Designing to Sample the Unknown: Lessons from OSIRIS-REx Project Systems Engineering

David Everett
NASA Goddard Space Flight Center
8800 Greenbelt Road
Greenbelt, MD 20771
301-286-1596
David.F.Everett@nasa.gov

Ronald Mink
NASA Goddard Space Flight Center
8800 Greenbelt Road
Greenbelt, MD 20771
301-286-3524
ronald.g.mink@nasa.gov

Timothy Linn
Lockheed Martin Space Systems
Company, P.O. Box 179
Denver, CO 80201
303-977-0659
timothy.m.linn@lmco.com

Joshua Wood
Lockheed Martin Space Systems
Company, P.O. Box 179
Denver, CO 80201
303-977-3199
joshua.l.wood@lmco.com

Abstract—On September 8, 2016, the third NASA New Frontiers mission launched on an Atlas V 411. The Origins, Spectral Interpretation, Resource Identification, Security-Regolith Explorer (OSIRIS-REx) will rendezvous with asteroid Bennu in 2018, collect a sample in 2020, and return that sample to Earth in September 2023. The development team has overcome a number of challenges in order to design and build a system that will make contact with an unexplored, airless, low-gravity body. This paper will provide an overview of the mission, then focus in on the system-level challenges and some of the key system-level processes. Some of the lessons here are unique to the type of mission, like discussion of operating at a largely-unknown, low-gravity object. Other lessons, particularly from the build phase, have broad implications.   The OSIRIS-REx risk management process was particularly effective in achieving an on-time and under-budget development effort. The systematic requirements management and verification and the system validation also helped identify numerous potential problems. The final assessment of the OSIRIS-REx performance will need to wait until the sample is returned in 2023, but this post-launch assessment will capture some of the key systems-engineering lessons from the development team.

TABLE OF CONTENTS

1. INTRODUCTION

On September 8, 2016, the third NASA New Frontiers mission launched on an Atlas V 411. The Origins, Spectral Interpretation, Resource Identification, Security-Regolith Explorer (OSIRIS-REx) will rendezvous with asteroid Bennu in 2018, collect a sample in 2020, and return that sample to Earth in September 2023. This complex mission has required a team of well over 1500 people to get this point, with flight hardware on its way to Bennu and a ground system and operations team controlling it. Effective systems engineering is essential for the success of any endeavor of this scale.

OSIRIS-REx was competitively selected in May 2011. The Principal Investigator (PI), Dante Lauretta, is at the University of Arizona, and he partnered with Lockheed Martin (LM) for the flight system development and with Goddard Space Flight Center (GSFC) for project management, systems engineering, and mission assurance. At confirmation, $611 million was budgeted for development through phase D outside of launch costs. The project actually spent $566 million through phase D, $45 million under what was planned at this point in the project. The project met all the major schedule milestones laid out in the proposal, including launch within the 39-day period available to reach Bennu.

This successful performance was not easy. The team has learned a lot about the complexities of operating near and contacting a small, unexplored object. The sampling technology and the lidar sensor technology were major developments for the project. And the sheer size of such a project presents challenges in coordinating the effort and keeping everyone working in the same direction.

This paper provides the project-level systems engineering perspective of the development effort. It covers here some of the key challenges encountered and some of the key processes used to get where the project is today. The focus is on useful lessons for future systems engineers, both strategies to copy and things to avoid.

2. MISSION OVERVIEW

This section provides a brief overview of the mission design. For more details, see Beshore [1].

Design Reference Mission

The OSIRIS-REx Design Reference Mission (DRM) served as the operations concept for the entire mission, launch through sample return to Earth. The DRM aided in the flow down of science objectives to requirements at the system, subsystem, and lower levels; served as a validation tool for use across the project to ensure mission element designs were compatible with the operations concept; and provided a blueprint for conducting spacecraft operations, science operations, and navigation. The DRM was also used to develop scenarios for flight and ground system testing, and it served as a kernel from which the science and engineering teams developed detailed operations plans and procedures. The DRM provided a stable operations concept for completing the spacecraft development and testing and for operations planning.

OSIRIS-REx launched on September 8, 2016, and will flyby the Earth a year later to get a gravity-assisted delta-v that will put the spacecraft in the same orbital plane as Bennu. OSIRIS-REx then cruises for 11 months and starts the optical search for Bennu in August 2018, marking the beginning of the Approach Phase. Rendezvous occurs in October 2018, followed by a month of slow approach to allow the flight system to search for moons around Bennu and to refine its shape and spin state models. The spacecraft will spend over 900 days at Bennu before the first departure opportunity in March of 2021.

Beginning with the Preliminary Survey and Orbital A Phases in November and December 2018, respectively, OSIRIS-REx lays the foundation for navigating the spacecraft through the remainder of the encounter by estimating the mass of Bennu and transitioning from star-based to surface landmark-based optical navigation for precisely estimating the spacecraft's state relative to the asteroid. OSIRIS-REx then performs comprehensive global mapping of the texture, mineralogy, and chemistry of Bennu in the Detailed Survey Phase, resolving geological features, revealing its geologic and dynamic history, and providing context for the returned samples. The Orbital B Phase, beginning in April 2019, establishes the "Safe Home" orbit from which all sorties closer to the surface begin. In the Reconnaissance Phase, the instruments document the regolith at candidate sampling sites in situ at scales down to the sub-centimeter, providing the high-resolution information needed to confirm the presence of sampleable regolith and select the primary sample site. Following site selection, a sequence of Touch-and-Go (TAG) rehearsal steps are conducted to separately demonstrate each step in the sample collection sequence prior to sending the flight system to the surface for sampling. In the Sample Collection Phase beginning in June of 2020, OSIRIS-REx acquires a minimum of 60g of bulk regolith and a separate 26 cm² of fine-grained surface material from Bennu. Additional information on TAG is in Lorenz [2].

After departing from Bennu, the spacecraft spends two and a half years on a ballistic cruise back to Earth. On September 24, 2023, the Sample Return Capsule (SRC) lands at the Utah Test and Training Range. From there the SRC will be transported to NASA's Johnson Space Center, where the samples are removed and delivered to the OSIRIS-REx curation facility.

Mission Architecture

Because OSIRIS-REx plans to sample an object that has never been approached, the mission must be designed to quickly process collected data into knowledge about Bennu, then use that knowledge to plan the next step. The science processing aspect of the ground system is tightly coupled to the navigation, with the science team delivering shape models and other key information related to possible sample sites. The DRM described above led directly to the ground and flight system architectures.

Ground System—OSIRIS-REx's ground system mission operations are performed primarily by three elements: the Mission Support Area (MSA), the Flight Dynamics System (FDS), and the Science Processing and Operations Center (SPOC). OSIRIS-REx also utilizes NASA's Deep Space Network (DSN) for command uplink and telemetry downlink. The OSIRIS-REx architecture diagram (Figure 1) captures the major ground elements and interfaces between them.

The MSA, located at LM's facility in Littleton, Colorado, serves as the spacecraft mission operations activity hub for command integration and uplink, maneuver implementation, spacecraft health & safety monitoring, and telemetry trending and analysis.

The FDS is responsible for trajectory analysis, orbit determination and analysis, and for maneuver and verification during the mission. A joint FDS team of personnel from KinetX and NASA's GSFC are co-located in the MSA during the mission.

The University-of-Arizona-developed SPOC monitors instrument telemetry and performs instrument trending analyses and health & safety monitoring. OSIRIS-REx science and ancillary flight data received by the SPOC is "ingested" into the SPOC Database. The SPOC then "pipelines" instrument data using algorithms developed by OSIRIS-REx instrument providers to produce calibrated and validated products to be stored in the database. Science data processing and analysis functions hosted by the SPOC use instrument products to generate science data products that are integrated to produce maps used for sample-site selection.
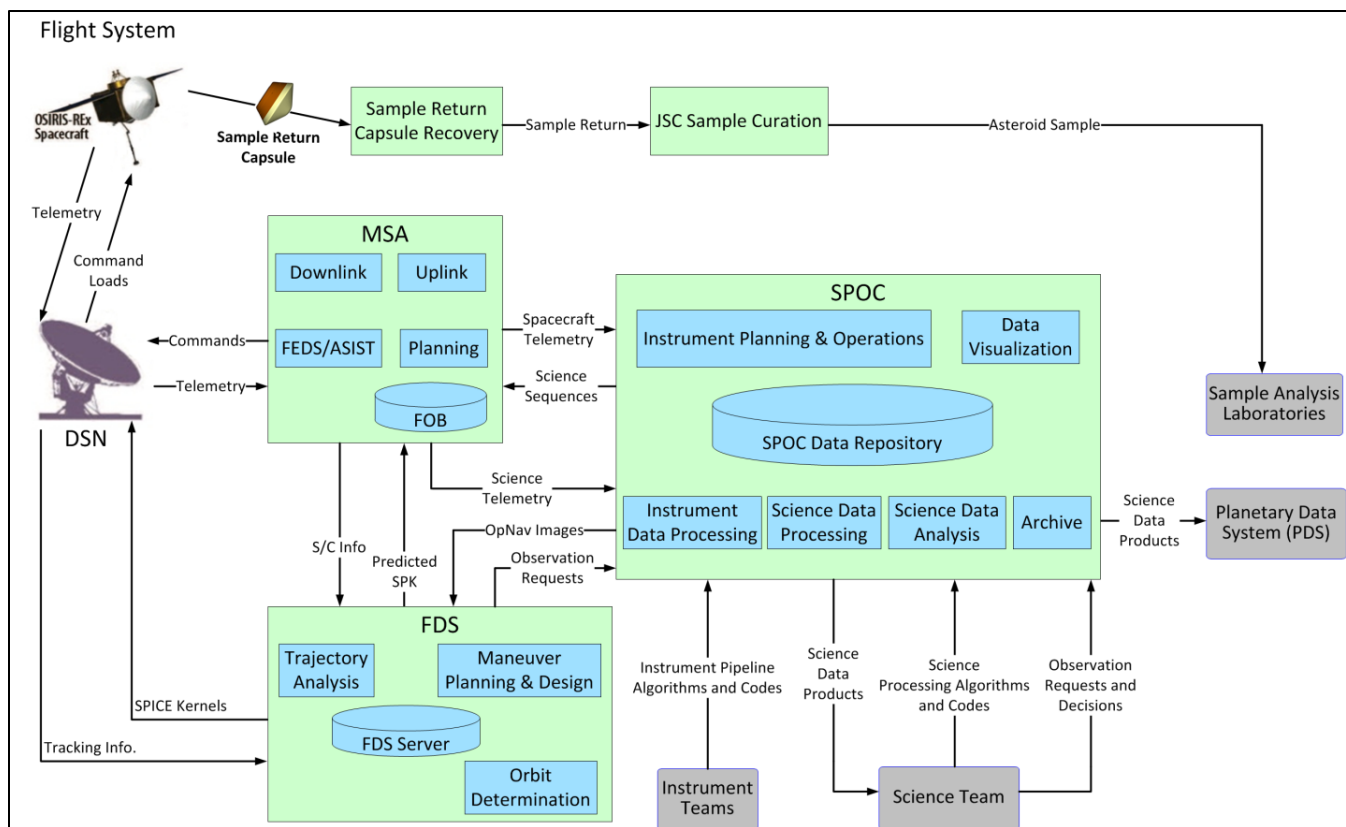
Figure 1: OSIRIS-REx Ground System Architecture

Launch Vehicle—An Atlas V 411 delivered the 2104 kg OSIRIS-REx to its targeted outbound trajectory with a C3 of 29.3 km²/s². As the 65th launch on the Atlas V, the interfaces were well-defined. But there are always some features unique to the mission. The OSIRIS-REx team worked early with the Launch Services Program (LSP) at the Kennedy Space Center, even before the procurement of the launch vehicle, to reduce the acoustic spectrum. The standard vehicle environments are designed to envelope all configurations, so some configurations will have more benign environments. The team came up with a compromise acoustic spectrum that reduced random vibration levels on components to a level consistent with some of the key heritage components.

It is not uncommon for launch vehicles to release margin in the last few months before launch, after the mission profile and any possible changes to the rocket are well defined. At this point in the flow, it is generally too late for the spacecraft to use the extra launch mass capability. OSIRIS-REx pursued with LSP an early release of some (150 kg!) of the launch vehicle margin. With the wet mass limit and expected dry mass at Critical Design Review (CDR), the spacecraft fuel tank had more capacity than we could use. The extra launch mass enabled the project to nearly fill the tank. The change in the CDR time frame gave the team time to adjust the test program to accommodate the greater launch mass. The extra fuel bought the mission 179 m/s of extra delta V, margin that greatly improves the likelihood of success should some problem develop during the mission.

3. FLIGHT SYSTEM OVERVIEW

LM integrated and tested the flight system at their Waterton, CO facility. GSFC oversaw the development of the five science instruments and held the development contract with LM. LM developed the single-fault-tolerant spacecraft bus based on several previous planetary missions. LM also developed all of the hardware necessary for acquisition and return of the asteroid sample.

Instruments

The flight system carries five instruments designed to thoroughly map and characterize Bennu prior to the TAG event. Their positions on the spacecraft are shown in Figure 2.

OSIRIS-REx Camera Suite (OCAMS)—The University of Arizona designed and built OCAMS. This instrument package includes 3 separate single-string cameras and a fully-redundant control electronics. PolyCam is the long-range camera, with a 630 mm aperture, f/3.15, and a 0.8-degree field of view. MapCam provides a majority of the mapping images, with 4 different color filters, a 125 mm aperture, f/3.4, and a 4-degree field of view. SamCam provides a wide-angle field of view that includes the sample head during TAG. It will capture 3 images every 5 seconds with a 24 mm aperture, f/5.5, and a 22-degree field of view. Although each camera is different, the capabilities overlap enough so that all mission objectives can be accomplished with any two of the three cameras.
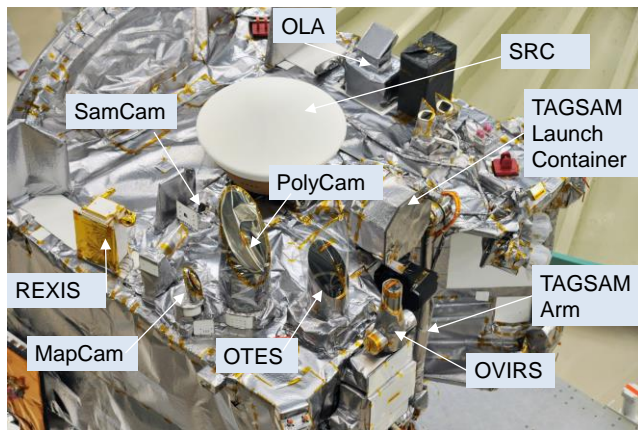
Figure 2: Flight System Instrument Deck

OSIRIS-REx Thermal Emission Spectrometer (OTES)—Arizona State University designed and built OTES. This point spectrometer has a 0.46-degree field of view, covering 5 to 50 μm wavelengths with 0.025 to 2.38 μm (10 cm^{-1} wave number) resolution.

OSIRIS-REx Visible and near-InfraRed Spectrometer (OVIRS)—GSFC designed and built OVIRS. This point spectrometer has a 0.23-degree field of view, covering 0.4 to 4.2 μm wavelengths with 140 to 410 ($\lambda/\Delta\lambda$) resolution.

OSIRIS-REx Laser Altimeter (OLA)—The Canadian Space Agency provided this instrument that was designed and built by MacDonald, Dettwiler and Associates Ltd. OLA is a scanning laser altimeter operating at 1064 nm and covering ranges from 0.5 to 7 km.

REgolith X-ray Imaging Spectrometer (REXIS)—The Massachusetts Institute of Technology and Harvard College students designed and built REXIS. The primary purpose of REXIS was providing an opportunity for students to get hands-on experience with space-flight hardware development. The spectrometer covers 0.5 to 7.5 keV, and it will provide additional information on the surface composition of Bennu.

Sample Acquisition and Return Assembly (SARA)

The SARA sub-assembly (shown in Figure 3), designed and built by Lockheed Martin, is a simple composite panel structure that integrates the Touch-And-Go Sample Acquisition Mechanism (TAGSAM) and the SRC into a single unit. The SARA panel consists of both electrical and mechanical assemblies, since all the harnessing and separation components for TAGSAM and SRC are included in the panel design.

By building up the SARA panel early during development, the team could perform an early checkout of the full sample acquisition and stow operation using simple EGSE, rather than requiring expensive Assembly, Test, and Launch Operations (ATLO) time. This testing included a full range of environments prior to System-level environments to confirm no change in performance, even over temperature, avoiding the need for expensive MGSE within the system-

level thermal vacuum chamber. In addition, the SARA assembly allowed alignments of the TAGSAM with SRC for stowing of the sample prior to integration on the bus, further reducing risk of discovering latent problems late in the ATLO flow.
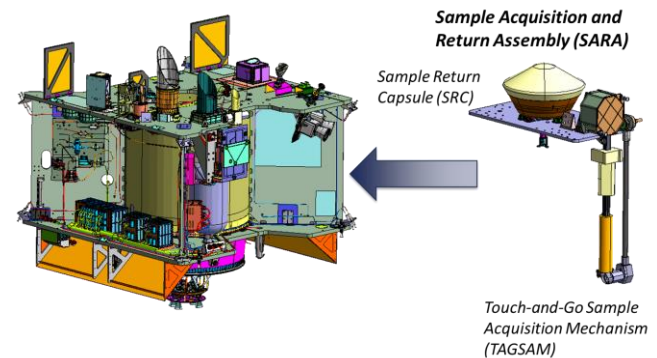

Figure 3: SARA Assembly and its Location on the Flight System

Brief descriptions of TAGSAM and the SRC follow below. For more detailed descriptions of this hardware in its planned operation, see Ajluni [3].

Touch-And-Go Sample Acquisition Mechanism—The TAGSAM sub-assembly was a new development consisting of a simple, single-plane, articulating arm. By utilizing single-axis motors, the design remained simple and compact while meeting all mission requirements.
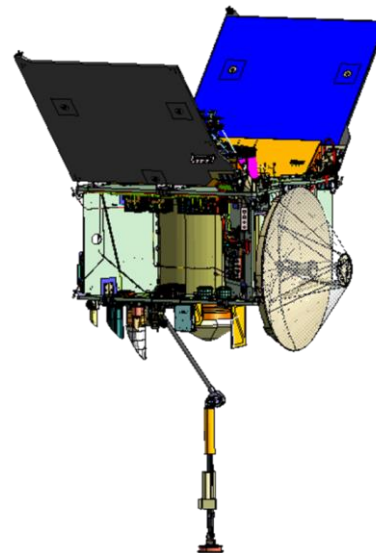

Figure 4: TAG Configuration

To collect a sample, high-grade nitrogen stored in three bottles (to support up to 3 TAG attempts) is injected into the asteroid surface at detection of contact with the surface. Liberated regolith on the surface of the asteroid is drawn into a collection basin via a screening system as the nitrogen vents to space, thus collecting the sample. A simple constant-force spring mechanism is utilized to stop progression of the

spacecraft towards the asteroid in a soft and controlled fashion after making contact. Absolute potentiometers in the shoulder and elbow motors are used to track location of the arm, along with hardstops located to position the sample head along the center of mass of the spacecraft during TAG. Figure 4 shows the TAGSAM configuration during TAG.

Sample Return Capsule—The SRC sub-assembly, shown in Figure 5, consists of a structure and avionics that were build-to-print from the successful Stardust comet sample return mission. Minor modifications were required to accommodate capture of the TAGSAM Sample Head, but the rest of the components required no modification (with exception – see Build Challenges). The SRC contains 3 motors – two for latches, and a third to open and close the heatshield to allow capture of the Sample Head. A simple avionics box with G-Switch and pressure-transducer-based triggers is powered shortly before release from the spacecraft, and is used for the deploy timing of the parachute (drogue and main chute). Batteries provide power throughout the Entry, Descent, and Landing sequence, which lasts up to 5.5 hours for the mission, starting at the point where the SRC is first powered.
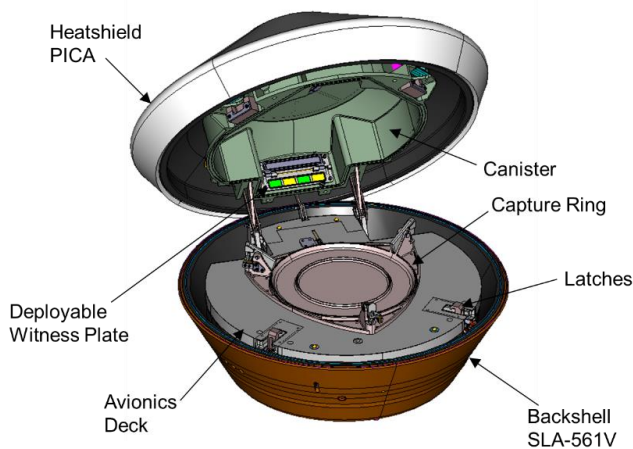

Figure 5: Sample Return Capsule

Spacecraft

The OSIRIS-REx Spacecraft, consisting of the structure and general subsystems is leveraged heavily from previous builds by LM, with most of the architecture taken directly from the Mars Atmosphere and Volatile Evolution mission (MAVEN). The basic architecture consists of an open-bus structure with components mounted in a distributed fashion to maximize real estate and reduce the need for ballast for center-of-mass balancing. To keep cost and schedule risks down, attempts to use heritage wherever possible was employed across all subsystems.

Several key components that were able to leverage full build-to-print from MAVEN include the Reaction Wheels, Inertial Measurement Units, Star Trackers, 4-Headed Sun Sensors, Thrusters (minus the Low Thrust Rocket Engine Assemblies (REA)), Small Deep Space Transponder, Traveling Wave

Tube Amplifier, and both the Low Gain and High Gain Antennas.

Minor modifications were required for the following components and is captured in more detail under the "Build Challenges" section: Command and Data Handling (C&DH) Unit, Power Distribution and Data Unit (PDDU), and Pyro and Propulsion Unit. The Solar Arrays are different dimensions from MAVEN, but same build structure and solar cell technology. The Spacecraft Batteries are composed of the same electronics as MAVEN, but utilize the smaller cell sizes from the Gravity Recovery and Interior Laboratory (GRAIL) program, as OSIRIS-REx does not experience any eclipses, thus no need for large battery capacity.

Finally, the OSIRIS-REx program was able to leverage hardware and designs from past programs as well, which include items like the Solar Array Two-Axis Gimbal, which was a minor modification from the units on Mars Odyssey, the 2-Headed Sun Sensors which are build-to-print from Mars Reconnaissance Orbiter (MRO), the Low Thrust REA's, which are build-to-print from Geostationary Operational Environmental Satellite R (GOES-R), and the Medium Gain Antenna, which was a minor modification from the unit used on the Phoenix Mars Lander.

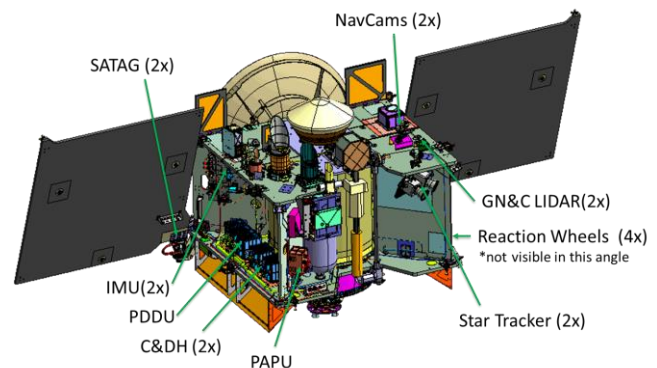Figure 6 shows the distributed nature of the components about the bus.

Figure 6: OSIRIS-REx Spacecraft

However, not all components could be leveraged from past programs. The OSIRIS-REx mission, with requirements to lightly touch an asteroid to collect a sample, drove some new development. This includes, in addition to the previously described TAGSAM, the mission-developed Guidance, Navigation, and Control (GN&C) Lidars, TAG Cameras (TAGCAMS), Natural Feature Tracking (NFT) Autonomy, and a Propellant Management Device (PMD).

GN&C lidar—While lidar applications in-space are not new, ranging against a very low albedo surface over a large range at potentially extreme angles drove a significant new development effort for the lidar system selected. The long duration operations requirements to support TAG after 4.5 years required more stringent hermetic sealing than heritage units, along with a need for higher radiation tolerant parts.

TAGCAMS—During early development for the program, it became apparent that navigation about a small body would be tricky. In addition, due to technical risks being raised about the complexities of the lidar system, an optical back-up to the GN&C lidar was desired (see Autonomous Docking under Design Challenges for more information). To that extent, optical cameras were added to the Flight System that could resolve both features on the surface of the asteroid while also acquiring stars in the same frame, all with short integration times. In the end, unique optical coatings were developed for the camera lens to ensure no issues.

NFT Autonomy—In addition to the Optical Navigation capabilities described above, the program desired a full autonomous back-up to the GN&C Lidar for performing TAG operations. Leveraged from past work on XSS-11, the program was able to develop a NFT capability that uses the TAGCAMS camera system to look for known features on the surface of Bennu (boulders, craters, etc). The system relies on early observations of the asteroid via the existing science campaign to build up a shape model of the surface of Bennu. Once the shape model is developed, reference features are defined and commanded back to the Spacecraft. The Flight System then uses the features to navigate during TAG.

Due to the processor intensive operation of parsing through images looking for features in real-time, care had to be given to ensure task priorities were configured correctly to avoid starving of higher priority tasks. By utilizing time stamping and state propagation, the system can take more time to process images if required, then apply the propagated state to determine the Flight-System's current location. During TAG, two predefined burns can be adjusted autonomously by Flight Software based on the propagated state from NFT. Bounds are placed on the allowable size of maneuver adjustments, thus if NFT shows that the spacecraft is too far off course, the Flight System can back away and return to orbit to allow ground to assess the situation.

PMD—Due to the pogo nature of the TAG operation, the Spacecraft needed to be able to perform maneuvers in both the +Z and −Z directions with minimal time between maneuvers. The size of the tank precluded the use of a bladder-type system, thus a unique fuel trap was integrated into the PMD to ensure propellant would be available during the back-away maneuver from the asteroid. The trap was sized to ensure the worst-case number and size of burns required to support TAG could be accommodated, with 2x margin above and beyond that.

4. DESIGN CHALLENGES

Unknown Asteroid Environment

The unexplored nature of Bennu presented some engineering design challenges. The instruments which will perform the early reconnaissance of the asteroid must handle the full range of surface brightness. For the cameras, a simple adjustment of the exposure times will compensate, although longer exposure times will drive the attitude control stability. And the mission requires enough dynamic range to support a reasonable range of surface brightness in a single image. The lidars, on the other hand, must have the proper gain control and protection to handle sudden reflective locations.

To contact the surface and collect the required sample, other characteristics of Bennu's surface come in to play. The grain size must be small enough to ingest into the sample head. The surface must be smooth enough to get proper contact. The system must be able to handle the range of possible surface strengths. And the planned trajectory must take into consideration variations in the gravity field as the flight system approaches the surface.

Early on, the project developed a "Design Reference Asteroid" (DRA) [4] based on the most up-to-date science data available. The science team made extensive observations of Bennu on its close approaches to the Earth in 2005 and 2011, and that data provided the basis for the DRA. The science team captured in the DRA all the parameters which were needed by the engineering team, with a best-estimate and an expected range for each. Examples of parameters include the orbital elements, size, the spin rate, photometric properties, thermal properties, and surface mechanical properties. Clearly, there is significant uncertainty in some of these parameters, but the infrared, radar, and optical observations of Bennu did provide some information. Also, other asteroids that have been visited, like Itokawa, provide some bounding of the problem. The engineering team recognized throughout development that the team could not place requirements on Bennu, but the engineers have placed requirements on the system to handle the expected range of environments at Bennu with margin. The collaboration between the science team and the engineering team in the development of the DRA enabled the project to clearly capture the knowns and unknowns, so the engineering team knew where best to apply the margin, maximizing the chance of success.

Significant Small Forces

Unlike a planetary orbital environment where the planet's gravity is by far the dominant force acting on the spacecraft, the "small forces" acting on the OSIRIS-REx spacecraft are of comparable magnitude to the asteroid's gravity, so much so that the accuracy with which the accelerations from these forces can be modeled determines how accurately the spacecraft's future orbital state can be predicted. Figure 7 shows the magnitude of these small forces relative to Bennu's gravity. In the Preliminary Survey Phase, the acceleration from the solar radiation pressure exceeds that due to Bennu's gravity, and during the Orbital A Phase the oblate shape of Bennu comes into play.
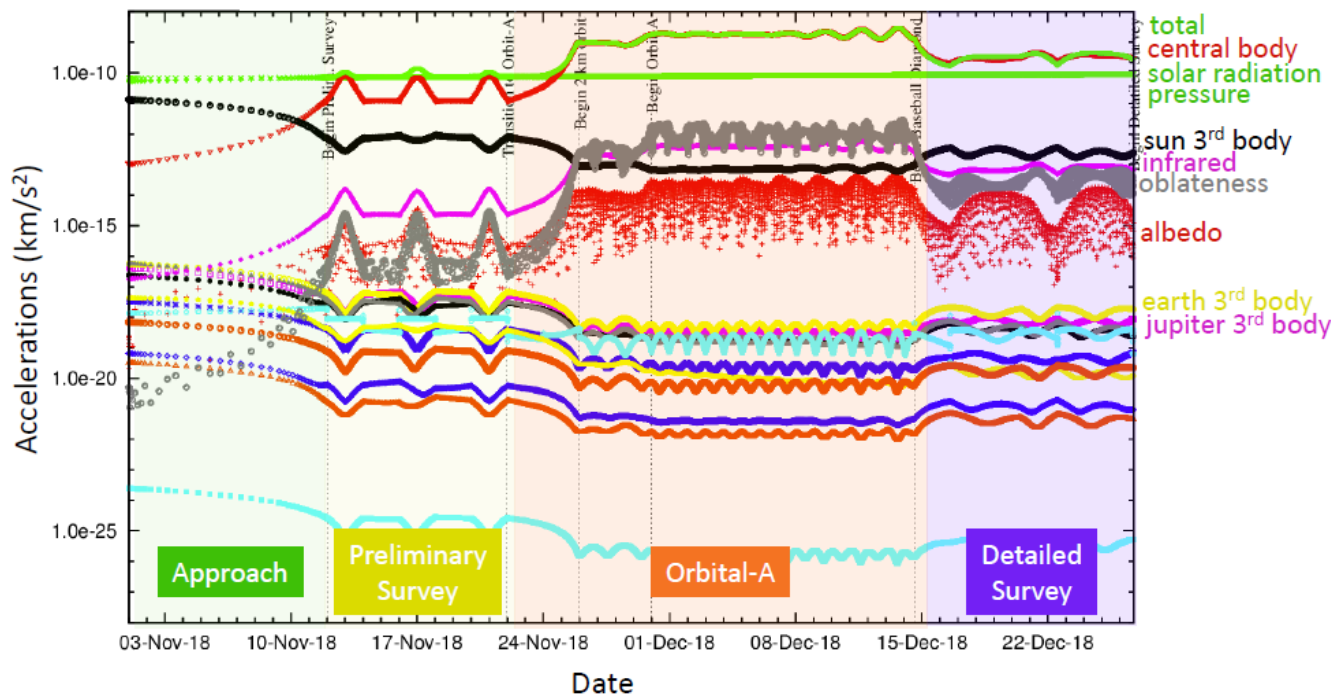
Figure 7: Small Force Accelerations Relative to Bennu's Gravity (central body)

Another small force acting to perturb the spacecraft's trajectory around Bennu is the residual delta-v from spacecraft thruster firings to unload system momentum. The magnitude and frequency of these momentum unload events must be carefully controlled, and the residual delta-v accurately modeled and factored into orbit determination solutions to maintain orbit and accurately predict the spacecraft's future state for the small delta-v maneuvers needed around Bennu for science observation targeting.

Autonomous "Docking" with a Natural Target

With a one-way light time of up to 18 minutes, spacecraft operations during the TAG execution needed to be completely autonomous. And with an expectation from the science side of being able to provide a safe landing ellipse of only 25 meters in diameter, significant effort early in the project design went into ensuring the spacecraft would remain safe throughout the TAG operation with no ground-in-the-loop interaction. As an example, one quick change was to articulate the arrays back away from the surface to add clearance, as the Spacecraft could tip by up to 45 degrees during ground interactions while collecting the sample. Another change removed large back-away thrusters that had risk of kicking up significant regolith from the surface, and instead using the smaller attitude control thrusters.

However, the bulk of the early focus was placed on how to get to the surface successfully. The baseline concept for performing TAG utilizes just a lidar system, and relies on an expected range and rate to the asteroid as the Flight System traverses across the surface. By starting ranging before ever even crossing the limb of the asteroid, the Flight System is able to detect the limb of the asteroid as a starting reference point. Then, using a pre-defined range/rate corridor, the Flight System is able to continue to track the progression over the surface until finally achieving a spot directly over the landing ellipse, where it then starts the slow decent to the surface. Analyses by the navigation and GN&C groups of all possible error sources was factored into the Range/Rate corridor checks to ensure no false trips, while guaranteeing spacecraft safety. The end result during early development was a very robust system that could detect and respond to a variety of faults, while ensuring the highest likelihood of collecting a pristine sample.

Unfortunately, delays in the development of the GN&C Lidar forced the team to look at alternative methods to safely get to the surface. After pulling together a trade, alternative options which included radar, other lidar sources, and even pure time-based maneuver solutions, a selection was made to go with NFT. The robust optical navigation system is able to use features on the surface of the asteroid (rocks, craters, etc) to determine position and make adjustments as needed. Similar to the lidar range/rate corridor solution, NFT required a new set of fault protection to monitor the health of the camera system, as well as track key parameters to ensure no spoofing of NFT that would result in a mission failure. Using simulated asteroid surface models, NFT has proven to be robust, as long as the asteroid surface has features that can be tracked.

Instrument Accommodation

In order to win a New Frontiers mission, the design must incorporate significant flight heritage to reduce the risk of development. Existing designs for instruments and the spacecraft must be adjusted somewhat to accommodate

difference from the past applications. The LM spacecraft was designed with this fact in mind. All of the instruments interface through a single circuit card design, with science data storage available right on the card. Any instrument-unique changes tend to be contained within this hardware and its associated software, leaving the rest of the spacecraft bus avionics relatively untouched. But there were some areas that were not fully envisioned by the original designers of the LM spacecraft architecture.

A topic that came up early and continued to follow the instruments was sun avoidance. The instruments do not include aperture covers, and some of the imagers can be damaged by relatively short exposure to the sun, so the spacecraft GN&C system, combined with ground planning of observations, must avoid staring directly at the sun. The system includes sun sensors facing in the direction of the instrument boresights to trigger a safing response should the spacecraft attitude deviate from plan. The safe mode includes a slow rotation to guarantee that the boresights will not dwell on the sun. Because the instruments are pointed along the primary thrust axis of the launch vehicle, the sun avoidance also put some constraints on the launch window. Future missions should better consider the implications of sun avoidance during phase A, when the structure and the mission architecture are still flexible.

Although the spacecraft bus has significant ability to autonomously detect and react to anomalies, fault protection for the instruments was largely assumed to be within the instruments themselves. This does reduce the interface burden, but it also results in significant duplication of effort across the instruments. Although thresholds can be easily modified late in development, the logic of safing checks is hard coded within the spacecraft software. If the logic were more easily configured, the spacecraft team could take on the instrument safing checks, reduce the cost of instrument software development, and improve the overall reliability and system-level visibility of instrument safing.

5. BUILD CHALLENGES

New Technology

The OSIRIS-REx flight system leverages heritage hardware and software, proven test campaigns and experienced personnel to ensure mission success. The OSIRIS-REx flight system builds on proven LM experience and uses established hardware, software, technology, and processes from Stardust, MAVEN, Juno, GRAIL, Odyssey, and MRO and lessons learned from twelve interplanetary missions launched by LM in the last seventeen years.

With that said, to meet the unique aspects of the OSIRIS-REx mission, there have been some flight system new technologies that did require extra oversight and proper mitigation of risk to get them fully developed and ready to support the OSIRIS-REx mission. As part of the overall OSIRIS-REx process, a full new-technology assessment was performed, to ensure proper steps were taken to implement these new technologies, as well as ensure the proper mitigation of risk.

One of the initial steps that system engineering did, with support/inputs from the subsystems, was the identification of new technologies. This evaluation and identification is a critical system engineering process, to ensure that those areas with build challenges, do have the necessary maturation plans in place, the necessary risks identified/tracked, and the necessary resources in place to support these new technology developments. To ensure the new technologies on the OSIRIS-REx program also got enough higher level oversight, an independent review team was organized to review those program "hot spots" where new technologies may exist, and to confirm the identification of a new technology. This new technology would then go through a formal process of risk burn-down, with maturation plans developed, along with a review of these plans with an executive level team review, to ensure the program has all of the plans/resources in place to execute the development and implementation of these new technologies. Figure 8 summarizes from where the key heritage components were leveraged (with New/MAVEN leverage shown as a separate block since it was the most recent interplanetary mission launched, prior to OSIRIS-REx launch).

The flight system new technologies identified for OSIRIS-REx were: GN&C Lidar, NFT, TAGSAM, and TAGCAMS. Here is a summary of each new technology and the challenges experienced during the developments:

GN&C Lidar—Advanced Scientific Corporation (ASC) provided a modified commercial lidar as OSIRIS-REx's RexEye space-rated GN&C lidar. The laser was defined by the program as new technology with TRL-6 achieved at Mission PDR in March 2013. OSIRIS-REx requires this GN&C lidar as the prime approach for successfully performing TAG event. OSIRIS-REx did a trade on other proximity sensors (range finder, scanning lidar, flash lidar) and the ASC lidar was deemed the lowest risk. The GNC lidar, at the highest level, supports not only providing range measurements during the TAG event itself to support the on-board navigation, but also during the reconnaissance passes, which are performed before TAG to characterize the TAG site, providing range measurements for PolyCam focusing. The transition from a short-duration, low-Earth commercial product to a long-duration planetary application demanded a more robust device and a laser life test. The total development associated with the lidar included the hardware development, qualification and delivery, the GNC algorithm development, fault protection updates, and thermal control and overall design considerations.

Figure 8: Heritage Assessment

The early development work on the laser module uncovered several issues. The first life-test unit failed early due to contamination that was traced to a heater. The hermetic package of the laser module was identified as a potential problem, but it passed its early testing, then later units showed excessive leaking. This problem resulted in a redesign of the laser module package.

Problems continued to plague the development as the complete lidar assembly came together. The first random vibration test caused foreign debris to show up on the inside of the laser window, pointing out the need for better contamination control during assembly of the lidar. The diffuser wheel position changed after testing, identifying an error in the mechanical design. And the shock testing showed an unexpected sensitivity. Overall, the problems were the type one would expect during the development of new flight hardware.

LM brought in many technical experts across the various technical fields from both LM and GSFC to help resolve these design challenges. This included electrical engineering, laser transmitter module experts, contamination control, optical, sensor module delamination experts, test subject matter experts and fellow, failure analysis labs and the large Space Operation Simulation Center support for end-to-end testing in a flight like environment.

Additionally, the project implemented development of NFT as a back-up navigation scheme for TAG. See below for more details on NFT.

Challenging developments provide many lessons. Here are the main lessons from the lidar development: (1) ensure a conforming prototype hardware test bed is available (and configuration controlled) throughout program; (2) tailor oversight methodology consistent with the vendor capabilities and limitations; (3) establish good control of hardware/software baselines, test procedure documentation, and GSE systems; (4) apply increased cost and schedule margin for new developments with small companies.

In the end, with strong system engineering oversight and strong LM/GSFC oversight of the GN&C lidar vendor, the team successfully delivered and integrated onto the flight system two fully verified flight GN&C lidars, with a fully qualified design.

Natural Feature Tracking (NFT)—LM developed and demonstrated NFT under Internal Research and Development (IRAD), with a PDR successfully completed in May 2013. OSIRIS-REx utilizes the redundant Malin NavCams (part of TAGCAMS) and LM IRAD based software to mitigate GN&C lidar risk. This NFT alternate technology was implemented as risk reduction to the GN&C lidar (risk of unexpected GNC lidar performance at Bennu). Since NFT

was not added until after mission CDR, the challenge was the integration and test of this optical image processing and software solution, and the associated fault protection, on schedule. The final design and implementation demonstrated a design that met full performance with margins in all key areas—FSW processing, memory, through-put and overall TAG footprint.

TAGSAM—OSIRIS-REx is flying TAGSAM, a single-plane motion arm, complete with the TAGSAM head that fluidizes/captures Bennu regolith. This TAGSAM design evolved over ~10 years, with hundreds of hours of testing completed (full TRL-6 completed at Mission PDR in March 2013. The TAG-based sampling approach selected for OSIRIS-REx is the result of over ten years of LM IRAD efforts. Inherent to this effort were exhaustive trade studies comparing landed and TAG strategies. Landing a S/C on a small body requires anchoring devices to stabilize the S/C, which is difficult because of the likelihood of low-strength regolith. Negating bounce upon landing requires hold-down thrusters, resulting in sample contamination. In addition, communication dropout is inevitable due to terrain masking, unfavorable attitude, and rapid day/night cycling. OSIRIS-REx eliminates these risks by using a short-duration contact TAG for sampling.

The team studied many different sample collection techniques, including sticky pads, claw and clamshell samplers, drive tubes, augers, coring drills, scoops, rakes, and gas-stimulated sampling. Gas stimulation is used because it minimizes moving parts, functions without motors during sampling, and keeps the sample pristine. The effectiveness and simplicity of the TAGSAM approach was initially verified through five years of intense development, including the production of five engineering development units, which have been used in an extensive test program.

OSIRIS-REx Sample Return Mission requires TAGSAM. GN2 bottles are utilized to fluidize regolith and TAGSAM head then captures regolith and traps it, utilizing a Mylar flap. LM Denver designed and manufactured TAGSAM and completed a full qualification program and flight acceptance program completed for TAGSAM flight head. Moreover, the approach to confirm the successful collection of the sample, once the S/C is safely away from Bennu, is achieved through a Cassini heritage algorithm for identifying S/C inertia /change in inertia, before and after sample collection.

TAGCAMS—Modified Malin Space Science Systems (MSSS) camera system utilizes hardware from MSSS modular imaging system product line. The TAGCAMS cameras use the C50 (5 Mpixel) camera body. TAGCAMS was at TRL-7 (System prototype demo in space environment) upon selection of this hardware for the OSIRIS-REx mission. OSIRIS-REx utilizes redundant MSSS NavCams (part of TAGCAMS) for both NFT and ground-based Optical Navigation inputs (the ground based navigation required wide FOV camera system as well). As a note, LM Space System Company has flown MSSS camera systems previously, but the challenge was OSIRIS-REx's critical uses

of this camera system for autonomous navigation. The heritage of this camera system includes a classified mission, the MSL rovers, and the Lunar Reconnaissance Orbiter.

In the end, by taking a very proactive approach at identifying the new technologies, and then developing plans to mature these new technologies into successful design and implementations, all of the new technologies on OSIRIS-REx were successfully designed, built, tested and integrated into the OSIRIS-REx flight system.

Heritage

As already outlined above, the OSIRIS-REx flight system builds on lessons learned from twelve interplanetary missions launched by LM in the last seventeen years. These lessons learned include:

- Perform a set of trades early looking at heritage that already exists to determine if it will meet the mission requirements; heritage is only value-added if it meets the requirements. However, look at opportunities where the mission design is flexible where requirements can be modified to fit within heritage capabilities. Often, requirements can be reworked that result in negligible impact to the L1 and L2 science requirements, while reducing cost and schedule risk by fitting within existing heritage designs and analyses. The OSIRIS-REx structure and the TAGSAM sample head were designed around the Stardust sample return capsule. The mission was designed around the heritage spacecraft capabilities.

- Stick with heritage that matches the mission. Trying to force fit heritage hardware causes cost overruns and increases risk. OSIRIS-REx made these decisions early.

- Scrutinize heritage hardware to guarantee it meets mission requirements, including environments (this includes full heritage reviews, prior to incorporation into the design). OSIRIS-REx held two weeks of heritage reviews in July 2012, shortly after the Mission Definition Review.

- Design for flight hardware testability. Building on planetary experience and incorporating recent lessons learned, the team has designed a systematic, thorough, and fully traceable component and system assembly, verification, and validation program.

- Provide robust power and telecom margins. Robust margins provide flexibility during operations.

With a Spacecraft system that is heavily tied to heritage architectures, there is added risk for any change. Full implications of change are not obvious to those who did not design the component. This was seen in a couple areas during development that led to some late discoveries during the build of the OSIRIS-REx Spacecraft. The following are some examples from the Spacecraft where heritage played a role, both good and bad.

Channelization—One of the more frustrating examples where heritage added risk is an error introduced into the channelization layout for the spacecraft that was not discovered until initial integration of components on the spacecraft. Following the heritage mantra, the harness subsystem used the same pin-outs for the avionics hardware as was used for previous builds, and was able to complete their design and drawing releases very early in the development phase for the program. However, due to some card modifications (see the "DTCI" section below), the Avionics team included a couple pin changes to increase robustness in the card. While the card changes were understood and fully vetted through appropriate Change Requests, with the harness drawings already released, the incorporation of the pin-out changes was not caught during harness manufacturing. During initial signal checkouts prior to mate of components, the discrepancy was discovered and corrected, but resulted in the loss of a couple days in ATLO to perform the modifications real-time. While use of heritage saved significant amount of time during the development phase, care must be given to ensure the proper reviews are still held later in the flow to ensure changes are fully captured and implemented correctly.

Power Modeling—Another area where heritage caused problems was in the modeling of the Power subsystem. For the heritage missions to OSIRIS-REx, the variation of incoming power is fairly stable for any given phase of the mission, thus the Power Model could be configured to simulate a single worst-case moment-in-time. For OSIRIS-REx however, the mission profile while at the asteroid span a solar range from 0.89AU to 1.4AU, with a wide variety of load cases occurring at different times. Taking the worst-case load for science observations planned to occur at 0.89AU would break the 1.4AU scenario if the heritage approach was used. Furthermore, huge changes in thermal states across the wide solar range change and various instrument configurations added to the challenge. To mitigate this, the model was updated to calculate the available power on a day-by-day basis throughout the mission. The initial attempts to use the heritage model and phase it to work for a 7.1 year period proved cumbersome, with the model taking multiple minutes to re-calculate when updates were required. In the end, new tools were created to optimize the model, thus getting calculation times down to under a minute per change. In this case, heritage was not helpful, and ended up creating extra work for the team compared to just starting from scratch at the beginning.

Data, Telemetry, and Command Interface (DTCI) Card—The C&DH DTCI is the payload interface card for the OSIRIS-REx spacecraft and heritage missions. The card contains the interfaces for not only the command and data interfaces to the instruments, but also the data storage capabilities, and provides some clear examples of the pitfalls associated with making "minor changes" to heritage hardware.

Where all the other C&DH cards interface with other heritage components (GN&C, Prop, Thermal, Telecom, etc), the DTCI card is used to communicate with instruments, which are typically unique for each mission. Thus, the design for the DTCI often needs to be tweaked to accommodate changes to payload interfaces or protocols. On OSIRIS-REx, the interfaces happened to align perfectly with MAVEN, and thus no changes looked to be required. The only open item from MAVEN involved some very minor memory corruption errors at extreme temperatures. These were considered a non-issue for the MAVEN program, but could cause problems on OSIRIS-REx due to operation within those temperature regimes. The short-burst nature of the science collection means that even minor memory corruptions will cause large perturbations to the data return. Thus, the decision was made early-on to fix the memory corruption errors, which MAVEN had tracked to requiring some simple rerouting of signals in the Printed Wiring Board.

Since the card was being modified for the memory corruption errors, the team decided to add robustness to the card pin-outs in the process. This was seen as a "free" change since it only involved moving around signals (no additions or subtractions), and ultimately led to the Channelization errors described above. Of bigger concern though, was that rerouting of the signals to fix the memory corruption did not actually fix the problem – indeed, the corruptions were worse after making the change. After several months of investigation, the issue was traced down to shifts in the timing circuit and drive current circuits that had been masked with the MAVEN implementation. The fix required minor modifications to all the DTCI boards, which were already installed on the Spacecraft and had gone through the initial dynamic environment testing at the System Level. The C&DHs had to be removed from the Spacecraft and penalty testing performed. The project experienced a significant amount of work, all for something that could have been resolved by optimizing the thermal design to stay out of the temperature regimes of concern, thus allowing the card to stay Build-to-Print from MAVEN.

Not that full Build-to-Print from MAVEN would have been without its own issues. One other late discovery with the DTCI card wasn't discovered until the first System Verification Test (SVT) on the Spacecraft running multiple instruments simultaneously. While the DTCI card was designed to handle simultaneous data transfers from instruments, the heritage card-level test program only simulated a single type of data interface with data always coming in 32-bit aligned. On OSIRIS-REx, multiple instruments transferred data without 32-bit buffering, thus requiring the DTCI card to handle this data buffering. Performing the buffering on one stream of data proved a non-issue, while multiple instruments proved problematic. Since the obvious solution would have required a FPGA change to the board late in the ATLO flow, the team sought alternatives. The final solution involved moving around instrument interfaces between the DTCI cards to ensure no simultaneous operations required 32-bit buffering. This required breaking

into the fully EMI-wrapped harness in the middle of system-level environments to make the change. The lesson here is that when using heritage components, the heritage test program should be scrubbed to ensure applicability to the program requirements.

Housekeeping Power Supply (HKPS) Card—The HKPS card is an example where heritage can be very helpful, even from unexpected sources. On MAVEN, the PDDU utilizes a High-Efficiency Power Supply (HEPS) board to provide secondary power for the box. The HEPS has heritage back to the 01' Mars Odyssey days, where power constraints required a power supply design with minimal internal losses to save every watt possible. After the development of the HEPS, future programs continued to use that design, as that became the new "heritage". The HEPS card was fantastic at what it was required to do, but also very expensive to manufacture.

Early during development on OSIRIS-REx, a cost opportunity was investigated to reduce the high manufacturing costs for the HEPS. Power was not a constraint on OSIRIS-REx, thus there was no need for a "high-efficiency" board. A trade was performed to look for other solutions, and the outcome pointed to resurrecting the HEPS pre-cursor, the HKPS board, which sacrificed efficiency for simplicity. The HKPS was a proven design, used on past programs including Stardust and Spitzer Space Telescope.

The original HKPS design required minimal modification for parts obsolescence, and was manufactured at a fraction of the cost for a HEPS. In addition, the card consumed less than 2W more compared to the HEPS due to increased efficiencies in the piece-parts. All-in-all, the HKPS has proven to be a very successful implementation of heritage, even from programs over 15 years ago.

Sample Return Capsule—While the HKPS proved to be a highly-successful implementation of old heritage, the SRC shows that extremely old heritage has its own level of pitfalls that need to be considered early in the development phase.

The SRC was originally designed for Stardust starting in 1991, with most of the design details completing in 1994 – still over 20 years prior to the OSIRIS-REx development. This alone proved the most significant challenge, for the following reasons:

- Microsoft is NOT backwards compatible – Going back one or two versions of Word and Excel may be achievable with Microsoft, but going back 20+ years poses a challenge. Special servers had to be setup that could run old 2000 versions of Microsoft just to allow import of heritage test reports, analyses, and data systems. Simple items like extracting original channelization for the SRC turned into tedious tasks of digging through files with name length limitations. And while hardcopies existed of all the data, finding exactly what you needed from 20+ year old archives is not a simple task.

- Hardcopies, while nice to retain, only tell part of the story – Stardust did a great job documenting everything, which was great… if you could find the data. Modern conveniences of being able to digitally search through a server for particular items is not possible when digging through boxes of drawings, even when scanned and converted to pdf. 2000+ pages of hardcopy documentation requires significant effort to sort through.

- Finding resources from Stardust days proved difficult – As could be expected, few of the original designers from the Stardust era are still around, and those that are have little recollection of the finer details of the design.

Beyond the general age of the Stardust SRC design, another factor that proved challenging was the change in engineering design and analysis expectations. Analyses and tests performed back in the "Better, Faster, Cheaper" days do not meet the rigor expected in today's low-risk aerospace industry. Significant effort was required to rework analyses and perform new testing that wasn't required back during the Stardust era, something that wasn't fully factored into the initial heritage assessment.

Note that the extra testing was not without value – two issues were uncovered during retest activities that may have been a problem for OSIRIS-REx if only the heritage test programs had been relied upon. The first revolved around the shock environment for the SRC Avionics Unit, specifically during the mortar firing to release the drogue parachute. On Stardust, the analysis showed that this was not a significant shock source due to the number of joints separating the parachute from the Avionics Unit. However, updated analyses and shock testing showed that the pressure wave from the mortar firing was able to directly interact with the SRC Avionics Deck, which produced a shock event 10x higher than expected. The Stardust shock isolation system was not sufficient to account for the increased shock, and needed to be enhanced. A new shock isolation system utilizing wire-rope units was developed to fit within the tight volume constraints of the SRC. The increased size of the shock isolation system required rebuild of the harnesses to account for the additional standoff heights, and special back shells utilized to ensure clearances with the heat shield hinge line.

The other issue uncovered in the additional testing involved the parachute. During a drop test, the Stardust heritage design was shown to have the cut-cord potentially bind in the bridle system, which would have resulted in failure to deploy the main parachute. An improved cut-cord system was implemented to ensure hang-ups of the cut-cord would be impossible, however, this introduced a single-point-failure if both primary and redundant cutters happened to execute simultaneously. Attempts to mitigate the SPF in the Avionics design proved problematic, with testing showing that there was a 1 in 20 chance of getting the cutters to release within the window of vulnerability. Based on the test data, the

parachute design was modified again to remove the SPF from the design completely.

Operations Processes—Another example of the adaptation of heritage for OSIRIS-REx is in the area of mission operations processes. Because of the challenges in predicting the spacecraft's future state due to the complex small force environment, spacecraft commands for activities like maneuvers and science observations that depend on those predictions must be updated shortly before execution to take advantage of the most recent navigation tracking information possible. In fact, a requirement was levied on the system to be capable of updating maneuver and observation parameters on the spacecraft within 24 hours of the activity execution. Heritage mostly-manual operations processes and software tools could not complete the tasks of preparing, reviewing, testing, and uplink of these "late updates" within the 24 hour window. This resulted in the automation of many of the heritage processes and tools used to develop and test late update command products. For example, the tool used to determine if any spacecraft pointing constraints are violated by a late-updated science observation design was modified to perform this checking autonomously.

Complex Subassemblies

While any complex system will experience its share of discoveries along the way, the TAGSAM and SRC subassemblies provide some interesting examples of note.

Although obvious after the fact, one key aspect of both the TAGSAM and SRC subassemblies is that, while mostly mechanical in nature, they contain electrical components. However, the responsible engineers for both subsystems were Mechanical Engineers by trade, and did not understand the full implications of some of their decisions on the electrical portion as the design matured. This caused disconnects between items like channelization and default sensor states that were not caught until very late in the development phase. One example of this was the wiring output to the phases of the motors, which were consistently backwards from how the avionics and FSW had been instantiated. Although documentation from both sides showed the design, because of the lack of electrical oversight, the significance was missed during data reviews. A similar issue arose with the wiring of the micro-switches, resulting in the state of the switches being ambiguous until the TAGSAM arm is deployed and the SRC is initially opened. In hindsight, any subassemblies that contain both mechanical and electrical components should have someone with more of a systems perspective leading the effort to ensure compatibility.

Another area where complexities of the subassembly caused consternation involved the pogo force during TAG. Constant force springs were sized based off the analysis to provide a soft push without introducing stresses in the shoulder and elbow joints of the arm. However, after initial testing of the actual force, it was discovered to be ~20% higher than expected. The cause was traced to the harness running down for the nitrogen bottles and wrist motor, which was stiffer than originally expected. Even though a model harness had been created to check stiffness for the model early in development, the flight harness contained extra shielding and tighter overwraps, resulting in the extra stiffness.

The final aspect of the TAGSAM design that proved challenging involved Electro-Static Discharge (ESD) mitigation. Analyses performed early in development showed that the asteroid could have a 100V differential compared to the Spacecraft at time of TAG. Modeling of the effects of this on the Spacecraft showed that this was a sufficient bump of the ground plane to potentially cause a reset of the processor, which would be catastrophic at time of contact. The mitigation was intended to be simple – isolate the TAGSAM arm from the Spacecraft with a bleed resistor to slowly discharge any voltage spike. However, due to the TAGSAM arm containing both conductive blankets and harnessing, the implementation proved harder than expected. The end result was to fully isolate the harness shields at both ends, as well as completely isolate the TAGSAM blankets from the Spacecraft at the shoulder interface.

Instrument Interfaces

Significant work was performed very early on the program to fully define and lockdown the instrument electrical and software interfaces to ensure full compatibility with the spacecraft avionics and flight software. Since many of the instruments were common with heritage programs built on the same architecture, little modification was required on either side of the interface to accommodate the vast array of payload designs. Indeed, due to the confidence in the early interface definition work, the program was able to realize early cost savings by deleting costly spacecraft simulation hardware, in lieu of delaying integration to the much higher fidelity Spacecraft Test Lab.

Where the electrical and software interfaces were worked heavily, the mechanical interfaces followed a more traditional pace, and as such, did experience some late discoveries on the program. Specifically, instruments that contained multiple components spread about the spacecraft with interconnecting harnesses held the largest room for error. With instrument harnesses manufactured by the individual payload teams, errors arose between the mock-ups and the assumptions of the spacecraft configuration at the various centers. Of particular note was a late discovery that multiple instruments, as well as the spacecraft, made an assumption that no other users required access through a particular pass-through in the decking. Each party had confirmed that their harness bundle easily fit, but during integration, it was quickly discovered that not all the connectors could fit through the hole as it filled with bundles. This forced increasing the hole's diameter at time of harness installation on the spacecraft to accommodate all the users.

Significant effort would have been prevented with better modeling of the harness routing from all parties, including proper harness diameters after all shielding is included. In addition, ensuring that each build center was constantly

aware of changes, no matter how minor, would have avoided late discoveries at time of integration.

One other instrument interface issue that showed up late in the build of the flight system was sensitivity of OTES to many different sources of mechanical noise. The team captured a requirement on the mechanical sensitivity at a particular frequency, but the instrument turned out to be sensitive to other sources at other frequencies, including vibrations produced the Miniature Inertial Measurement Unit (MIMU) and by whistling in the purge system. The MIMU interaction was only found during concurrent operations of OTES with the MIMU several months after integration. Future missions should be sure to operate all spacecraft components with the instruments shortly after they are integrated. A slight adjustment of the OTES mirror drive table took the MIMU vibration out of its band of operation.

Mass and Center-of-Mass (CM) Tracking

While OSIRIS-REx had very healthy mass margins throughout the program, CM knowledge on the mission was critical as determination of the sample mass was accomplished via detecting the change in inertia during a spin of the Spacecraft. Having a CM offset from the center of the fuel tank could lead to fuel slosh, which would be disastrous when trying to detect as low as 60g of sample.

Due to the high-heritage nature of the Spacecraft, most masses were well understood for the project. However, some key mistakes were made in the design that resulted in needing significant ballasting to correct. The first error was in the modeling and tracking of the harness mass. Heritage values for the final flight harness had always come in within 5% of the model predictions, thus lower-than-normal contingencies were carried in the model. However, when the OSIRIS-REx harness completed construction and was delivered, the weight was almost 30% over the model predictions – a significant mass increase. This was traced to a cost-savings opportunity that had been exercised to use wire from a common buy with another mission. The wire contained extra shielding, resulting in the increased mass. To further exasperate the problem, the extra shielding produced a thicker harness bundle, requiring more overwrap, resulting in even more mass gain, as well as having to reroute in some locations due to space constraints. This included needing to build harness ramps over propulsion lines due to insufficient clearances under the tubing as was originally proposed. In addition, the harness mass is not evenly distributed about the Spacecraft; approximately 75% of the harness mass is near the C&DH and PDDU, however the component placements to balance the CM had been done with the significantly lighter harness. Even after moving a battery to help offset the mass, ballasting was required on the order of 2.5% of the total Spacecraft mass to bring the CM back within requirements. While OSIRIS-REx had sufficient mass margins to accommodate this, this would have been disastrous for other missions.

Figure 9 shows an image of the OSIRIS-REx harness bundle which caused the bulk of the mass and CM issues that arose on the project.
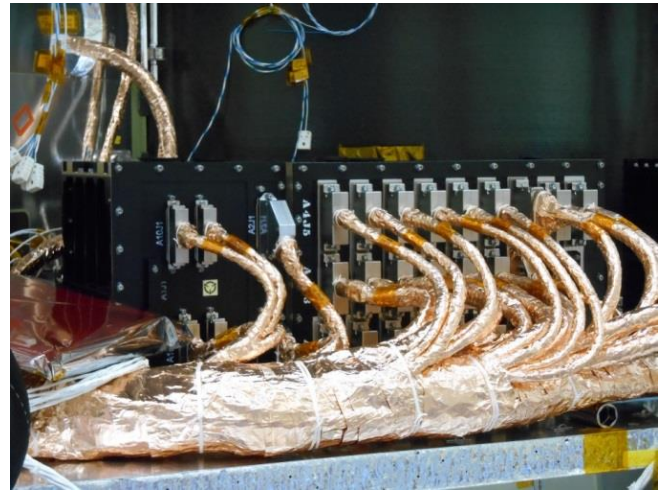

Figure 9: Harness Bundle

Performance Trending

The main performance measurements tracked and reported on a monthly basis throughout the development of the project were dry mass and power. Both trends followed heritage profiles. For mass, the trend continued to increase slowly, with big discrete changes at major deliveries. The trend of mass throughout the project is shown in Figure 10, with the main mass drivers denoted.
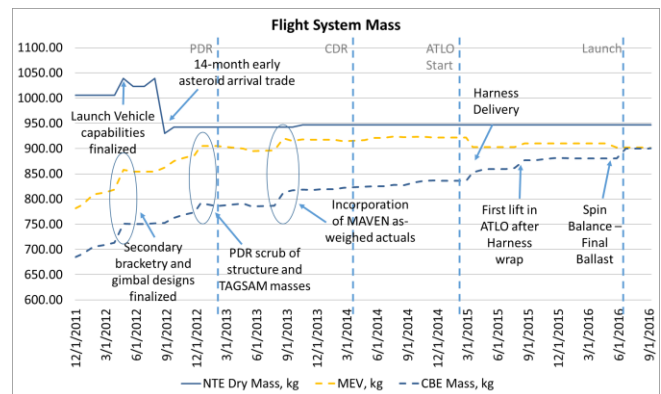

Figure 10: Flight System Mass Trend

Prior to PDR, the Not To Exceed (NTE) mass shifted up as the team established initial requirements for the launch vehicle procurement. The team then performed a trade and decided to provide additional margin for science by arriving at Bennu 14-months earlier than originally planned, thus resulting in a decrease in allowable dry mass. In addition, minor adjustments to the trajectory resulted in delta V budget changes, which impacted the allowable dry mass for the fixed wet mass.

As the Spacecraft, TAGSAM, and payload designs matured

prior to PDR, discrete jumps in the Current Best Estimate (CBE) and Maximum Expected Value (MEV) of mass for subsystems can be seen, with the largest jumps attributed to better definition of secondary bracketry requirements and gimbal sizing. In parallel with the early OSIRIS-REx development, MAVEN was into their Flight System integration, and incorporation of their as-measured masses for common components with OSIRIS-REx resulted in another discrete CBE mass increase between PDR and CDR. Finally, after the start of ATLO for OSIRIS-REx, the heavier harness mass was discovered as noted above.

For power, the trend is very flat throughout the project. Due to the high-heritage nature, most components had very solid power consumption numbers available. Because the project had healthy power margins throughout development, the team was able to ease mission operations by allowing the arrays to stay in different fixed orientations throughout the flight, rather than continuously tracking, thereby avoiding implementation of a sun tracking algorithm.

The biggest variables for power were: available power from the arrays, and average heater load. As the Solar Array model continued to be refined, some gains in available power were made, but the main change is decrease in load following system-level thermal vacuum where a lot of the conservatism in the thermal model is able to be removed. Power modeling has historically held significant margin. This is due to the use of using worst-case predictions for available power, as well as worst-case predictions for heater load. While more complex methods could be implemented to Root-Sum-Square these model inputs, traditionally, as was the case on OSIRIS-REx, it is easier to use the stacked-worst-case assumptions during development. Figure 11 shows the trend in available power versus load throughout the development of the project.
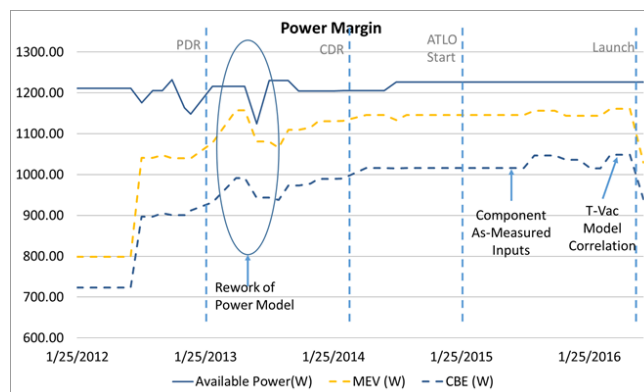

Figure 11: Flight System Power Trend

Mission Plan

On the operations side, two key challenges for implementing the OSIRIS-REx mission plan included 1) the tight coupling between ground elements, particularly in executing the late update process described above, and 2) the operational need for and use of science data products. As an example of the tight coupling between ground elements, a late update to science observation pointing parameters requires activities by all ground elements during the 24-hour late update period. After the DSN downlinks the latest optical navigation images, FDS must reconstruct the trajectory and predict the state of the spacecraft at the time of the observation. The new predictions are delivered to the SPOC to determine if the observation pointing design needs to be updated to accommodate changes in the spacecraft's predicted state, and if so, modify the design. The updated observation pointing design is then delivered to the MSA to check against spacecraft constraints, prepare and verify the updated products, and send these to the DSN to uplink them to the spacecraft. This need for the late update capability drove thorough definition of the interfaces between elements and significant testing of those interfaces and processes, testing which is ongoing and will continue until the spacecraft reaches Bennu in August of 2018.

The second operations challenge involves having science in the operations loop. The primary modus operandi for science in planetary missions is to receive data from the science instruments and process and analyze that data over relatively long time scales to support scientific investigations and publish results. For OSIRIS-REx, products developed by mission scientists using tools and techniques developed for scientific inquiry are needed to navigate and guide the spacecraft and to select a suitable sampling site. This requires the SPOC and science team to process instrument data on much tighter timelines than science operations personnel are accustomed to. As an example, significant investment was made to develop the software, processes, interfaces, and team to use imagery from OSIRIS-REx's suite of cameras to produce detailed shape and rotation state models of Bennu, along with accurately located and modeled surface "landmarks" for use by the FDS and on-board NFT software to navigate the spacecraft around Bennu and ultimately guide it to the surface to collect the sample. To demonstrate this capability and all of the elements involved, a series of tests was performed in which simulated Approach and Preliminary Survey Phase imagery generated by the FDS team was used by the science team to construct shape and rotation state models, and landmarks, which were then passed back to the FDS team and used in a detailed simulation of landmark-based navigation during the Orbital A Phase. Further testing of the shape product generation and utilization processes by both FDS and NFT are planned during Outbound Cruise.

6. RISK MANAGEMENT

The biggest challenge for a planetary mission is staying on schedule. Orbital dynamics of the solar system does not wait for the system to be ready. If the project can stay on schedule and hit the first launch period, it will most likely stay in the cost box. The problem for the systems engineer is maintaining the technical integrity of the mission despite strong management pressure to stay on schedule. If the project is constantly reacting to issues, the chances are high that some fixes will be poorly thought out, merely addressing

the symptom without robustly correcting the cause. The faster a project must respond to a problem, the more likely the fix will jeopardize mission success. Proper risk management, where potential problems are identified early and mitigated, is essential to mission success.

OSIRIS-REx developed a culture of risk identification and mitigation early on in the design phase. The formal risk management process was structured to reinforce the culture, enabling the project to identify many potential problems and head them off before they became major issues. Over 200 risks were formally tracked at the project level, with only about 30 becoming problems, usually with substantially mitigated impact. There were very few completely unexpected problems that hit the project during development, and the team had the time to fix those few problems without consequence to the mission objectives or reliability.

The OSIRIS-REx risk management process was built on monthly meetings between the project systems engineer (PSE) and each element lead, during which they would update the status of the current risks and talk about any other concerns that might need tracking. These informal meetings provided an opportunity for the PSE to talk with each lead one-on-one, gaining insight into the complexities of each element. The environment fostered a frank discussion of concerns and risks without jumping into solutions or political posturing. The project risk manager facilitated the discussion and logged the updates in the risk database. After the one-on-one meetings, a monthly risk management board meeting, chaired by the project manager, would gather all the leads, project management, the Principal Investigator, and the PSE together to go through the top risks and decide priorities for mitigations.

OSIRIS-REx's under-budget and on-time performance is a direct result of the project's active risk management. Project management used the risk management process to determine spending priorities, applying resources to mitigations with the greatest value. The dollars spent on mitigating risks resulted in a net savings over time, allowing the project to spend additional dollars on additional risks. The project had a healthy management reserve when it was selected, and management used that reserve effectively to *prevent* problems, rather than holding the money and spending it when problems manifested.

7. REQUIREMENTS MANAGEMENT AND VERIFICATION

Early in the design phase, the lead scientists embraced the requirements development and worked hard with the systems engineering team to define clear, verifiable requirements. The entire development effort revolved around the mission requirements defined in phase A, with very little change throughout the mission. This stability was a key factor in achieving on-time delivery under cost.

As the project proceeded into phase B, the individual instrument teams, the spacecraft team, and the ground elements each developed their element-level requirements which they linked back to the mission level. At the project systems level, the verification engineer looked at the flow-down from mission-level to element-level for every single mission requirement. This was an extensive effort to validate the element-level requirements, with engineering judgment applied, and it uncovered a number of misunderstandings and disconnects across the team. Detailed design proceeded in parallel with this effort, so there were times where the team needed to make some design changes, but the effort was early enough that the major disconnects were captured long before they significantly impacted cost or schedule.

As soon as the flow-down validation was complete, the team turned its focus on the detailed plans for requirements verification. With over 4000 requirements tracked at the project level, the work ahead needed to be carefully planned. The verification lead worked with each of the elements to establish appropriate tests and analyses. They agreed on appropriate verification evidence up front. And the verification lead established a clear process with a second set of expert eyes on every piece of evidence. No requirement was formally closed until all the evidence was reviewed and in place under configuration control. The team closed an average of 250 requirements a month over a 15-month period leading up to launch.

At a high level, the requirements management and verification effort was very much like a text-book example. But the devil in the details that is hard to covey is the importance of the knowledge and perspective that the verification engineer brings to the task. Our verification engineer took to heart the fact that a mistake in the element requirements development early on will lead to large cost and schedule impacts later, and a mistake in requirements verification at the end can result in loss of mission. He applied engineering judgment to all his work, constantly asking, "Does this make sense?" He set the tone for the team, and brought in the appropriate experts to review the work. During system-level integration, the project had very few problems because the right verification work was done early. In flight so far, the system is performing very well.

More details on the OSIRIS-REx requirements verification effort can be found in Stevens [5].

8. VALIDATION

First, a few definitions, to help differentiate verification from validation: Verification is the process of determining that a system satisfies its requirements and specifications; it answers the question "did we build the system right?" Validation is the process of determining that a proposed system solves the problem for which it is intended; it answers the question "did we build the right system?"

For the OSIRIS-REx validation, the project took a two phased approach. The first phase of requirements validation precedes the design solution and shows that a design based on the requirements will be the right system design prior to

anything being built from the design. This has largely been accomplished as part of the requirements walk-through / tabletop / peer review process. The second phase of validation is concerned with showing the designed and tested system meets the intended function originally specified. This largely consists of putting the flight system through targeted testing with most, but not all of the validation activity being performed at the flight system level.

For the flight-system validation level testing, the team created a representative set of operational scenarios against which the OSIRIS-REx flight system was validated. These scenarios were derived from the DRM. Level 3 spacecraft requirements were then linked to the individual scenarios where they were validated by the scenarios. The validation program exercised and replicated, to the extent feasible, a representative set of nominal and off-nominal / failed states and modes of operation.

The SVT's used in the flight system validation were:

1) Launch
2) Launch – Off Nominal
3) Outbound Cruise
4) Science Reconnaissance
5) Science Survey
6) Science TAG
7) Science TAG– Off Nominal
8) Science NFT TAG
9) Science NFT TAG – Off Nominal
10) Earth Return (and SRC Deploy)
11) Earth Return – Off Nominal

These tests directly correspond to specific phases of the DRM.

For OSIRIS-REx, many of the specific validation events occur during the SVT's on the spacecraft, during the Operational Readiness Tests (ORT's), and during the Mission Readiness Tests (MRT's). The SVT's used flight-like, MSA / Ground-Data-System-generated sequences in tests that were run on the flight system, using realistic (nominal) scenarios. The project ground-system organization managed the ORT's. Launch ORT was conducted prior to launch to validate proper readiness for the event. The ORTs utilized test beds in lieu of the spacecraft for post-launch activities. Finally, the two MRTs validate the SPOC's ability to develop instrument commands using the operations tools, and then this test was performed with the entire OSIRIS-REx mission operations system in the loop.

Additionally, the OSIRIS-REx project conducted a Risk Reduction Test (RRT) program, beginning toward the end of the flight software system integration process. Risk reduction testing was almost entirely performed on SoftSim test beds (no hardware in the loop). RRT scenarios were based on Fault Tree branches and some test cases took the form of multiple-fault "stress tests." In the end, the RRT program was used to ensure the flight system behaved as expected to support the OSIRIS-REx mission.

Finally, in terms of validation closeout, every validation event was individually assigned to a responsible system engineer. In addition, the validation procedures include the details for close-out, including which organization is responsible for analyzing the results (i.e., subsystem, fault protection, etc.), and which individuals must sign off on the declaration of a validation activity's success or failure.

Validation event was considered closed when the following occurred:

1. Data review has been completed and signed-off.

2. All steps in the test procedure were followed, or written rationale provided for why any steps were skipped

3. All applicable configurations in the Mission Plan are included in the test setup

4. Successful execution of command sequence(s) and scripts has been verified. This includes the absence of rejected commands, proper operation of the involved sequence engine(s), proper command execution times, correct command counters throughout the test.

5. Correct fault protection execution has been verified. This includes any expected fault protection execution, as well as the lack of unexpected execution (i.e., correct fault counters throughout the test).

6. Availability of adequate telemetry data both in real-time and in recorded / playback form has been verified (i.e., that we have access to the data we need in order to judge whether things worked). Also, any telemetry anomalies are fully understood or documented in writing.

7. Any relevant problem reports opened during test execution that are test liens have been closed.

8. Compliance with flight rules, constraints, and mission rules has been verified. Any idiosyncrasies must be explained in writing, as needed.

9. OBSERVATIONS AND CONCLUSION

OSIRIS-REx has been a huge development success. The flight system launched on schedule and under budget without compromising on the reliability or performance of the system. The operations phase is well planned out, with significant funding available. The efforts of over 1500 individuals came together to build a system that is well-placed to achieve all objectives, amazing the world with views of an object that has never been seen, teaching us something about predicting asteroid trajectories and how we might alter them to protect the Earth, and returning material that will be analyzed for decades.

During the five-and-a-half-year development effort, the team overcame numerous obstacles. But the effort was made much easier by a culture geared toward success. The Principal Investigator established a pattern of open communication early on. The Lockheed Martin team took a close look at the heritage hardware and software, and openly presented its weaknesses in front of their customer. When problems arose, team members were not reluctant to ask for help. Overall, the team looked forward toward mission success, anticipating problems and tackling them before they became insurmountable.

The project promoted a culture of listening to team members and reviewers. Whether or not the project changed course due to another perspective, the person who spoke up knew that the team had considered the idea. The project took all actions from reviewers very seriously, formally driving them to closure. The utilization of diverse perspectives helped the team identify risks long before they became problems.

OSIRIS-REx has demonstrated the positive effects of good risk management. The keys to this process were (1) good communication of risk from the lower levels through a culture of risk identification and mitigation; and (2) active use of the risk management information by project management when deciding where to spend reserves.

Heritage components provided significant cost savings and development risk reduction, but there are some clear cautions when using heritage. Heritage components and the *testing of those components* must be thoroughly reviewed for the new application. Old heritage designs will require significant effort to recover the documentation and some design information will be unavailable, so plan a thorough test program with plenty of extra schedule. Don't forget about the ripple effects across the system of any changes to heritage components.

The team stumbled across some previous lessons learned, too. It is difficult to not get complacent with heritage systems. And there will always be the unexpected problems. In general, though, the project kept its focus. A solid verification and validation program caught numerous errors before they became significant problems.

Designing to explore and sample a largely unknown object was made easier with a strong partnership between science and engineering. The Design Reference Asteroid codified the asteroid environment, and the Mission Requirements Document and Design Reference Mission provided clear direction to the team. That partnership will continue through operations as the project collects data and learns about Bennu.

OSIRIS-REx still has 7 years left before the Sample Return Capsule brings its extraterrestrial cargo safely to Earth. We cannot claim mission success. However we have reached a major milestone, and the flight system hardware cannot be changed. We have all learned a bit more about building space flight systems, and hopefully that knowledge will be of benefit to others.

REFERENCES

[1] Beshore, Edward, Brian Sutter, Ronald Mink, Dante Lauretta, Michael Moreau, William Boynton, Jason Dworkin, David Everett, Christopher Shinohara, Jonathan Gal-Edd, "The OSIRIS-REx Asteroid Sample Return Mission," *Proceedings of the IEEE Aserospace Conference*, Big Sky, MT, March 2015.

[2] Lorenz, David, Ryan Olds, Alex May, Courtney Mario, Mark Perry, "Lessons Learned from OSIRIS-REx Autonomous Navigation Using Natural Feature Tracking," *Proceedings of the IEEE Aerospace Conference*, Big Sky, MT, March 2017

[3] Ajluni, Thomas, Timothy Linn, William Willcockson, David Everett, Ronald Mink, and Joshua Wood, "OSIRIS-REx, Returning the Asteroid Sample," *Proceedings of the IEEE Aerospace Conference*, Big Sky, MT, March 2015.

[4] Hergenrother, Carl W., et. al., "The Design Reference Asteroid for the OSIRIS-REx Mission Target (101955) Bennu," arXiv:1409.4704.

[5] Stevens, Craig, Angela Adams, Bradley Williams, Colby Goodloe, "Cleared for Launch - Lessons Learned from the OSIRIS-REx System Requirements Verification Program," *Proceedings of the IEEE Aerospace Conference*, Big Sky, MT, March 2017.

BIOGRAPHY

David Everett is currently the Project Systems Engineer for the OSIRIS-REx. In his 25 years at NASA, he has led the design, build, and launch of four spacecraft, and he was a key player during the launch of three others. Before OSIRIS-REx, Mr. Everett led the technical effort for the Lunar Reconnaissance Orbiter as the Mission Systems Engineer, from design through early-orbit operations. Mr. Everett has actively supported NASA outreach activities through over 75 speaking engagements. He has received 37 individual awards, 25 group awards, and a patent for his efforts at NASA; he has published 18 papers; and he co-edited (and wrote the spacecraft design chapter for) the book Space Mission Engineering: The New SMAD. *He earned a BSEE summa cum laude, at Virginia Tech in 1986 and a MSEE at*

the University of Maryland in 1989. Before he joined NASA in 1991, Mr. Everett worked at Westinghouse Electric Corporation where he was awarded two patents for his designs of RF circuits.

Ronald Mink is a Senior Systems Engineer at NASA's Goddard Space Flight Center and served as the Deputy Project Systems Engineer for OSIRIS-REx during its development. His primary responsibility was as the "mission architect", formulating and maturing the Design Reference Mission and applying it to the design, development, and validation of the mission's system elements. Mr. Mink has been with NASA for over 15 years and continues to serve OSIRIS-REx as the Mission Systems Engineer during mission operations.

Timothy Linn is the Senior Manager of Space and Missile GN&C (Guidance, Navigation and Control) within Software Engineering at Lockheed Martin Space Systems Company. He has over 21 years of experience in Aerospace GN&C and System engineering. Previously Mr. Linn was the Chief Systems Engineer for the flight system of the NASA OSIRIS-REx New Frontiers 3 Asteroid Sample Return Mission. He worked on OSIRIS-REx program from the initial proposal phase through the launch of this spacecraft. Prior to OSIRIS-REx, Mr. Linn worked on numerous interplanetary spacecraft designs. He was the GN&C Subsystem Lead on NASA's GRAIL Discovery Mission and NASA's Phoenix Mars Lander. Mr. Linn started his career with Lockheed Martin as a GN&C analyst, working multiple interplanetary missions (Stardust, Genesis, MRO), in addition to working in Defense Systems/Special Programs. Mr. Linn graduated from the University of Colorado, Boulder, with his Bachelors and Masters Degrees in Aerospace Engineering and earned his MBA from Regis University.

Joshua Wood was the System Design Lead and System Architect for the OSIRIS-REx Flight System. In his 16 years at Lockheed Martin, Josh has worked on over 8 NASA missions, including Mars Odyssey, Genesis, Mars Reconnaissance Orbiter, Spitzer Space Telescope, Phoenix Mars Lander and GRAIL Lunar Gravity Mapper, covering a variety of roles from development, integration, test, and operations. Josh graduated in 1999 with a BS in Engineering with a Mechanical Specialty from Colorado School of Mines.