition, size distribution, and flux of dust emitted in the coma using the Comet Interstellar Dust Analyzer (CIDA) and Dust Flux Monitor Instrument (DFMI).

Imaging of the DI crater, while not a stated criterion for mission success, was a goal of great interest. This required controlling the encounter time so as to arrive at encounter at the point in the comet's rotational period that would place the crater underneath the S/C with desired geometry and lighting conditions, which could require changing arrival time by up to ~20 hours (half of the comet's 40.7-hr rotational period). Analysis found that the most fuel-efficient point in the trajectory to adjust arrival time was roughly a year before encounter (requiring ~2.9 m/s DV per hour of arrival time change). A maneuver at this time was included in the mission plan as TCM-28. A 20-hour change in arrival time at one year before encounter required ~7.2 kg, which was over 40% of the estimated propellant quantity onboard at the time the Stardust-NExT mission was proposed. The estimated propellant quantity at that time was sufficient to accommodate a 20-hour change with some margin. Adjustments closer to encounter required significantly more fuel to accomplish (e.g. 15 m/s per hour at E-30d, increasing closer to encounter) and could not be accommodated within the propellant onboard. This required the important decision of how much to adjust arrival time to be made a year prior to encounter.

Stardust NAVCAM performance

During the Stardust prime mission, contamination of the camera was observed at launch and during the two subsequent 1-AU trajectory passes, as well as prior to the Annefrank encounter and on approach to Wild 2. The source of the contamination remains unknown. A procedure was developed which was successful in removing much of the contamination. This involved turning on the CCD heater to raise the CCD temperature to 0–10 deg. C, and then performing a "bake maneuver", slewing the spacecraft to place the Sun on the radiator on the –Z side of the spacecraft to raise the CCD temperature to 20– 25 deg. C. Camera bakes were planned periodically during the Stardust-NExT mission, including four in the last 60 days prior to encounter, to mitigate the risk of image degradation due to contamination.

Also observed during the prime mission was a fixed pattern of radiation-damaged high pixels that reduced the sensitivity of the NAVCAM. Power-cycling the camera greatly reduced the density and amplitude of the pattern noise. However, there was concern that the NAVCAM would fail to turn on after being turned off. Consequently, the Stardust-NExT project decided not to power-cycle it more often than absolutely necessary, allocating twelve power-cycles to support the most important OPNAV images (those immediately preceding maneuver data cutoffs and the image used to initialize Autonav prior to encounter).

Stardust Autonav system

Stardust has a unique camera/mirror system to track the comet nucleus through closest approach. This system contains a hard-mounted camera on the spacecraft bus, with the boresight bouncing off a mirror that rotates through 180 degrees so the comet can be imaged on approach, periapse, and departure. The predicted delivery uncertainty for ground-in-the-loop navigation was roughly 30 km in the crosstrack directions and about 90 seconds in downtrack (all 1 sigma). These accuracies are not sufficient to maintain visual lock on the comet through closest approach where the highest resolution is obtained. In particular, the crosstrack uncertainty is too large to determine the flyby plane in which the scan mirror needs to sweep, and the downtrack uncertainty is too large to determine the mirror angles vs. time through the encounter. For this reason, onboard closed-loop tracking is needed (as it was for the primary mission). This is provided by the autonomous navigation (AutoNav) system. Details of this system are provided in [1].

AutoNav was initiated at E-24 min and seeded with the best ground-based ephemeris knowledge. At 30 second intervals, images were taken and centroided, then used to perform a least-squares solution of the spacecraft state. The updated state was passed on to the Attitude Control System to compute the correct attitude for aligning the scan mirror plane and to the mirror controller to point the mirror. The first update occurred at E-10 min; subsequently updates occurred after every image. At E- 5 min, the spacecraft performed a roll maneuver to put the nucleus in the scan mirror plane. The encounter imaging sequence of 72 images was initiated at E -4 min, with images taken every 6 or every 8 seconds, with every second or every third image used by AutoNav. AutoNav terminated 90 seconds past the nominal encounter time.

Effects of unbalanced thrusters on navigation

Thrusters were mounted in four groups of 4 (two 1-N RCS and two 5-N TCM thrusters) on the Stardust S/C, all firing nearly in the same direction but canted to provide control authority about all three axes. This unbalanced thruster configuration produced a nonzero net DV whenever thrusters were fired. Thrusters provided the sole means of attitude control. Frequent pulsing was required to maintain pointing within a desired deadband, producing non-gravitational acceleration that over time had the largest effect on the trajectory other than gravity (Ref. 2). In addition to ACS deadbanding, propellant was expended for slews / walks to change attitude, and TCMs (using TCM thrusters). The unbalanced thruster configuration made it necessary to model DV due to both ACS deadbanding and planned attitude changes (slews or walks) to accurately propagate the trajectory and ensure successful targeting. Any changes to the mission plan that added or deleted attitude changes caused errors in the predicted trajectory. This changed the standard mission paradigm in which instrument calibrations and other spacecraft activities have no effect on navigation. In addition, unplanned events that caused thruster firings (e.g. safings) caused the trajectory to diverge from predictions.

The order of magnitude of DV due to attitude changes frequently was roughly the same as TCM DV. Consequently, the B-plane corrections at TCM's were dominated by non-random events; at every maneuver on Stardust-NExT, the design B-plane correction was greater than the 1-sigma relative orbit determination error. As a result, the TCM DV allocation was not well predicted by traditional tools using formal statistics. Another undesirable side-effect of the unbalanced thruster configuration was that cancellation of fuel-consuming activities to save propellant moved the predicted trajectory in the B-plane, which could increase the DV required at the upcoming maneuver and prevent realization of all the expected propellant savings. The impact of the above factors over the course of the Stardust-NExT mission was magnified because of the propellant situation, discussed below.

Changes to propellant budget

When the Stardust-NExT mission was proposed, the mission was planned using an estimate of remaining fuel onboard which was based on two methods: bookkeeping of thruster-on times since launch, and calculation from tank pressure, volume, and temperature (the "PVT" method). Later measurements using a third method based on thermal inertia of fuel in the tank yielded a substantially lower estimate of onboard propellant. This "Propellant Gauging System" (PGS) method relied on measuring the thermal response of the tank to heating and comparing the observed temperature rise to simulation results obtained from a thermal model of the tank (Ref. 3). Previous experience with Earth-orbiting spacecraft had shown that when propellant quantity is low, PGS is more accurate than PVT (because the sensitivity of the estimate to pressure changes decreases as propellant is depleted) or bookkeeping of thruster pulses (due to uncertainty in Isp and the amount of propellant expended per pulse). The accuracy of the PGS estimate is, however, heavily dependent on the fidelity of the thermal model of the tank.PGS testing was performed on three occasions: October, 2008; May, 2009; and November, 2010. The onboard propellant estimates from these tests were ~3kg lower than the estimates obtained from PVT and bookkeeping. (The accuracy of the PGS estimate was confirmed at the end of the mission on 3/24/11 when a decommissioning burn exhausted all remaining propellant at which time the PGS-based prediction was found to have been ~0.2 kg higher than the actual propellant quantity.)

Due mainly to this significant reduction in the estimated propellant quantity onboard, the maximum arrival time change at TCM-28 was reduced to 7.75 hr. by the time a decision on arrival time adjustment had to be made in January, 2010. At that time, the science team reported the results of an intensive, 2-year effort to predict the comet's rotation state at arrival and to recommend an arrival time adjustment if needed (Refs. 4, 5). This effort concluded that without an arrival time adjustment, the mission would not satisfy its performance floor objective of imaging 25% of terrain imaged by DI, and that a delay of 8 hr. would be necessary to meet the performance floor (and would also result in viewing the DI crater on approach, instead of at closest approach as originally planned). As a result, the project made the decision to delay arrival time by 8 hr. at TCM-28, with the knowledge that close monitoring of propellant consumption and predictions would be necessary for the remainder of the mission as a consequence (Ref. 6).

Subsequently, it was found that the propellant cost of continuous operation of the spacecraft in 0.25-deg. deadbanding (needed to support science imaging at 2-hr. intervals as planned for the last 60d before encounter) was higher than predicted. A preliminary estimate put the number of pulses per day during 0.25-deg. deadbanding at 400, for a total consumption of ~390 g at 18 mg/pulse for 65 days (E-60d to E+5d, allowing 5 days after encounter for transmission of science data). There was no extended period in 0.25-deg. operation to check this estimate until a NAVCAM bake and calibration in August, 2010 (several months after TCM-28 in February, 2010) at which time the S/C was operated for ~4 days in 0.25-deg. deadbands. This experience, acquired in an attitude different than the attitude the S/C would fly on approach, yielded ~630 pulses per day on average, an increase of ~50% over the preliminary estimate. This prompted an intensive effort over several weeks to reevaluate the estimated consumption in tight deadbanding. This effort resulted in an even higher estimate of 675 ± 75 pulses per day, totaling ~790 ± 87 g from E-60d to E+5d (Ref. 7). In response to this "hit" to the propellant budget, the project deleted a planned NAVCAM calibration and delayed the start of science imaging and associated 0.25-deg. deadbanding until E-40d to save propellant. In addition, TCM-30 was delayed from 10/13 until 11/20 to be able to incorporate the results of the above effort into the design of the maneuver.

The spacecraft experienced three safe mode entries in a period of roughly three weeks in late fall of 2010. A safe mode entry on 10/28/10 was caused by an unrequested reboot, believed to be a single event upset. Another entry on 11/11/10 was caused by a false IMU failure indication. The last of the trio of entries was on 12/10/10, due to a MEEB (memory error external bus) upset, which required a cold reboot of the S/C. This was accomplished on 1/4/10. These, like all safe mode entries, expended propellant and imparted DV that altered the trajectory (which had to be compensated for at the next maneuver).

The delay in detection of the comet (discussed below) also adversely impacted the propellant budget. A B-plane correction of up to several thousand km was anticipated to correct comet ephemeris error at the first maneuver after measurements from OPNAV images became available to augment radio data. The original plan assumed TCM-31 on 1/14 (E-32d) would be the first maneuver after OPNAV became available. In early January when the comet had not yet been detected, TCM-31 was delayed to 1/31 (E-14d) and TCM-32 was delayed from 2/4 (E-10d) to 2/7 (E-7d). Due to the decreased time to go to encounter, these delays necessitated an increase in propellant allocated for these maneuvers. To accommodate this increase, the start of science imaging and 0.25-deg. deadbanding was delayed further to E-7d.

Optical navigation and comet detection

An early science analysis predicted that the comet would be detectable as early as 60 days before encounter. Consequently, science imaging at 2-hr intervals was planned to begin at E-60d, requiring continuous operation in 0.25-deg. deadbands until encounter. The duration of science imaging was subsequently reduced from 60 to 40 days to save propellant as discussed above;

however twice weekly OPNAV imaging (which required operation at 0.25-deg. deadbands only for short periods of time) was begun at E-60d.

Initial OPNAV images did not reveal the comet, even with extensive image processing (co-adding and filtering of images, with software developed during operations). This prompted a reexamination of the early prediction of comet detection during which comparisons with the Wild 2 encounter scenario forced a revision of previous assumptions relating to comet brightness, resulting in a revised prediction of comet detection at ~E-20 days.

In addition to pattern noise and periodic requirement for camera bakes, two other issues impacted optical navigation making some images unusable: image smear and stray light contamination. Image smear of up to twenty-five pixels was observed due to S/C attitude motion within the 0.25-deg. deadband, with smear greatest near thruster firings. Attempts were made to recover useful images from smeared ones by characterizing the stellar point-spreads in each picture and sharpening them and the comet images with de-convolution techniques,. These attempts were unsuccessful due to high background noise and low comet signal. We relied instead on taking enough pictures so that we could reject those with smear greater than ten pixels. During the last week of approach the comet was bright enough to allow reducing OPNAV exposures to five seconds to minimize smear.

During the prime mission, it was observed that at some S/C attitudes and scan mirror angles, stray light scattered into the camera from undetermined spacecraft structures and produced increased background noise that varied in a complex way over timescales of several minutes. A calibration had been done at a representative attitude that showed no stray light at a mirror angle of 176 deg. However, a significant increase in background noise was noted in images taken between 1/8 and 1/17 at mirror angles > 168 deg. which made those images unusable. A test confirmed that background noise was at acceptable levels for a scan mirror angle or 160 deg.; subsequently the attitude profile was changed starting 1/18 (E-27d) to fix the mirror angle at 160 deg., which produced usable images again. Subsequently, a background estimation technique was developed in which the median value of each pixel across a set of 8 images shuttered close together in time, and that median subtracted from each of the images. This virtually eliminated the scattered light pattern, and also completely removed the pattern noise spikes.

Detection of the comet was reported on 1/20 (E-25d), after changing the attitude profile to fix the scan mirror angle at 160 deg. However, at E-20 days a combination of small comet prediction errors and ACS pointing errors caused the comet to be lost in a number of pictures. A temporary solution was planned and up-linked in one day.

Co-adding of images was required to detect the comet until E-7 days. The final week of imaging at 2-hr. intervals required the efforts of four opnav team members staffing 2-1/2 shifts per day. In total, 638 pictures were received (18 more were not down-linked) and 552 of them were used. The remaining 76 were unusable due to a variety of causes discussed above.

Final maneuver design decision

Images were scheduled at 2-hr intervals prior to E-106h, following which imaging was halted to accommodate a camera bake (the last one prior to encounter). Imaging resumed at E-82h and continued at 2-hr intervals until E-52h. Imaging was then halted for the execution of TCM-33 at E-48h, with one final image shuttered at E-42h. The data cutoff for TCM-33 was set at E-78h after a detailed re-examination of the entire Stardust-NExT maneuver design process in the months before encounter showed that the minimum time to design this maneuver was 30 hours.

As shown in Fig 1., optical navigation residuals over the last several days had developed a roughly sinusoidal shape, which was interpreted as evidence of a center-of-brightness / center-of-mass offset. However, the three image sets shuttered at E-82, E-80, and E-78h aroused a great deal of interest in the possibility that this pattern had been broken, possibly an indication of

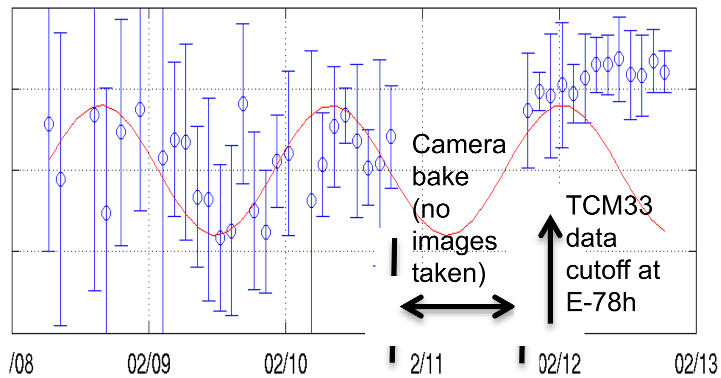

Fig. 1: Optical navigation residuals in R.A. over last several days before encounter

"seeing the nucleus" clearly through the coma. The decision meeting to select a TCM-33 design occurred at E-71h. Although the data cutoff had already passed and the last orbit determination solution before the decision meeting had already been designed, a special effort was made to process the next few imaging sets as quickly as possible to gather all available data before the decision meeting to see if the apparent trend continued. The next several points did indeed confirm the indication of "flattening out" of the previous sinusoidal trend, and after some discussion, the decision was made to design TCM-33 based on the assumption that the E-82, E-80, and E-78h data points were giving reliable information. Consequently, a new orbit determination solution was produced after the meeting, with these three points weighted heavily; this was used for the design of TCM-33.

Contingency maneuver cancellation

An opportunity for a contingency maneuver (TCM-34) was built into the schedule at E-18h, based on the limited amount of data available after the execution of TCM-33 including the image shuttered at E-42h. This contingency was to be used only to move the encounter aimpoint further from the comet. Since time available was insufficient to accommodate the full 30-hour maneuver design process, exercising this contingency would have amounted to selecting one of three maneuver designs that had been "pre-canned" (including testing of S/C sequences) weeks in advance for uplink to the S/C. These maneuvers were designed to move the aimpoint 35, 70, and 150 km radially away from the comet in the direction of the nominal aimpoint in the B-plane.

The B-plane "wedge plot" shown in Fig. 2 was devised to illustrate the decision criteria adopted by the project for exercising this contingency option. If the best orbit determination solution after TCM-33 indicated the flyby would occur in the green region in the B-plane, no contingency maneuver would be executed; in the yellow region, the decision would depend on the specifics of the situation; in the red region, one of the three contingency maneuvers would be chosen and executed to best satisfy the science requirements of the mission. As always, navigation recommendations were sent to Principal Investigator Dr. Joseph Veverka of Cornell University and Project Manager Tim Larson of JPL who would make the final decision. The lower altitude limit of the yellow region was driven by the 3.6 deg/sec angular rate limit of the scanning mirror.

At our flyby speed of 10.9 km/s, this limit was reached at an altitude of 174 km. The science team set the lower altitude limit at 155 km, accepting a low probability of image smear to get a high probability of images at or near the highest possible resolution. The upper altitude limit of the yellow region was set by the requirement to obtain one stereo image pair at 320 km altitude or less (20m/pixel or better resolution). The angular limits of the yellow region were driven by the desire to keep the solar phase angle at closest approach between 0 – 40 deg., with the nominal trajectory targeted to 20 deg. As Fig. 2 shows, the nominal aimpoint ("target" in the figure) coincidentally lay within ~0.1 deg. of the comet equator in the B-plane.

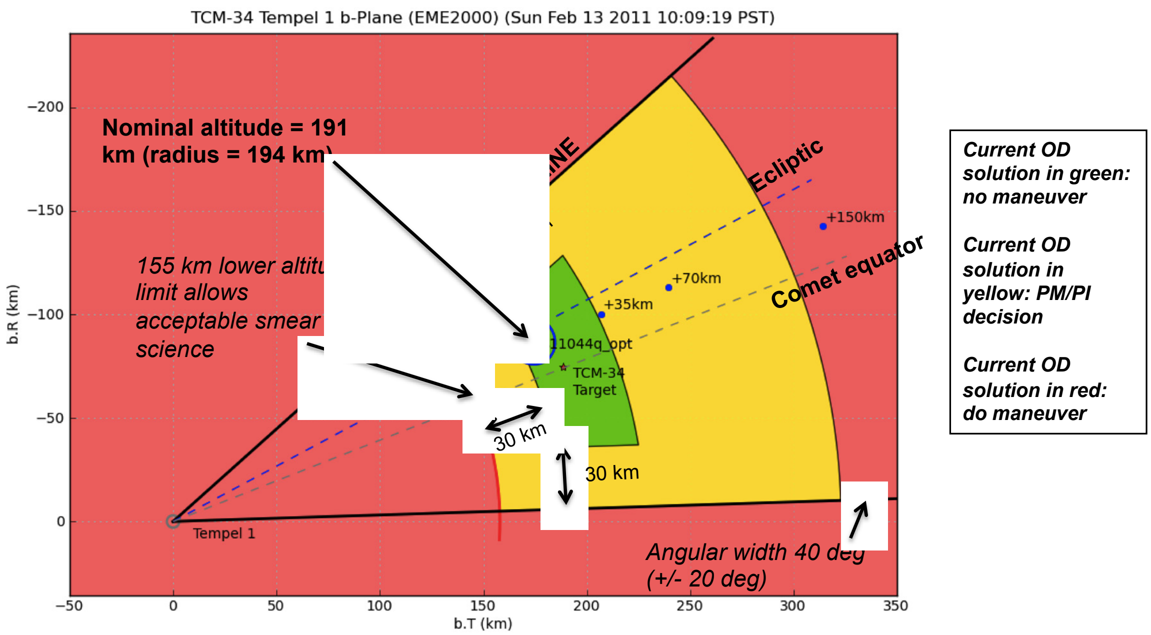

Fig. 2: Diagram ("wedge plot") showing decision criteria for TCM-34 contingency maneuver

Fig. 2 also shows the post-TCM-33 orbit determination solution (the point labeled s11044q_opt, surrounded by a nearly-circular blue ellipse which shows the 1-sigma B-plane uncertainty). Since this point was in the green region, no contingency maneuver was necessary. This confirmed that the decision to base the design of TCM-33 on the assessment that the sinusoidal pattern had disappeared in the E-82, E-80, and E-78h images was indeed the correct decision (with high-fives exchanged all around).

Autonav and operations during encounter

The Stardust autonomous navigation system successfully tracked the comet through closest approach, capturing the nucleus in the camera field-of-view in all 72 planned images. Post-encounter analysis showed that AutoNav performed as expected, with the final state correction amounting to about 13 km crosstrack, and 16 seconds in the encounter time. The crosstrack correction was almost entirely along the radial direction to the comet such that the mirror alignment attitude adjustment was less than 0.5 deg. Post-flyby reconstruction put the final estimate of the flyby altitude at 178 km (radius 181 km), yielding nearly the best imaging resolution possible without incurring image smear. The S/C team confirmed that telemetry showed that the scanning mirror was driven right up to its angular rate limit but did not exceed the limit.

Conclusions

The stunning success of Stardust-NExT provides an outstanding example of how much “bang for the buck” can be derived from extended missions using existing assets. However, the story of Stardust-NExT, of which an incomplete account is provided here, is also a cautionary tale providing a reminder (if any was needed) of the value of robust resources and margins, and above all an experienced team capable of responding with flexibility and ingenuity (as this team did with replans of activities at an accelerating pace all the way to encounter, responding to simultaneous problems such as delayed comet detection while operating with thin propellant margins, and implementing innovative image processing techniques on the fly), working together and backed by institutions with long experience in planetary missions.

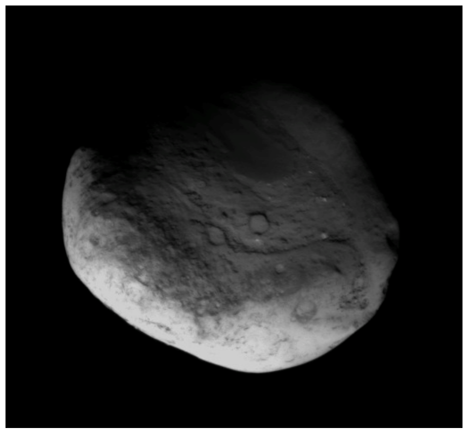

Fig.3: Stardust-NExT image of comet Tempel 1 near closest approach

Acknowledgements

The authors would like to acknowledge the support of Stardust-NExT project management, particularly Project Manager Tim Larson, Deputy PM Don Sweetnam, and Mission Manager Al Nakata, as the issues described in this paper arose. Also, the LMA spacecraft team, led by Allan Cheuvront, responded quickly and effectively to anomalies and worked closely with navigators. The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[1] S. Bhaskaran, J. E. Riedel, and S.P. Synnott, “Autonomous Target Tracking of Small Bodies During Flybys”, Paper AAS-04-236, AAS/AIAA Astrodynamics Specialist Conference, Maui, Hawaii, February 2004.

[2] T. McElrath, “Stardust Spacecraft Design Implications for Mission Operations Cost and Risk,” JPL Interoffice Memorandum 3430-11-v18, 16 May, 2011

[3] B. Yendler, “Results of PGS Fuel Estimation #3 for Stardust Spacecraft,” Lockheed Martin internal memorandum, December 20, 2010

[4] M. J. S. Belton, “STARDUST-NExT Encounter Time-of-Arrival Determination and Adjustment Strategy”, presentation to JPL management, January 20, 2010

[5] Belton, Michael J. S., Meech, Karen J., Chesley, Steven, et al, Deep Impact, Stardust-NExT and the accelerating spin of 9P/Tempel 1, *Icarus,* Volume 213, Issue 1, p. 345-368 (2011)

[6] Larson, T., “Time of Arrival Decision Briefing,” presentation to JPL management, January 20, 2010

[7] S. Ardalan, T. McElrath, P. Thompson, and B. Young, “Tight Deadbanding Predictions,”, presentation, November 3, 2010