Lessons Learned from OSIRIS-REx Autonomous Navigation Using Natural Feature Tracking

David A. Lorenz
SGT Incorporated
7701 Greenbelt Road
Greenbelt, MD 20770
301-286-6457
david.a.lorenz@nasa.gov

Ryan Olds
Lockheed Martin Space Systems
Company, P.O. Box 179
Denver, CO 80201
303-971-9694
ryan.olds@lmco.com

Alexander May
Lockheed Martin Space Systems
Company, P.O. Box 179
Denver, CO 80201
303-977-6620
Alexander.j.may@lmco.com

Courtney Mario
Draper
555 Technology Square
Cambridge MA 02139
617-258-2384
cmario@draper.com

Mark E. Perry
The Johns Hopkins University
Applied Physics Laboratory
11100 Johns Hopkins Road
Laurel, MD 20723-6099
240-228-0699
Mark.perry@jhuapl.edu

Eric E. Palmer
Planetary Science Institute
1700 E Fort Lowell, Suite 106
Tucson, AZ 85719
520-301-4517
epalmer@psi.edu

Michael Daly
York University
4700 Keele Street, Toronto
Ontario, Canada, M3J 1P3
416 736 2100 x22066
dalym@yorku.ca

Abstract—The Origins, Spectral Interpretation, Resource Identification, Security-Regolith Explorer (OSIRIS-REx) spacecraft launched on September 8, 2016 to embark on an asteroid sample return mission. It is expected to rendezvous with the asteroid, Bennu, navigate to the surface, collect a sample (July '20), and return the sample to Earth (September '23). The original mission design called for using one of two Flash Lidar units to provide autonomous navigation to the surface. Following Preliminary design and initial development of the Lidars, reliability issues with the hardware and test program prompted the project to begin development of an alternative navigation technique to be used as a backup to the Lidar. At the critical design review, Natural Feature Tracking (NFT) was added to the mission. NFT is an onboard optical navigation system that compares observed images to a set of asteroid terrain models which are rendered in real-time from a catalog stored in memory on the flight computer. Onboard knowledge of the spacecraft state is then updated by a Kalman filter using the measured residuals between the rendered reference images and the actual observed images. The asteroid terrain models used by NFT are built from a shape model generated from observations collected during earlier phases of the mission and include both terrain shape and albedo information about the asteroid surface. As a result, the success of NFT is dependent on selecting a set of topographic features that can be both identified during descent as well as reliably rendered using the shape model data available. During development, the OSIRIS-REx team faced significant challenges in developing a process conducive to robust operation. This was especially true for terrain models to be used as the spacecraft gets close to the asteroid and higher fidelity models are required for reliable image correlation. This paper will present some of the challenges and lessons learned from the development of the NFT system which includes not just the flight hardware and

software but the development of the terrain models used to generate the onboard rendered images.

TABLE OF CONTENTS

1. INTRODUCTION

The Origins, Spectral Interpretations, Resource Identification, Security, Regolith Explorer (OSIRIS-REx) mission will characterize and map the surface of asteroid Bennu (Figure 1) and return sample to Earth. Bennu is both the most accessible carbonaceous asteroid and one of the most potentially Earth-hazardous asteroids known. Bennu is a B-type asteroid and represents an important source of volatiles and organic matter to Earth as well as being a direct remnant of the original building blocks of the terrestrial planets [1]. Knowledge of the nature of asteroids like Bennu is fundamental to understanding planet formation and the origin of life. The return to Earth of pristine samples with known geologic context enables precise analyses that cannot

be duplicated by spacecraft-based instruments, revolutionizing our understanding of the early Solar System. Bennu is one of the most well characterized asteroids known using ground and space based telescopes to view Bennu across the entire light spectrum. There is strong evidence to
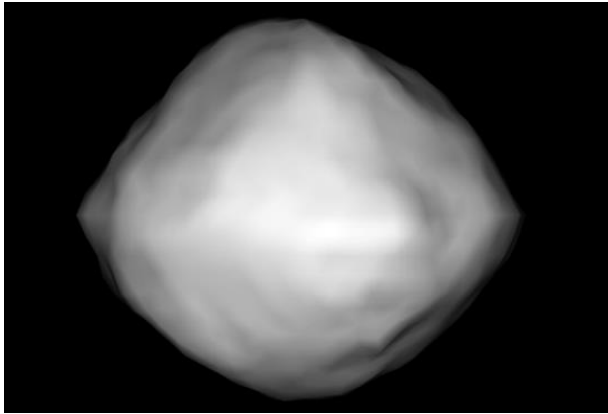

Figure 1 - Near-Earth Asteroid Bennu (shape derived from ground-based planetary radar data) [2]

support the presence of regolith (loose material) available for sampling. The study of Bennu will assist NASA and the scientific community in understanding the origin of the Solar System and life itself. It will also serve to better understand the hazards and resources needed to for future missions small near-Earth celestial bodies. [1]

Mission Overview

OSIRIS-REx was launched on September 8, 2016. It will rendezvous with Asteroid Bennu in November of 2018 after performing a series of orbit maneuvers and Earth Gravity Assist in September, 2017. Once it reaches Bennu, OSIRIS-REx will spend a year photographing and studying the asteroid so as to fully characterize its surface and its gravitational field. These operations will allow the OSIRIS-REx team to select a candidate sample site that will both ensure the safety of the spacecraft while also providing the most likely place to obtain a scientifically valuable sample. Once the sample site is selected, the spacecraft will enter the Touch And Go (TAG) phase of the mission that will culminate in the spacecraft gently making contact with the asteroid and collecting a sample of asteroid regolith (minimal requirement 60g). Once the sample is collected, the quantity of sample will be estimated and verified. The sampling head along with the sample will be placed in the Sample Return Capsule (SRC) for its return journey to Earth, arriving in September, 2023. As the spacecraft approaches Earth, the SRC will be jettisoned from the spacecraft and re-enter the Earth atmosphere landing at the Utah Test and Training Range. It will then be transported to Johnson Space Center, where the collected regolith samples will be removed and delivered to the OSIRIS-REx curation facility. The rest of the spacecraft will continue in its orbit around the Sun. [2]

The OSIRIS-REx flight system is made up of the spacecraft bus (which includes the structure and all of the various subsystem components to control and operate the vehicle), the TAG Sample Acquisition Mechanism (TAGSAM), the SRC, two Guidance, Navigation & Control (GN&C) Lidars, two TAG Cameras (TAGCAMS) and the five science instruments (OSIRIS-REx camera suite (OCAMS), OSIRIS-REx Visible and Infrared Spectrometer, OSIRIS-REx Thermal Emission Spectrometerj, OSIRIS-REx Laser Altimeter (OLA), and Regolith X-ray Imaging Spectrometer) [3].

Touch and GO Description

The TAGSAM is the key flight system component used for making contact and acquiring sample from the surface of Bennu during the TAG mission phase. TAGSAM is designed to collect greater than 150g to provide margin to the 60g mission requirement. Prelaunch tests using a variety of soil samples indicate that the TAGSAM could acquire 1000g or more. The TAGSAM functions by fluidizing regolith with a high pressure gaseous nitrogen flow to transport it to a sample container located in the TAGSAM 'head'. The TAGSAM is made up of a single planar, articulating arm with redundant motor windings at the shoulder, elbow, and wrist, providing large structural, torque, and alignment margins, ultimately ensuring successful sample acquisition and stowage of the

Figure 2 - Single Plane of Motion TAGSAM with potentiometers for simple & reliable positioning

TAGSAM head into the SRC (Figure 2). The section of the TAGSAM arm between the elbow and the wrist is a spring loaded telescoping arm (Pogo) that will compress as the spacecraft contacts the surface. This will limit the forces experienced by the spacecraft during initial contact and also assist in rebounding the spacecraft away from the asteroid after the sample is collected.

OSIRIS-REx utilizes a TAG approach for sample collection to minimize spacecraft complexity and reduce risk of being in contact with the surface for extended periods of time. The TAG trajectory profile (Figure 3) uses a series of three maneuvers to take the spacecraft from a circular orbit at 1km radius down to the surface in about four hours.
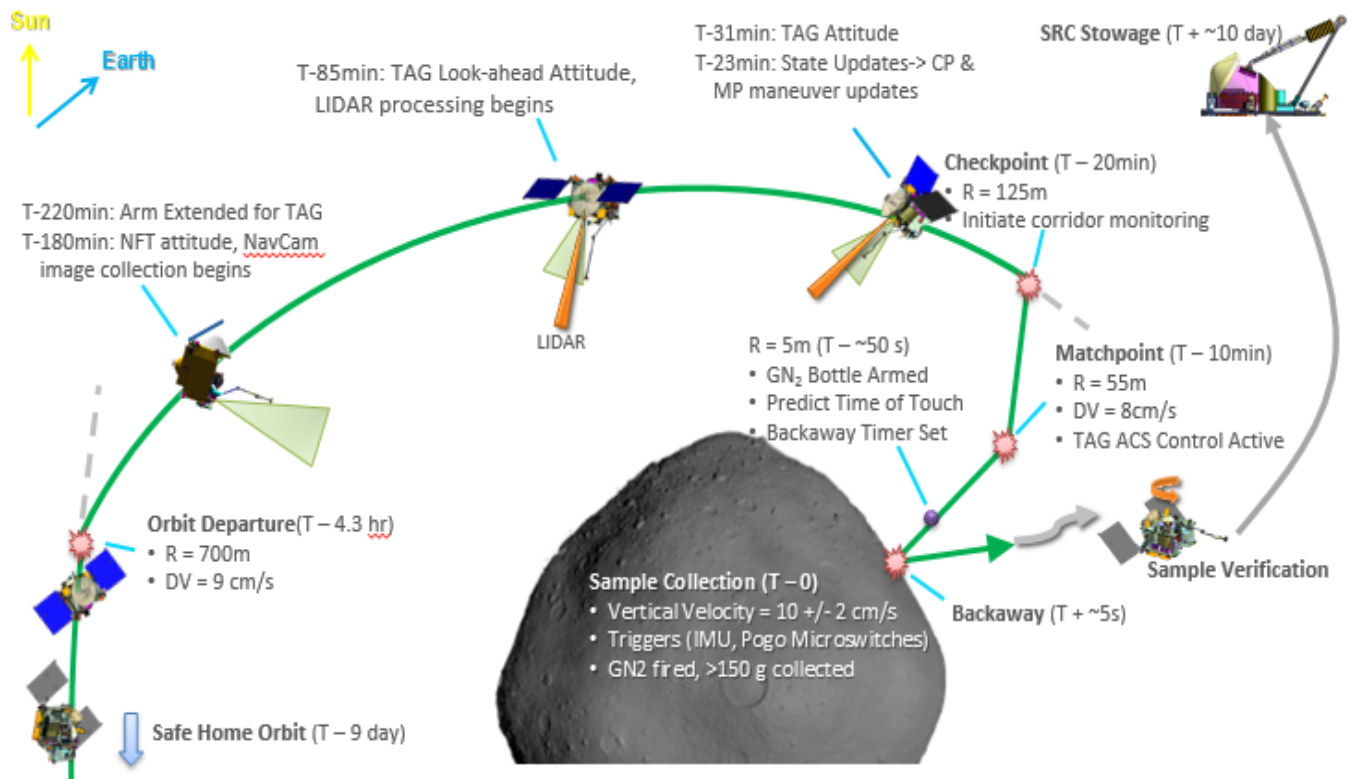
Figure 3 - TAG Phase Event Overview

The orbit departure maneuver (ODM), about 4.3 hours prior to sampling, puts the spacecraft on a trajectory that has a perigee altitude of 125m. Once it reaches this Checkpoint (CP) altitude, it performs a maneuver in order to set it on a course to intersect the asteroid. Approximately 10 minutes following the CP maneuver, the spacecraft reaches an altitude of about 55m where it performs a third maneuver, known as Matchpoint (MP), that change adjusts the horizontal velocity so as to match the horizontal velocity of the asteroid. It also adjusts the vertical velocity so that the spacecraft will contact the surface at 10 cm/s.

After MP, the spacecraft is in free fall to the surface with the attitude control system engaged to hold the desired TAG attitude. Because of the microgravity environment, over the 10 minutes free fall the velocity only grows by a few cm/s to the desired contact velocity of 10 cm/s. At contact, the TAGSAM arm compresses, the TAGSAM sampler head performs regolith sample collection and then the arm helps rebound the vehicle away from the surface. The spacecraft then fires thrusters to maneuver away from the surface and safely escape Bennu's gravity.

OnBoard Navigation System Overview

Because the round trip light time will be on the order of 30 minutes, the spacecraft is required to be fully autonomous as it navigates from ODM through TAG, collects a sample, and then performs a backaway maneuver that will send the spacecraft on a trajectory safely away from the asteroid. To accommodate this, the onboard navigation system performs a

number of functions following the ODM. First, it provides an update to the spacecraft state just prior to the CP maneuver. This update will be used to adjust both the CP and MP maneuvers to account for variances in the ODM performance and unknown forces such as solar pressure and higher order gravitational terms only experienced at low altitudes. Secondly, the navigation system monitors the spacecraft state as it descends toward the asteroid comparing the measured state to the expected state. If the deviation is too great, the onboard system will issue an abort and the spacecraft will perform a backaway maneuver to get it a safe distance from the asteroid. Thirdly, the navigation system will determine when the spacecraft crosses several key altitudes where the spacecraft configuration will be changed. The first altitude crossing is at 25m where some of the spacecraft fault protection is disabled and preparations are made for asteroid contact. The second altitude is at 5m where the navigation system calculates the expected time-of-touch (TOT). If the spacecraft detects contact prior to the opening of the TOT window or does not make contact prior to the closing of the window, it will perform an abort.

The driving requirements for TAG navigation are to contact the surface within a 25 m radius of the desired sample collection site with velocity errors less than 2 cm/s. To meet these requirements it was proven early on in the mission design process that the CP and MP maneuvers would need to be adjusted to clean up the dispersions after ODM due to navigation and burn uncertainties. The baseline navigation method was to use Lidar range measurements to determine

the spacecraft orbit state prior to CP in order to use a guidance algorithm to update both the CP and MP Burns [4]. The Lidar baseline had risk due to hardware development challenges and there was serious concerns that the Lidars may not be available in time. The project team decided a backup to the Lidar navigation was prudent to mitigate this risk. The decision was made to develop an autonomous optical navigation system called Natural Feature Tracking (NFT) as this backup using input from the existing GN&C TAGCAMS.

Natural Feature Tracking

Multiple backup techniques were considered to estimate the spacecraft state prior to CP in order to implement the necessary onboard guidance for meeting the TAG accuracy requirements. Of the many options, it was clear that optical navigation was likely the most promising and mature technique to implement. Prior to the Preliminary Design Review (PDR), the project team decided that navigation cameras would be added to the hardware baseline to permit onboard optical navigation software capability at a later date until further analysis could be done to fully define what the backup would be. From Mission PDR to Critical Design Review (CDR) the backup trade continued and converged on utilizing NFT from Lockheed Martin's (LM) previous experiences with optical-based navigation. NFT was shown to be able to meet the driving TAG accuracy requirements as well as perform navigation beyond CP to enable safety monitoring of position and rate during final descent.

The NFT system is an autonomous optical navigation software suite. This system was developed by LM Space Systems Company and was qualified for flight following its addition to the program after completing verification, ground testing and being demonstrated through flight-like scenarios. The NFT software has been designed to perform autonomous high-precision orbit determination by tracking the location of known features on the surface of small bodies such as Bennu. This capability can be tailored for any number of specific applications and underwent some modification for the OSIRIS-REx mission. In the case of the OSIRIS-REx mission, the NFT system is used to update and predict the orbital state of the spacecraft at the time of the CP burn and also at the time of asteroid sample collection.

The NFT system estimates the orbital state by matching "features" found in imagery collected by the Navigation Camera (referred to as NavCam). This imagery is matched against the predicted appearance of the feature, which is rendered on-board as depicted in Figure 4. NFT predicts which features it expects to see in the image, renders the

expected appearance of these features and finally matches these predictions to the actual image data. The locations of these matches are used to update the on-board knowledge of where the NavCam is and how it is oriented.

For the onboard rendering, NFT renders each feature using shape model data generated from data collected earlier in the mission. Each feature is represented by an array of terrain
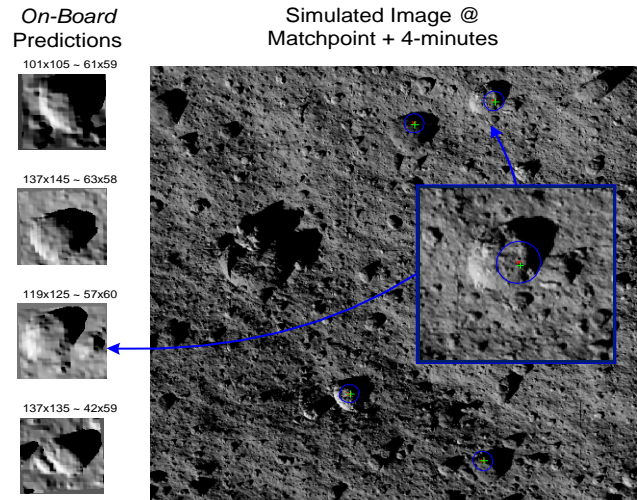

Figure 4 - NFT Feature Matching Depiction

information and associated albedo data that corresponds to a patch on the asteroid surface. NFT then uses this data along with the predicted sun angle and camera pose to render the expected appearance of the feature, which is matched against the onboard collected image using normalized cross correlation.

The quality of a feature match can be quantified by a correlation score and the location of the corresponding correlation peak in the camera focal plane. These two metrics quantify how strongly the NFT rendered image matches the actual imagery collected by the NavCam as well as how far from the predicted location the correlation peak was found (see Figure 5). These two metrics of correlation score and pixel error of the associated correlation peak location will be referred to again as different tests involving correlation performance are described. The NFT software has been matured through TRL 8 and will reach TRL-9 after being used on-board OSIRIS-REx in the 2018/2019 timeframe at asteroid Bennu. A checkout will first be performed in a 1 km orbit to test the system performance and feature prediction accuracy. Following this checkout, NFT will be ready to support the OSIRIS-REx TAG event.
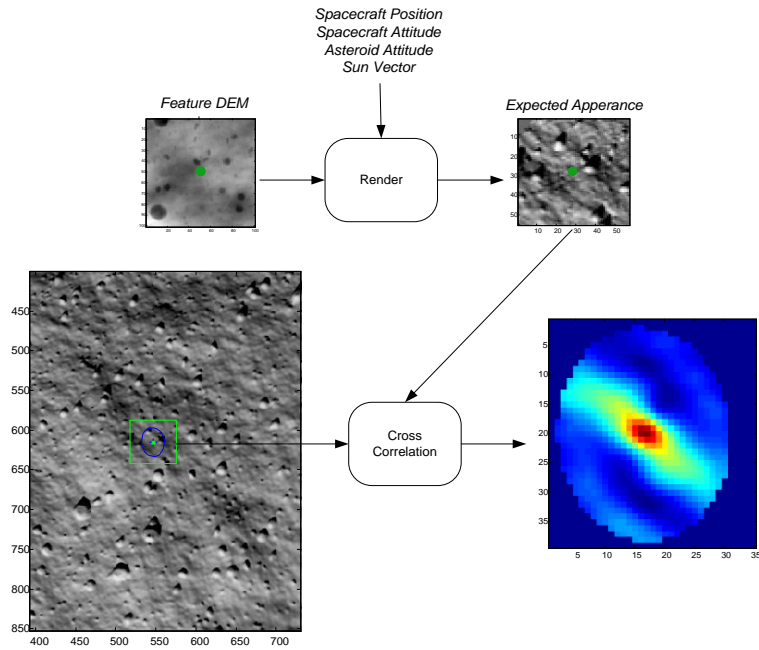
Figure 5 - Example Correlation Surface [5]

Shape Model

As a result of how NFT renders the expected appearance of features onboard using shape model data, NFT's performance is directly related to how well the shape model data represents the asteroid surface. NFT renders from a feature catalogue of approximately 300 Digital Terrain Maps (DTMs) that contain information about the features shape and relative albedo. The DTMs are created by the OSIRIS-REx Altimetry Working Group (ALTWG) using one of two capabilities: the OLA, or Stereo Photoclinometry (SPC). Both processes are iterative; with each step producing higher-fidelity products until requirements are satisfied or successive iterations yield no improvement. The availability of two options for the shape models provides redundancy for both science and operations.

An OLA shape model is built from range measurements collected from OLA and begins by generating low-resolution surface maps from a global set of OLA measurements. The OLA data are then separated into a suite of local maps that are compared to the global shape, and the OLA data are adjusted to the global map using an iterative closest-approach algorithm, producing the required products, including a shape model.

An SPC shape model is built by processing of images from the OCAMS. Using multiple images with different emission and incidence angles, SPC combines standard stereo techniques with photoclinometry, which derives the slope of the surface at each pixel [6][7]. This is a computationally intensive process, and a global shape model takes one to two weeks, depending on the available processing assets. SPC then uses the surface slopes to produce topography in local regions and then collates the local maps to produce a global

shape model. For the albedo data needed by NFT, SPC processing automatically produces albedo, which is a component of surface reflectance and required to extract the slope from multiple images. However, surface albedo is convolved with topography, and SPC albedo can have topography-related artifacts, particularly if the emission and incidence angles of the component images are less than optimal.

OLA and SPC shape models are different, each with advantages that can be exploited, depending on the use. SPC and altimeter models differ primarily in 1) the accuracy of the elevations and 2) the completeness of the terrain. The accuracy of elevations in altimeter-based shape models depends on the range uncertainty of the individual altimeter measurements, which are more accurate than image-based methods such as stereo or SPC. Elevations are suppressed in SPC products due to regional averaging that smears steep terrain [8]. With four to eight overlapping images required for SPC, its shape models are usually contain few gaps. SPC also has better coverage than nominal stereo as all pixels are used, not just the control points. The completeness of both altimetry models depends on the coverage of the altimetry measurements. For OLA, coverage will be thorough due primarily to OLA's scanning capability, which can saturate (the separation between laser spots is similar to the size of a laser spot on the surface) a 6x6-deg region in one pass. For SPC, OCAMS has a long-range camera that provides higher spatial resolution than OLA at the same distance for additional coverage. Although SPC processing is more time-consuming than OLA processing, SPC products can be available before OLA shape models with the same resolution

5

because the data will be available sooner in the mission timeline.

2. SYSTEMS AND INTERFACES LESSONS

Systems and interfaces

Involving the science team data products in subsequent spacecraft operations is a somewhat unique situation for OSIRIS-REx as compared to many other missions and as such presented a set of unique challenges. NFT is a complex system that relies on a detailed set of DTMs that allow it to correlate image renderings with actual images. This interplay between the engineering effort to develop the NFT system that is part of the Flight Software (FSW) (and spacecraft operations) and the ALTWG effort to develop adequate DTMs requires a significant systems engineering effort right from the beginning. Initially, great effort was put into defining requirements that could be well understood and adequately constrained the system. However, the complex interfaces used to produce shape models needed by the NFT rendering software consisted of technical details that were not easily communicated and represented by the requirements put into place. This led to a process that had difficulty in properly defining and interpreting the constraints on the system. As issues began to arise as a result of this environment, systems team members in the program office were added to help resolve these issues and the requirements were altered to make the verification process much more integrated between the NFT and ALTWG teams. This collaborative verification process allowed the requirements to be written at a higher level which made defining and interpreting the requirements much simpler. The downside to this was that the verification was now made more complex.

.

Figure 6 below shows the complex, iterative process for building DTMs and the NFT onboard feature catalog. As ALTWG generates DTMs for NFT from sensor data, NFT will test each feature and either accept it for the final catalog or continue working with ALTWG to improve the feature quality. While this is ultimately the best solution to ensure NFT features will successfully correlate in flight, it requires a lot of data to be passed back and forth between teams.

Backup System Challenges

Because of its backup status and late addition to the program, there were several unique challenges the NFT development had to overcome to produce a fully validated and verified system. These challenges are briefly described below.

The first challenge was the unforeseen complexity in defining requirements on shape model accuracy between the NFT subsystem and the ALTWG interface, which will be discussed in greater detail in a later section. The fact that NFT was considered a backup system exacerbated this situation as initially, NFT was not given the priority needed to efficiently resolve the issues. The second challenge is related to the first and involved ensuring that all of the teams and subsystems which shared a new interface with the new NFT subsystem properly impacted their cost and schedule to accommodate the increase in scope. This was difficult for teams outside of the NFT subsystem to do given the NFT system status as a backup. This backup status made any NFT related work for external teams as low priority (per design). These interfaces included low-level interfaces such as spacecraft FSW but they also included high-level interfaces such as production of special shape models from
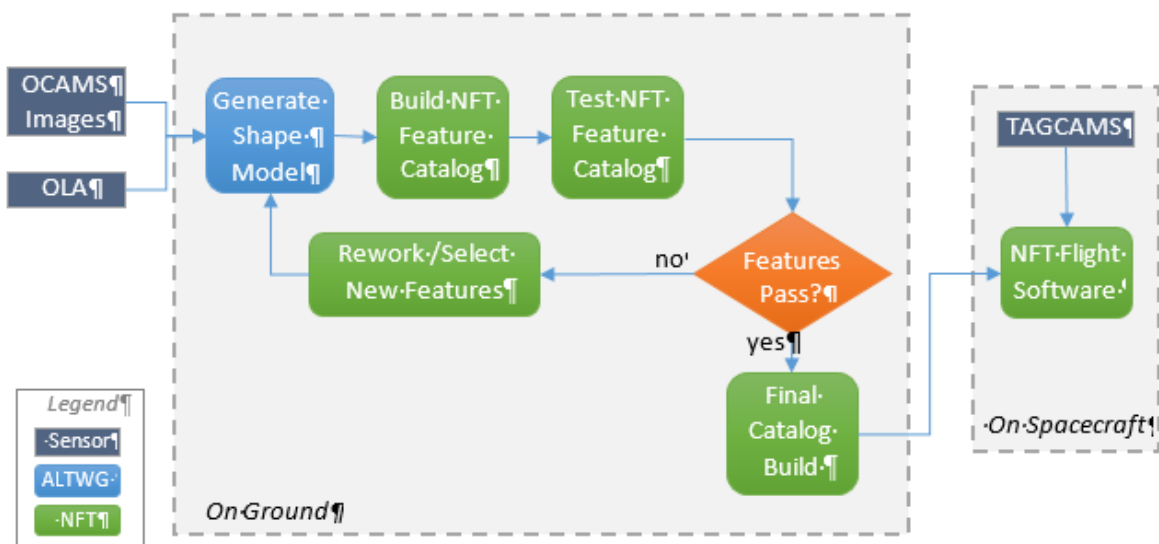

Figure 6 - NFT Shape Model Product Flow

the ALTWG and Navigation team data products. At times, this made the NFT schedule hard to keep on track. The external teams were eventually augmented to accommodate the additional work but it did cause delays.

One thing that was done very well was that a dedicated subsystem was created with the sole purpose of developing, validating and verifying the NFT system. This team was charged with verifying the system as if it was the prime navigation system and this amount of focus drove out many high level issues on the road to producing a robust system for the mission to utilize.

3. SHAPE MODEL

Ground Sample Distances — Initially, shape model requirements were defined as two sets of numbers for accuracy, precision, and resolution or ground sample distance (GSD) – one set for the global shape model and the second set for the TAG site area. The global model was the entire shape model and defined to have a GSD of 35 cm/pixel, while the TAG site area was the 25x25m area around the TAG site and defined to have a GSD of 5cm/pixel.

Upon further testing however, these requirements proved to be insufficient for NFT to be able to build features from the shape model and successfully correlate them in flight. This was primarily due to the fact that the NFT trajectory requires smaller GSD shape model data (ie higher resolution) than the original binary requirements provided. Figure 7 shows how the GSD of the NFT camera decreases closer to the surface, to a final GSD ~1cm/pixel. On the global scale, several images have GSDs much higher than the global 35cm/pixel
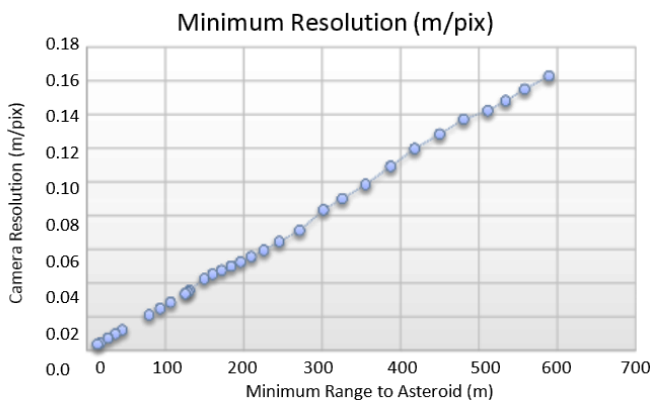

Figure 7 - Example Image Resolution during TAG Trajectory

shape model would provide. Similarly, the camera GSD drops below 5cm/pixel starting around altitude 150m, yet the TAG site area at 5cm/pixel is only visible in the final few images. As a result, with these requirements, a large majority of the features used by NFT would not be adequately supported by the shape model. Testing has shown that the shape model needs to have nearly the same GSD as the NFT image in which it will be used in order to successfully correlate.

Given the high resolution of shape model data needed for NFT, and the large surface area covered by NFT images, it was not practical to define a global shape model resolution requirement as it would require a significant amount of processing by the ALTWG to produce it. In reality the areas of the shape model used for NFT features would amount to only a small percentage of the total shape model. In addition, without a truth shape model of the asteroid, which will not be available in flight, there is no way in flight to test that the shape model meets the accuracy and precision requirements. This resulted in the NFT shape model requirements evolving from accuracy and precision at the global and TAG site level, to a single set of correlation requirements that the shape model must meet for all NFT features. The correlation requirements are now defined such that the shape model data used for NFT features must meet given correlation pixel error and correlation score thresholds for the primary peak, while all secondary correlation peaks score below the correlation score threshold to avoid accepting false matches. With the shape model requirements defined in this manner, not only are the requirements testable in flight using flight images from OCAMS, but they ensure that the shape model data used to build NFT features will enable successful and robust correlation.

SPC vs OLA — Since the shape model data is so critical to NFT's performance, a lot of time was spent testing both SPC and OLA shape model products, which provided a lot of insight into the strengths and weakness of both models for input to NFT.

One of the main advantages of SPC is that albedo data is also derived as part of the solution. NFT utilizes albedo in its rendering algorithm to better match the variations on the asteroid surface which often results in features with higher correlation scores. In addition, albedo is especially useful for correlation for images with little to no shadows (i.e. low incidence angle). On the processing side, SPC can also support early testing of the high altitude NFT features since the image data needed by SPC for those features will be available earlier in flight.

However, one of the risks associated with SPC is that the SPC solution can inadvertently smooth out the shape of features. For example, the sharp edges of a rock or crater are often smoothed, which negatively impacts how NFT renders the feature as resulting shadows from the smoothed out terrain will often be the incorrect size and shape. This causes larger correlation pixels errors due to either broader correlation peaks or the correlation locking onto the edge of the shadow rather than the center due to the size mismatch. In addition, accurate SPC models require a variety of camera angles and sun angles used for deriving SPC model to build robust features, which can impact the data collection part of the mission.

For OLA, one of the main advantages is that the shape of the terrain is well-maintained. Unlike SPC, the edges of boulders and craters are not smoothed out, which results in sharper correlation peaks with smaller pixel errors. However, albedo is not automatically included as part of the solution. Adding

albedo that is properly registered to the OLA shape model proved to be quite problematic.

The testing of OLA-generated DTMs showed that this uncertainty in OLA measurements created unacceptable errors in the elevations of smaller features, the ones near MP and the TAG site. With a 3-cm measurement error for each OLA data point, the smaller rocks and craters, which frequently had topography variations on that order, were masked by the OLA measurement error. For larger rocks and craters, the OLA DTMs met the correlation pixel error requirements but often struggled to meet the correlation score requirement. The accuracy of OLA-generated DTMs was further degraded by the mis-match in resolution: the source OLA data had spot sizes of 7 cm and the required resolution was less than 3 cm/pixel.

The test team evaluated several approaches to improve the accuracy of the small-feature OLA DTMs. One approach was to oversample the area, anticipating that averaging the overlapping OLA measurements would damp the errors in the individual OLA ranges and improve the DTMs. This technique produced only a marginal improvement as the OLA spot size was large compared to the required spatial resolution. A second approach, to add measurements from lower altitudes, had more success as the spot size of each OLA measurement was reduced. Since reducing the altitude does not reduce the range errors, the DTMs were still impacted by the elevation errors, although more features were able to meet the correlation requirements. However, the improvement was sufficient to meet the NFT correlation requirements for the sample features.

Another factor in creating DTMs for operational use is the time required to produce the features. OLA processing, which is usually faster than SPC, slows substantially when required to solve for errors in spacecraft` position, a requirement necessary for the highest-resolution DTMs. (In flight, spacecraft parameters such as position and attitude will not be perfectly known.) The ALTWG team is developing a second approach that is an order of magnitude faster, and is adequate for NFT timelines.

Software Maturity Level — There are two stages of processing to generate the shape models. For SPC, the initial processing has been certified to Class B software and is mature. This is the processing that solves for spacecraft errors and combines multiple images to create local topography maps. To use this software requires extensive training, and the ALTWG team has several trained personnel. For OLA, the initial processing produces a low-resolution global shape model from the raw data. This processing is less mature, and, as mentioned above, is under development to improve processing time that is required to adjust for factors such as errors in spacecraft position. The team could not validate the NFT capability using these DTMs until the software became mature enough for testing.

4. SPACECRAFT CONSIDERATIONS

Binning — NFT testing with shape model data has shown that the underlying shape model resolution directly impacts NFT's ability to successfully correlate, which is especially relevant for the features used in low altitude images where the GSD can be as small as ~1cm/pixel (see Figure 7). While collecting more data would correct this problem, it would significantly impact the mission timeline. Another alternative would be to leverage image binning to reduce the resolution of the image itself. Binning combines adjacent pixels and therefore reduces the spatial resolution. For example, in the case of 2x2 pixel binning, 4-1cm pixel are combined into a single 2cm pixel. Using binned images would allow NFT to navigate all the way surface using shape model data that on full resolution images would only work through CP. This could also reduce the computation time required for rendering and correlating features and could also enable the use of more than 5 features per image.

However, while binning would support the low altitude images, it could also negatively impact performance of the higher altitude images. On a binned image, 1 pixel of error corresponds to a larger absolute distance error than it would on the full resolution image. In addition, initial testing also shows that binning increases the correlation score of the secondary peaks. This could especially impact higher altitude images that are already at risk of false matches due to the correlation search regions including a larger search area. In the future, dynamic binning, where only certain NFT images would be binned, could leverage binning at lower altitude images to help improve correlation with limited shape model resolution, while still maintaining high fidelity performance for high altitude images using full resolution images.

Depending on the camera modes available, binning can either be implemented on the camera itself or by averaging within FSW. When binning is done by the camera sensor, overall noise is reduced, increasing the signal-to-noise ratio of the image. The time to transmit the image from the camera to FSW is also reduced due to the image containing fewer pixels. If binning is done by averaging pixels from the full resolution image within FSW, this can give more flexibility in how binning is done. However, this still requires transmitting the full resolution image from the camera to FSW. It can also be especially computationally intensive to average pixels over the entire image, although this could be reduced by only averaging pixels within the correlation search region areas of the image.

Gravitational Uncertainties — The NFT software has the capability to predict the orbital state of the OSIRIS-REx spacecraft at future epochs in the trajectory. This capability is used to propagate between optical measurements, but also to inform other modules of the spacecraft guidance software. When navigating around small bodies, several models are needed to perform this propagation, one of which includes an accurate gravity model. Building an accurate gravity model of a small body like Bennu has many challenges. Higher order terms in a spherical harmonics representation may not be observable when flying the spacecraft at altitudes typical with a quiescent orbit around the asteroid. This leads to relatively large uncertainties in the modeled gravity field at

low altitudes. In addition, the uncertainties in many cases may not even be completely quantified by the estimated covariance matrix for the model.

Because OSIRIS-REx will make contact with the surface to collect a sample, NFT must deal with these higher uncertainties as the spacecraft descends.

Feature Selection — Selecting features for use with the NFT software requires keeping several metrics in mind to ensure good performance. A high level summary of some of these guidelines is provided in this section. Recall that a feature in this context is comprised of a relatively small area on the surface of the asteroid that is represented in a rendered image patch (see Figure 5).

First, the characteristics of the feature should be considered. Features which have steep slopes and a large vertical relief may be more difficult to model accurately. On the other hand, features like this usually are comprised of areas of high contrast which make them very identifiable and attractive to a correlator. A balance must be found between these considerations.

A feature must also be unique when compared to its surroundings (see Figure 8). A single rock for example can

be more easily confused with another rock whereas a grouping or distinct pattern of rocks is much more unique. With the pattern, NFT is less susceptible to shadow errors caused by smoothed out terrain as the correlation algorithms will lock onto the center of the pattern of shadows. The NFT software has been designed to sort out misidentified features, but a robust feature choice from the beginning enables better performance overall.

Furthermore, the lighting conditions under which the feature is to be observed must be considered. Lighting conditions can partially be described by the incidence angle. This is the angle between the feature-to-Sun vector and the surface zenith vector. At high incidence angles, the feature may be too dark to use, but at very small incidence angles the feature may become "washed out" and dominated more by relative albedo (see Figure 9). All of these effects are taken into consideration by the NFT software. Additional SPC robustness testing results are described in [9]. Feature arrangement in the camera focal plane must also be considered. A feature arrangement that is clustered into one corner of the camera focal plane my lead to poor observability into the NFT state estimate. This is avoided by selecting features that are well distributed and adjusting the ordering in the catalog.
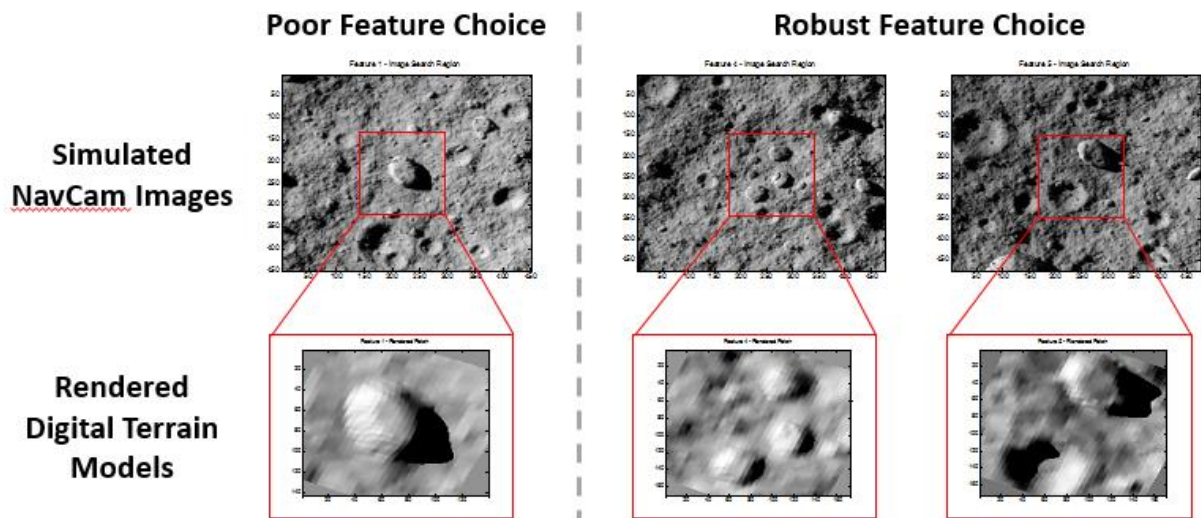
.

Figure 8 - Comparison of Robust and Poor Features
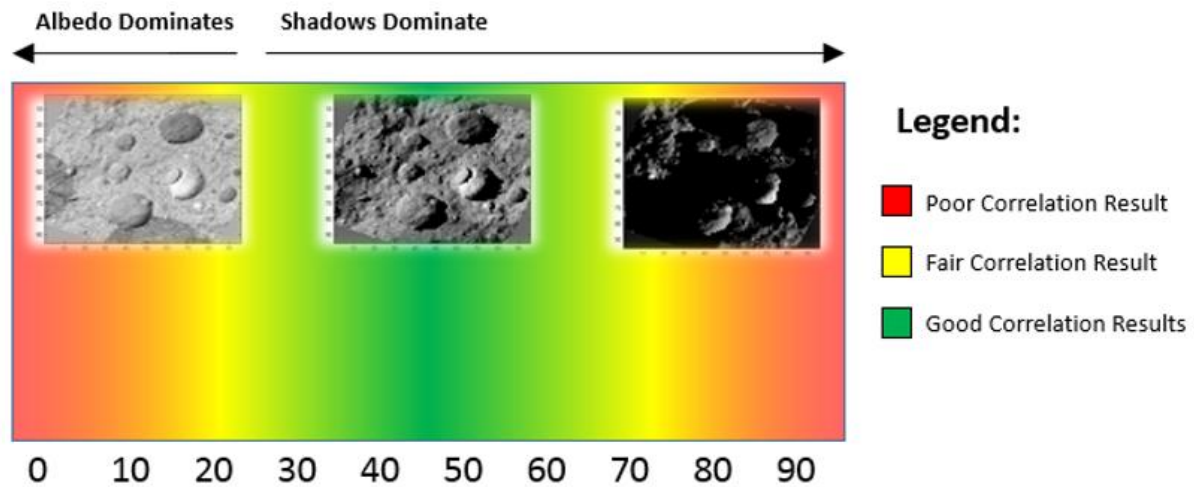
Figure 9 - Solar Incidence Angle (Deg)

NFT has the benefit of using a catalog of known features, unlike other odometry-only systems which detect features on the fly and can therefore define their locations incorrectly. However, the success of using known features is directly tied to how well the shape model represents the asteroid surface so that NFT can reliably correlate. The shape model requirements for NFT are defined to ensure that the primary correlation peak has a pixel error below a pixel error threshold and a correlation score above a score threshold, as well as that the correlation secondary peaks are all below the same correlation score threshold, eliminating the risk of false matches. The correlation score for NFT was set to 0.6. The correlation pixel error threshold has been set using error budgets to determine the max allowable pixel error per feature needed for NFT to still meet TAG requirements. Since the shape model data may be produced using either SPC or OLA, testing has been done using features from both models and pros and cons for both types of models have been identified.

5. TEST DATA/TRUTH MODEL

In order to test the capabilities of shape models being generated by the ALTWG, a set of truth models were developed that modelled the asteroid topography and albedo. The models were generated to mimic both an expected case (based on previous small body encounters) but also stressing cases for both the topography and albedo variations. Using these models, sample images and Lidar signal returns were generated. These data were then provided to the ALTWG in order to generate the shape models required for validation. During early formulation of the mission, the specifications for this truth model were generated. The accepted GSD for the truth model was generally between 1cm and 5cm. A reevaluation of the truth model requirements used for testing the ALTWG products was not performed following the late arrival of NFT as a backup system.

However, as shown in Figure 7, the final image used by NFT for TAG navigation, taken approximately 3 minutes prior to TAG, has a resolution of about 0.7 cm. This image resolution

is the stressing case for generating DTMs and is also the stressing case for validation. Unfortunately the truth models available for verification did not fully support testing to this level. Several of the truth models had small regions with high enough resolution but these also had some unrealistic characteristics that made testing difficult. Additionally the regions with high resolution were somewhat smaller than the NFT camera Field Of View. With the available truth models, a test program was cobbled together that was satisfactory but it involved combining results from various tests rather than having one clean test program. The test program could have been made much simpler and a significant amount of time and effort would have been saved if during the NFT formulation, the test capabilities for the requisite DTMs would have been reexamined.

In addition, the testing of ALTWG's ability to properly generate DTMs involved multiple groups. It involved the NFT team and several groups within the ALTWG because the DTMs not only included topographic information but albedo information as well. Testing of the DTMs involved one group generating the shape of the feature and passing this to the group that would generate the albedo information. Based on which technique was being used, the DTMs sometimes then went back to the topographic group to help register the two pieces of information. Once that was completed, the DTMs were then passed to the NFT team at LM for evaluation. If the DTMs did not correlate, it required that the teams work together remotely to understand where the problem may lie. Given the multiple organization in generating and evaluating the test data, this test became rather complex and difficult to manage.

One simplification to the testing did help immensely. Rather than testing the final DTMs as a whole, the generation and testing of the DTMs was broken up into its component parts. First just the topographic data was provided in the DTMs without any albedo information. It was assumed that there was no variation in albedo across the feature DTM. The NFT then correlated the images assuming against images generated assuming there was no albedo variation. This

allowed the team to solve any issues associated with the topographic information before adding in another factor that would complicate the testing results. Once the topographic issues were resolved, the albedo information was added to the test data set.

6. CONCLUSION

The NFT system was added after spacecraft CDR as a backup navigation system to augment the existing Lidar navigation baselined for OSIRIS-REx. This optical navigation system relies heavily on the ability to correlate real-time images against rendered images generated from a catalog of feature shape models. Bringing this system on late in the process and as a backup system contributed to the fact that it was not viewed as a dedicated subsystem until much later in the development of NFT. The involvement of multiple organizations within the project needed for success was not fully realized.

The capabilities of developing the requisite shape model were not fully appreciated especially for the small GSDs required near TAG. Multiple techniques for developing the shape model had to be evaluated. Each with different strengths and weaknesses that were not fully appreciated until the stressing requirements of NFT were placed upon them.

The proper criteria for selecting features was not fully understood until testing began necessitating an enhancement in the ground tool used for identifying features. There were a number of other enhancements that may have proved useful during the development of NFT such as image binning. Finally, the specification of the asteroid truth model used for system testing should have been reviewed once NFT was accepted as a backup navigation system.

ACKNOWLEDGEMENTS

REFERENCES

[1] Lauretta, D. S., and OSIRIS-Rex Team. "An overview of the OSIRIS-REx asteroid sample return mission." In Lunar and Planetary Institute Science Conference Abstracts, vol. 43, p. 2491. 2012.

[2] Alexander May, Brian Sutter, Timothy Linn, Beau Bierhaus, Kevin Berry, Ron Mink. "OSIRIS-REX TOUCH-AND-GO (TAG) MISSION DESIGN FOR ASTEROID SAMPLE COLLECTION." IAC-14-A3.4.8 2014.

[3] May, A., Sutter, B., Linn, T., Bierhaus, B., Berry, K., Mink, R. "OSIRIS-REx Touch-And-Go (TAG) Mission Design for Asteroid Sample Collection." International Astronautical Federation (IAC) Conference Abstract-65th Annual, IAC-14-A3.4.8, 2014.

[4] Kevin Berry, Brian Sutter, Alex May, Ken Williams, Brent W. Barbee, Mark Beckman, and Bobby Williams. "OSIRIS-REx Touch-And-Go (TAG) Mission Design and Analysis." AAS 13-095 2013.

[5] Ryan Olds, Alexander May, Courtney Mario, Reid Hamilton, Chris Debrunner, and Kalle Anderson. "The Application of Optical-Based Feature Tracking to OSIRIS-REx Asteroid Sample Collection." AAS 15-114 2015.

[6] Gaskell, R W, O S Barnouin-Jha, D J Scheeres, A S Konopliv, T Mukai, S Abe, J Saito, et al. "Characterizing and Navigating Small Bodies with Imaging Data." Meteoritics & Planetary Science 43, no. 6 (September 2008): 1049–61. doi:10.1111/j.1945-5100.2008.tb00692.x.

[7] Gaskell, R.W. (2011). Optical navigation near small bodies. Spaceflight Mechanics 2011, Pts I-III, vol 140, 1705-1717.

[8] Perry, M. E., Kahan, D. S., Barnouin, O. S., Ernst, C. M., Solomon, S. C., Zuber, M. T., Smith, D. E., Phillips, R. J., Srinivasan, D. K., Oberst, J. and Asmar, S. W. (2011). Measurement of the radius of Mercury by radio occultation during the MESSENGER flybys. Planetary and Space Science, 59, 1925-1931.

[9] Courtney Mario and Chris Debrunner. "Robustness and Performance Impacts of Optical-Based Feature Tracking to OSIRIS-REx." AAS 16-087 2016.

BIOGRAPHY

Alexander May was the TAG Phase Lead for OSIRIS-REx through the design, development, and test of the spacecraft. Prior to ORISIS-REx he was a GN&C engineer specializing in spacecraft navigation and autonomous systems. Mr. May has been with Lockheed Martin for 14 years and has experienced working missions from proposal through operations from LEO to GEO and soon asteroid Bennu. He has also led several IRAD efforts developing rendezvous and proximity operation technologies. Alex graduated from Lockheed Martin's Advanced Technical Leadership Program in 2014. He received his aerospace engineering bachelor's degree with honors from Iowa State University and master's degree from University of Colorado. He is on the board of directors for the American Astronautical Society and has been a chairperson the past 12 years for the annual GN&C conference.

David Lorenz is the TAG Campaign Lead for the OSIRIS-REx project at Goddard Space Flight Center. Mr. Lorenz has worked at GSFC for over 34 years starting in attitude operations for the first on-orbit repair of a satellite – Solar Maximum Repair Mission. He has worked as a GN&C engineer for multiple spacecraft as well as mission operations for Landsats-4 & -5, Terra and Aqua. Mr. Lorenz developed the operations plan and FSW changes that allowed Landsat-4 to be returned to nearly full operations following the failure of 2 of 4 of its solar panels. He served as shuttle integration manager for the UARS mission and then moved to mission operations for this mission to assist in maintaining the full mission despite the failure of the solar array drive and battery anomalies. Mr. Lorenz also was part of the launch and early check-out team for multiple GOES satellites (I, J, L, N, & O) as well as assisting in the development of the INR system for GOES-R before coming to OSIRIS-REx. Mr. Lorenz received his bachelor's degree in Physics from Thomas More College and his masters in Physics from the University of Maryland.

Dr. Michael Daly is the York University Research Chair in Planetary Science. His research interests focus on answering a variety of planetary science questions using custom instrumentation in the laboratory or in-situ. Dr. Daly is currently leading the science contribution of Canada's OSIRIS-REx Laser Altimeter that was launched in September. He also works in the area of deep-UV Raman spectroscopy and is currently building a $3.5M planetary surface simulation facility. Mike is also the Undergraduate Program Director for York's unique Space Engineering Program. Prior to joining York University, he led the engineering of Canada's first instruments to operate on Mars and the design of the cameras in the International Space Station's Dextre robot's end-effectors.

Courtney Mario is a Perception Systems Engineer at Draper in Cambridge, MA. She is currently a member of the Natural Feature Tracking team for OSIRIS-REx, focusing on the image correlation performance for the different shape model inputs. Her other work at Draper has involved developing vision-inertial systems for GPS-denied applications for ground vehicles, UAVs, and pedestrians. Prior to Draper, her grad school research involved image-based vehicle lane detection algorithms as well as analysis of airport ground surveillance systems for use in automated operations. She received a BS and MS in Mechanical Engineering from Tufts University in 2009 and 2011, respectively.

Mark Perry, a planetary scientist at the Johns Hopkins Applied Physics Laboratory, performs system engineering and science analysis for the OSIRIS-REx Altimetry Working Group, which provides the topographic data used for both optical navigation and science investigations. Over 25 years, Dr. Perry has worked as scientist, engineer, and manager on deep space missions, Earth-orbiting satellites, the space station, and astrophysics telescopes. Recently, he analyzed RF occultations to conduct geophysical studies of Mercury for the MESSENGER mission. Other current work includes using observations from the Cassini Ion and Neutral Mass Spectrometer to characterize Saturn's magnetosphere and the dynamic nature of the plumes of Enceladus.

Ryan Olds is the Natural Feature Tracking CPE and team lead for the OSIRIS-REx Guidance Navigation and Control design and operations team at Lockheed Martin Space Systems Company. He oversaw all design, V&V and test activities for the Natural Feature Tracking subsystem for the OSIRIS-REx mission. Prior to Natural Feature Tracking, Ryan developed and designed the OSIRIS-REx attitude control system. He continues on the OSIRIS-REx program as both GN&C and NFT team lead during operations and will oversee all GN&C efforts to maneuver the spacecraft around the asteroid Bennu and ultimately collect a sample from the surface. Ryan has been with Lockheed Martin Space Systems for 12 years and prior to OSIRIS-REx he has worked in Guidance Navigation and Controls for several NASA missions including a Lunar gravity mapping mission (GRAIL), Mars Entry Descent and Landing (Phoenix) and precision pointing of a NASA great observatory (Spitzer). Ryan received both his Bachelor's and Master's degree in Aerospace Engineering from the University of Colorado at Boulder.

Eric E. Palmer was the lead for testing the photoclinometry software suite to generate the global shape model and specific navigational features to be used by NFT during the OSIRIS-REx mission. He is a research scientist that focuses on thermal modeling and developing topography of asteroids, comets, the Moon and the Earth. Dr. Palmer works at Planetary Science Institute for the last six years and is the principle investigator for the Asteroid and Dust Subnode of the Planetary Data Systems. Eric received a computer science degree from Purdue University and his Ph.D. in Planetary Science from the University of Arizona.