SSC24-III-10

NASA PRISM's Lunar Vertex Mission – Lessons Learned in Establishing a New Low-Cost Science Mission Paradigm

Ann L. Cox, William F. Ames
Johns Hopkins University/Applied Physics Laboratory (JHU/APL)
1110 Johns Hopkins Rd Laurel MD; 240-228-7583
ann.cox@jhuapl.edu

ABSTRACT

Lunar Vertex (LVx), a lander-rover science investigation of Reiner Gamma, was selected as NASA's first PRISM (Payloads and Research Investigations on the Surface of the Moon) mission in June 2021. The PRISM program is under NASA's Planetary Missions Program Office (PMPO). To facilitate quick delivery at low cost, NASA is managing PRISM missions to the requirements for research and technology projects as documented in NASA Procedural Requirements document NPR7120.8. It is JHU/APL's first science mission developed under the NPR 7120.8 requirements. NASA and JHU/APL worked together to accept a higher risk posture, to allow for less stringent requirements, significantly reduced documentation, and reduced oversight and review than required for flight projects managed to NASA Procedural Requirements document NPR 7120.5. This program structure facilitates use of commercial products, providing a pathway for new suppliers to participate in NASA scientific investigations. Lunar Vertex (LVx) is a pathfinder to establishing a new low-cost science mission paradigm and has been a learning experience for everyone involved. Lessons learned have been identified and recommendations are provided.

LUNAR VERTEX MISSION OVERVIEW

NASA designated Reiner Gamma as the first Payloads and Research Investigations on the Surface of the Moon (PRISM-1a) destination. In June 2021, NASA selected the Lunar Vertex payload suite for the PRISM-1a mission. Lunar Vertex will investigate the origin of lunar magnetic anomalies; the origin of lunar swirls; and the structure of the mini-magnetosphere that forms over the magnetic anomaly.

APL is responsible for science, project management, systems engineering, safety and mission assurance, the payload, and science operations of the Lunar Vertex mission. APL also provided two magnetometer instruments and performed integration and test of the rover.

The Lunar Vertex payload suite is comprised of three instruments to be hosted on the lander and a rover carrying two instruments. PRISM payloads will be delivered to the lunar surface by commercial landers as part of NASA's Commercial Lunar Payload Services (CLPS) program.

The Lunar Vertex lander instruments, shown in Figure 1 through Figure 3 are:

- The Vertex Camera Array (VCA) provided by Redwire Aerospace
- The Magnetic Anomaly Plasma Spectrometer (MAPS) provided by the Southwest Research Institute
- The Vector Magnetometer – Lander provided by Applied Physics Laboratory

Figure 1: VCA Instrument

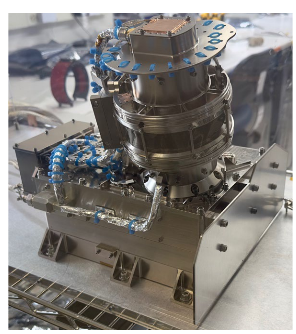

Figure 2: MAPS Instrument

Figure 3: VML Instrument

The lander will deploy the Lunar Vertex rover, carrying a Rover Multispectral Microscope provided by Canadensys Aerospace and the second APL magnetometer, Vector Magnetometer – Rover (Figure 4). The roving vehicle is provided by Lunar Outpost and is based on their Mobile Autonomous Prospecting Platform (MAPP) product line. This was designated as MAPP-C, the C defining it as the rover chassis. The Lunar Vertex "rover" is the integrated MAPP-C vehicle, RMM instrument, and VMR instrument. The Lunar Vertex rover is shown in Figure 5.

Lunar Vertex modified Redwire Aerospace (VCA), Canadensys (RMM), Bartington (magnetometer sensors), and Lunar Outpost (MAPP) commercial products to provide the capabilities needed to deliver the required science data and to survive the environments. Modifying existing commercial products allowed Lunar Vertex to develop a meaningful, groundbreaking science investigation for a fraction of the typical costs.

The Lunar Vertex three lander instruments were delivered in June 2023, less than 2 years after project start. The instrumented rover was delivered 2 ½ years after project start. Launch and completion of operations is currently planned for 2025.

LESSONS LEARNED THEMES

At the time of the Systems Integration Reviews/Acceptance Reviews an initial set of Lessons Learned were solicited from all team members and their institutions as well as APL's Chief S&MA, Chief Engineer Office, and Program Management Office. More than 100 individual lessons learned were submitted.

This paper describes some of the most beneficial lessons learned collected under the following themes:

- Implementation as a Research & Technology Mission
- Project Management
- Communication
- Technical Requirements and Verification
- Interface Management
- Integration and Test
- Safety and Quality Assurance
- Risk Posture, Rating Risks, and Descopes

IMPLEMENTATION AS A RESEARCH & TECHNOLOGY MISSION

PRISM missions are governed by NPR 7120.8 NASA Program and Project Requirements for Research & Technology. The requirements for research and technology missions reflect acceptance of higher risk than on flight projects. There is less documentation, deliverables, project reviews, and oversight than

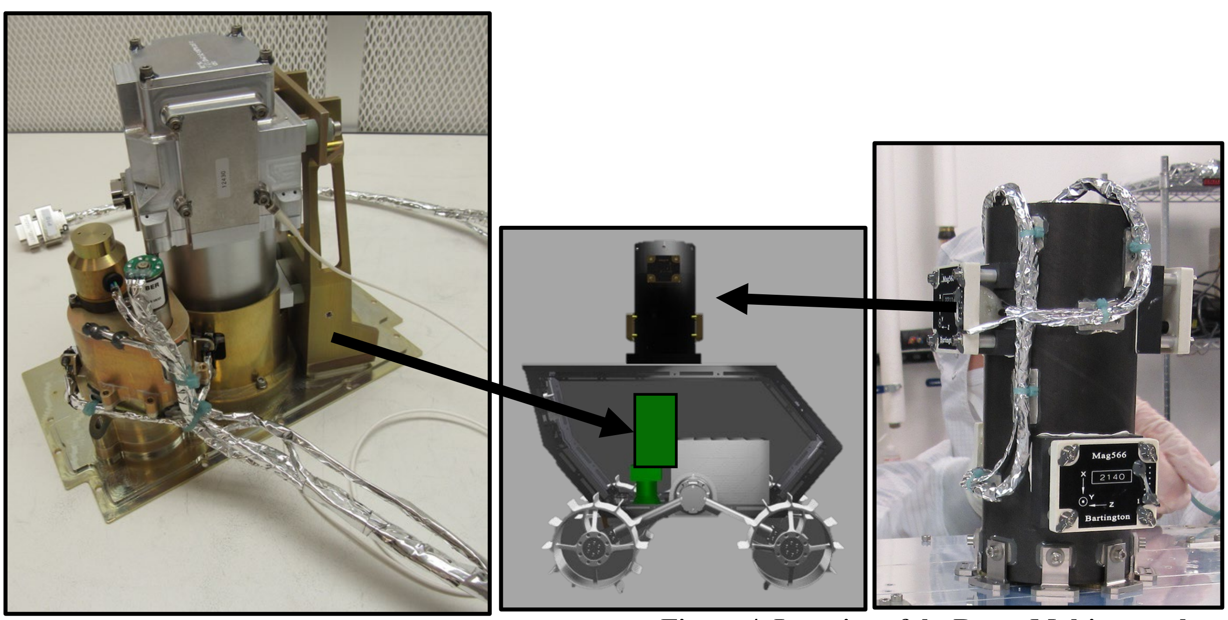

Figure 4: Location of the Rover Multispectral Microscope and Vector Magnetometer – Rover instruments on the MAPP-C vehicle.

required for flight projects managed by NPR 7120.5 NASA Program and Project Requirements for Flight Projects. However, NPR 7120.8 does not contain the same level of specificity found in NPR 7120.5. Each program office determines the details for their research and technology missions. As the first PRISM program mission, NASA and Lunar Vertex worked together to generate the statement of work, document deliverables, NASA-led project reviews, risk management, and oversight.

There is a single Key Decision Point (KDP) and only 3-4 NASA-led project reviews. Rather than a Standing Review Board (SRB), an Independent Assessment Team, comprised of less than 10 people, conduct these reviews. The foundation for the streamlined PRISM requirements was twofold: to Do No Harm (DNH) to the lander and its payloads and to demonstrate performance in the expected environments.

People who are external to the program office and project team and not familiar with the PRISM requirements frequently reverted to the requirements for flight projects as defined in NPR 7120.5. There was a resistance to deviate from business as usual which created a disconnect with expectations. Our NASA Mission Manager and Technical Authority were consultants to the IAT. They provided information and insight regarding risk posture, requirements and expectations to the IAT. Interactions between the Mission Manager and Technical Authority with the review team during NASA-led project reviews was very beneficial.

Lunar Vertex was APL's first science mission implemented using NPR 7120.8 Imposing standard practices, developed for flight projects, has significant schedule and cost impacts. Lunar Vertex needed to educate and guide their institutions on the differences in executing a project per NPR 7120.8 rather than NPR 7120.5 throughout the project life-cycle. Implementing institutions need to adapt and tailor their standard practices, for flight projects, to successfully execute projects governed by NPR 7120.8.

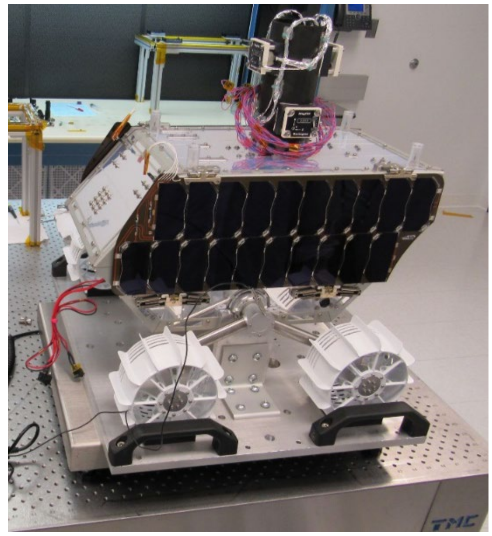

Figure 5: Lunar Vertex rover

PROJECT MANAGEMENT

PRISM investigations are not dedicated missions. They are one of a number of payloads that the CLPS lander is delivering to the lunar surface. Our contract is to deliver Lunar Vertex to the lander provider's integration and test facility by a specific date. Meeting the delivery date is paramount and is independent of the launch date. Failure to meet the contractual delivery date puts a payload at risk to be de-manifested. The delivery schedule drove many technical decisions. This schedule pressure required a "good enough" approach, expedient decision-making, acceptance of workarounds rather than fixes, and implementation of descopes.

Payloads should plan for a storage period in the event of a launch date slip. Our schedule drove many technical decisions, including descopes. A detailed, integrated project schedule must be developed and maintained to reflect realistic schedule expectations and performance. Subcontractors may not have sufficient schedule tools and experience in which case, the project scheduler should provide them with schedule support. The Lunar Vertex scheduler attended the weekly technical meetings to obtain the current technical status ensuring the integrated master schedule reflected the actual performance to date. Integrating the scheduler with the technical team led to open communication and honest reporting, which reduced our schedule risk.

For projects with tight schedules, subcontracts need to be in place quickly. Payloads should generate baseline subcontract documents based on the proposal to get partners started immediately. Waiting for development and maturation of project requirements is not recommended, they can be captured later with subcontract modifications. Projects should consider early delivery incentives clauses in their subcontracts.

Cost constraints limit team size. Payloads should assemble a small team of experienced staff that have the capabilities to perform multiple roles. Payload teams should then identify areas where subject matter experts (optical engineering, magnetic cleanliness, structural analysis, etc) are needed to supplement the core team and get commitment for their support, and allocate budget to cover their specific tasks. Given Lunar Vertex's limited budget, and fast-moving schedule, the financial manager generated a detailed weekly cost analysis and met with the project manager to review and determine any needed corrective action(s). This level of monitoring and adjusting helped Lunar Vertex to stay on plan. It also resulted in maintaining a realistic, current Estimate at Completion (EAC).

COMMUNICATION

The vocabulary used across organizations may vary making it critical to clearly communicate needs and expectations such that all parties are aligned. In conversations and correspondence with individuals outside of their organization, teams should use clear and distinct communication. Our team was comprised of institutions and personnel with experience in a number of different industries and a wide range of customers. Terminology amongst the team was not uniform, which led to misinterpretation of the conversation or document. Use off acronyms and abbreviations, heavily used on space flight missions, required a learning curve to those unfamiliar with them.

The team, and their institutions, needed to keep an open mind and be receptive to doing things differently. Schedule and budget constraints required taking a "good enough" approach to meeting requirements, as opposed to achieving the best performance. It necessitated a collaborative approach to merge commercial practices with processes developed for flight projects.

Specific content of deliverable documents must be defined upfront so that the needed information is obtained. Commercial suppliers are reluctant to share proprietary/ competition sensitive information. Procurement of commercial off-the-shelf and product line items do not typically get delivered with additional technical data. This limits the information for verification of requirements as well as debug and root cause analysis of anomalies.

Project information and data must be readily available on platforms accessible to the entire team. This is essential to ensure the team is on-the-same page with respect to the current project details. It prevents spending resources performing tasks that used incorrect parameters.

TECHNICAL REQUIREMENTS AND VERIFICATION

Lunar Vertex requirements were structured with the same considerations as most spaceflight missions, with the particular challenge of delivering requirements quickly for a large number of subsystems relative to the size of the systems engineering team. Overall requirements structure was a simplified version of a standard NASA science mission, with Payload-level requirements defining the measurements required to meet mission objectives and instrument/rover level requirements defining the subsystem performance required to meet the mission's measurement objectives. Instrument and rover requirements were in turn flowed down to their own subsystems according to the engineering practices of the particular institution responsible for the instrument. Both baseline and threshold requirements were defined with their own flowdowns. Baseline requirements defined the performance required to meet nominal mission objectives, while threshold requirements defined the minimum level of performance for which the mission was worth flying. This overall structure worked well, but there were compromises and difficulties in implementation.

The scale of the requirements development effort was underappreciated at program start, leaving the systems team with a dilemma – spend time producing high quality requirements for all subsystems, at the cost of slowing development, or produce rougher requirements as quickly as possible? We opted for the latter, which provided subsystems with design goals early but at the cost of leaving several difficult requirements definition problems lingering for an extended period. To mitigate this tension between quality and speed, programs should put a higher priority on requirements development in the late stages of proposal writing, rather than leaving the effort until formal program start. While some instrument performance requirements were provided from the proposal team, these were based on "wish lists" or the performance of instruments on other missions which were not comparable to Lunar Vertex. Directly levying these requirements on instrument vendors would have resulted in costs and schedules much too large for our mission. At program start we reworked the requirements to flow from what was required to meet mission objectives, but found the instrument science and systems engineering teams were not adequately staffed to perform these analyses quickly. For future programs, staff generating mission requirements should be trained in requirements writing and the particular needs of low cost/low schedule missions, and adequate staff support for requirements development should also be provided immediately at program start.

In order to provide maximum design flexibility for instrument teams and minimize the procedural burdens of requirements updating and verification, the number of requirements imposed on any subsystem was kept low and restricted in scope to the final performance of the instrument/rover while minimizing design assumptions. While this was successful in allowing subsystem teams to freely update their designs as needed, there were unanticipated costs. As an example, the Rover Multispectral Microscope (RMM) instrument consisted of a microscope assembly and an illuminator board. Multispectral functionality was provided by changing the wavelength of the illuminated pairs of LEDs. Because the instrument contained both the camera and light source, specific camera performance requirements could not be levied without assuming a particular light source performance, and vice versa. The systems team wished to maintain the ability for the instrument team to be able to trade, for example, illuminator brightness against effective quantum efficiency, in accordance with the systems engineering principle that requirements should not specify design choices. To do this, the instrument requirements defined a nominal scene representative of the anticipated lunar environment and specified performance metrics such as signal-to-noise ratio to be achieved in various measurement modalities in the target scene. The cost was a less straightforward and specific requirements set and more complicated verification. As the camera portion of the RMM instrument was small modifications to an existing commercial-off-the-shelf (COTS) product, the benefit to maintain design flexibility in the camera was low. A more straightforward but less flexible set of instrument requirements would have simplified development and verification for the instrument team. Future programs should consider maintaining the separate Payload and Instrument levels of requirements, but in cases where a major subsystem of an instrument is a COTS product, assume the performance of that subsystem will be fixed and specify the rest of the instrument around that subsystem. While "bad practice" in the sense of assuming a solution, it allows both requirements and instrument architecture to close quickly with straightforward verifications, and the cost of lost flexibility is minimal, given the COTS nature of the product. Should redesign of a COTS subsystem be required, the higher level (Payload) requirements provide the basis for a rewriting of the instrument-level requirements.

In highly cost and schedule constrained missions, scope management is a vital tool. Initial formulation of baseline and threshold mission requirements resulted in only small differences between baseline and threshold requirements sets. This resulted in few options for using reductions in scope to correct budget and schedule problems. When threshold requirements were renegotiated in response to integration and test difficulties, this required a large reworking of requirements and verification while integration and test was ongoing. To accomplish this without driving schedule, the mission requirements philosophy was changed from a flow-down from science objectives to a "flow-up" based on expected performance of the finished product and analyses from the mission science and systems engineering teams which showed that the mission would answer key science questions. Verification then focused on functional, rather than performance, testing, due to time constraints both on integration and test itself and the time available to write new performance requirements at the subsystem level. Future missions could reduce the effort required in requirements definition and verification planning late in the program without compromising on rigorous verification by early definition of mission threshold requirements which provide ample opportunities for descopes. This will provide more flexibility in descope options and allow the engineering team to prepare for reductions in scope before they are needed. Functional, rather than performance, verification should also be considered for systems which have a very high probability of exceeding mission performance requirements should their basic functionality be verified. This streamlines verification testing with minimal risk.

INTERFACE MANAGEMENT

The Lunar Vertex mission includes three instrument systems mounted to the lander and two instruments mounted to the rover. There was also a large schedule offset between the lander delivery and the Lunar Vertex delivery – lander Critical Design Review was roughly contemporaneous with Lunar Vertex delivery. This presented some notable interface management challenges at all stages of the project.

The importance of keeping interfaces as simple as possible and complex only as necessary is well-understood, but it is important to evaluate interface complexity for more than technical challenge. In particular, we discovered that interfaces which cut across three different institutions are much more difficult to manage smoothly than interfaces which cut across two. A three-party interface introduces new trade-offs which are not present in two party or within-institution interfaces. An ideal interface management approach requires interfaces to be matured sufficiently quickly to support the development of subsystems on both sides of the interface, is responsive to the needs of all parties, and is clearly documented and understood by all parties. With three-party interfaces, either interface definition only occurs with all three parties present, limiting speed and responsivity, or pairs of institutions can meet as-needed to resolve issues, which is a threat to universal understanding and documentation. For this reason, three-party interfaces should be avoided, even at the cost of somewhat higher technical complexity. For example, the VCA instrument required a tilted surface to be constructed to adjust camera pointing such that the field of view came sufficiently close to the lander. Neither the instrument nor the lander provider was interested in adding such a mounting surface to their scope of work. In the resulting trade study, the APL Lunar Vertex team decided that the additional complexity of adding an APL-built "mounting wedge" in between the VCA cameras and the lander was low enough to justify the cost savings of building the mounting wedges at APL. This had several consequences. By inserting APL hardware in between the instrument and lander provided hardware, any proposed changes to the instrument mechanical interface would have to propagate through the interface wedge design before the lander team received a fully-specified proposal, and vice-versa. This slowed the update process and roughly doubled the opportunities for mistakes and miscommunications. Additionally, the "simple" interface wedge became more complicated as the instrument design matured. The volume of the wedge was needed to support supplemental instrument hardware and the wedge integration process was complicated by changes in the camera body coating selection. The mounting wedge's secondary role as a form of volume growth contingency was not appreciated during the initial trade study, and keeping the wedge design and build within the instrument provider's scope of work would have simplified both instrument design and integration. In the end, Lunar Vertex transitioned to alternative mounting wedges provided by the lander team due to the challenges created by the three-way interface, demonstrating the costs of under-rating the additional complexity of adding a third institution to an interface.

Another contributor to the interface challenges was the difficulties created by the large offset in schedule between the Lunar Vertex payload and the lander. The initial plan was for instruments to provide NASA with fully or nearly-fully specified interfaces to include in the Lander Request for Proposal. As the Lunar Vertex subsystem interfaces were not as mature as initially expected, substantial updates to the interface specification were required over the course of instrument development, updates which the lander team was not yet prepared to fully evaluate. By the time the lander provider was prepared to support final interface definition, the instrument subsystems were in final integration and test or fully delivered. Some lander-proposed interface updates were not presented until after instrument providers had delivered their hardware and their period of performance ended, making proper evaluation difficult. For future missions, interface management can be dramatically simplified by keeping the payload and lander schedules sufficiently similar that the payloads still have some design flexibility and subcontractor support at the point when the lander is finalizing its interface definitions. If this is not possible, payloads should be selected to be as close to off-the-shelf as feasible in order to reduce risk.

Despite the challenges posed by the complexity of the Lunar Vertex interface problems, these issues were solved through active interface management with small variations from standard systems engineering practices. The primary tool for interface resolution was semi-monthly meetings between the lander and Lunar Vertex systems engineering teams with interfaces recorded in Payload Integration Plans, or PIPs. Routine issues were resolved in this meeting with each side's relevant design engineering consulted independently as needed. For more complex or urgent issues, relevant experts from both sides of the interface, including from subcontractors, joined the meeting to work through the issue in real time, with follow-ups as needed. In a few cases, technical experts from both sides of an interface were put directly in touch with each other and defined the interface independently. This worked well for disciplines such as software which were isolated from other design elements, but in other cases it led to confusion when subject matter experts agreed to a solution verbally or recorded the interface in a separate document. The most notable deviation from standard systems engineering practices was in documentation approach. The Payload Integration Plans were not subject to a standard review/approval control approach, and on occasion redundant interface definition was described in alternative documents. While the intention of this approach was to reduce the communications overhead in interface definition, on more than one occasion avoidable errors or miscommunication took place due to confusion in where controlling interface information was recorded or how it was approved. While review/release cycles may appear time consuming, future programs should nonetheless consider conforming to standard practice in this regard. The effort required in reviewing and approving documentation may be thought of like insurance payments – it may appear an unnecessary expense, but it is an expense that can be anticipated and planned for, and the alternative is to risk large expenses unpredictably occurring at inconvenient times.

INTEGRATION AND TEST

PRISM payloads are required to complete all acceptance testing, including environments, prior to delivery. The only lander system level tests are an aliveness test and functional test of the lander-to-payload interface.

The Lunar Vertex mission initially took a traditional approach to development, integration, and test, with engineering models (EMs) planned to be developed and tested, followed by flight model (FM) assembly, integration and test. The two exceptions were the MAPS instrument and the rover. The MAPS instrument built an EM only of the instrument's re-closable door for technology maturation. This decision was made in light of the maturity of the instrument and Southwest Research Institute's extensive experience with delivering scientific instruments. As for the rover, the flight model MAPP-C was to be delivered to APL, where the instruments were to be integrated and tested and the final integrated rover delivered to the lander provider.

Due to schedule constraints and challenges encountered by each of the subsystem teams during integration and test, many subsystems saw substantial deviations from these initial plans in order to save schedule with minimal cost in risk or science return.

To enable on-time delivery of the MAPS instrument, the re-closeable door was descoped in favor of a remove-before-flight cover, which removed long-lead parts from the instrument bill of materials. The door was intended to protect the MAPS instrument from contamination impingement on the sensor head during flight. Analyses from the lander and MAPS instrument science teams indicated the risk of damage to the sensor head in flight was low so long as sufficient time to allow dust to settle between rover drive-off and MAPS high voltage bring-up were allowed, which the mission operations team was able to accommodate and meet mission science return requirements. Additionally, calibration of the MAPS electron sensor was descoped, reducing the scientific return of the instrument, though without threatening threshold requirements. The electron sensor was still installed and will be operated in flight, but with increased measurement uncertainty. Through these measures, the MAPS instrument was delivered on-schedule with a minor increase in risk and reduction in scientific return, consistent with the research and technology (NPR 7120.8) nature of the Lunar Vertex project.

Similar efforts were required to achieve an on-time delivery of the RMM instrument to rover I&T. The first decision was to procure a COTS optical assembly in parallel with the long-lead custom-designed lens assembly to reduce schedule risk. Subsequent testing showed that, while the COTS optical assembly had lower performance in some wavelengths than the custom assembly, it met performance requirements and issues with environmental qualification were highly unlikely. In the end, the custom assembly did not arrive in time, and the COTS assembly was used. The second schedule compromise made was to prioritize the flight model (FM) build over the engineering model (EM). Rather than the engineering model being delivered first and the flight model build starting afterwards, two identical flight-grade RMM instruments were built in parallel, and the first to be completed would be shipped to APL as the flight model, and the second would become the engineering model. While the engineering model would thus have limited value for informing the flight build, it would still be used in ground testing and for in-flight support of the RMM instrument. The resulting risk was accepted as the RMM instrument team had a thorough prototype-level test campaign with minimal design changes between the prototype and flight models. The final major compromise was waiving environmental acceptance testing of the flight model RMM at the instrument level. Flight-model RMM environmental testing was performed at the integrated rover level after integration of the RMM instrument into the MAPP-C. Environmental testing of the EM unit separately was performed in parallel. The combination of all of these measures was an on-time delivery of the RMM flight instrument which passed all environmental and performance tests.

The MAPP-C underwent similar efforts to ensure an on-time delivery as the RMM instrument, with environmental testing being waived until integrated rover testing and the FM rover delivered in advance of the EM rover. Some software development was also deferred, being completed during the rover integration process. However, the MAPP-C is a much more complex system which had not undergone the same level of prototype testing as the instruments, resulting in a less seamless integration and test program for the rover than for the instruments. Successful delivery of the rover required an integrated effort across Lunar Outpost, APL, the lander provider, the instrument teams, and several MAPP-C subsystem vendors. The primary coordination tools were a daily integration and test meeting between the APL and Lunar Outpost integration and test teams, with LO providing a combination of on-site and virtual support, and a continuous Zoom call between the APL and LO teams for real-time task coordination. Outside of these working meetings, Lunar Outpost coordinated anomaly resolution efforts with its subcontractors, while APL coordinated with instrument teams, the lander provider, and NASA, with both sides providing updates at the daily integration and test meeting. This separation kept each meeting's attendance small enough to be productive and allowed key information to flow where it was needed while respecting proprietary data concerns. Due to budget and schedule constraints, not every anomaly could be fully resolved. This required frequent coordination with the science and mission operations teams, as well as NASA, to prioritize efforts and ensure the final product was the best possible fit for the mission's science objectives.

The varying experiences with aggressive schedule reduction approaches, such as building engineering and flight models in parallel and deferring environmental testing to the next level of assembly, indicates that future programs should consider these options, but do so with caution. For systems where prototypes have been thoroughly tested and few changes are planned in the flight model these approaches should be strongly considered, but for more complex or less mature systems the risks are much higher. The mix of on-site and virtual support, with continuous virtual support and as-needed physical presence of varying subject matter experts (SMEs) depending on the task, worked well and should be repeated. Telemetry logging software was a vital tool for rover integration and test, and in cases where different software tasks need to be prioritized, logging telemetry and making it available to all collaborators as quickly as possible should be a high priority. If a decision is made to take delivery of a key subsystem before all nominal acceptance testing is completed, a set of pre-delivery requirements should still be negotiated and verifications performed to ensure all parties have the same expectations upon delivery, even if the performance of the final product is not expected to change. Finally, the joint Lunar Outpost/APL leadership, with Lunar Outpost managing their suppliers and APL communicating with the lander and NASA teams effectively balanced speed of information transfer, meeting sizes, and proprietary data concerns. Similar processes should be considered on future programs, but tight integration of the two leading teams and careful attention to maintaining a close working relationship is required.

Equally important to performance and environmental verification was interface testing with the lander, verifying that Lunar Vertex subsystems will interface with the lander successfully from lander integration and test through the duration of the mission. The primary vehicle for lander interface testing was the Lander Virtual Payload Adapter (VPA), which allowed Lunar Vertex hardware to connect to prototype lander software hosted on Amazon Web Services. Due to delays in the delivery of the VPAs, the first VPA testing took place on the FM instruments at APL, predominantly performed in parallel with rover integration and test. Rover VPA testing was incorporated into the final stages of rover integration and test, and at the time of writing MAPS VPA testing is ongoing. Limited interface testing with the lander flatsat was also undertaken at the lander integration and test facility upon delivery, with additional flatsat testing planned in advance of lander integration and test. While software interface testing was actively undertaken interfaces matured rapidly, but the need for SMEs from the lander team, the Lunar Vertex ground software team, and the instrument teams to simultaneously support VPA testing limited periods when testing could be performed. While the VPA is an excellent tool for interface verification, earlier delivery is recommended in order to ensure adequate time for interface definition. Additionally, successful interface testing was enabled by early definition of software interfaces – software interface control documents (ICDs) were drawn up months before VPA testing started, enabled by most Lunar Vertex subsystems either having quite simple software interfaces, heritage from a mature system, or benefiting from work on one of the lander's earlier planned launches. In programs with less mature software, earlier delivery of VPAs is recommended to enable initial testing before payload development is complete, ideally during payload engineering model testing.

SAFETY & QUALITY ASSURANCE

Range safety requirements, levied by the range safety office, are applicable to all spaceflight systems and need to be incorporated in the project safety requirements. They need to be met and verified to ensure that safety approvals can be granted to allow entry, processing, and launch from a range. These same safety requirements need to be met during manufacture, assembly, integration and test activities to safeguard personnel, hardware, and facilities.

Using commercial products may engage companies that have not previous built hardware for space applications. Since their standard processes and procedures may not have been developed using NASA standards, a review of their quality documents and facility walk-through should be conducted. Based on the review, there may be quality and workmanship requirements that should be levied in the subcontract. As an experienced space flight hardware provider, APL supported our partners in understanding the basis for the requirement(s) and how to implement processes and procedures to meet these quality and workmanship standards.

Payloads should consider sending the project quality assurance staff to perform inspections prior to closeout of critical hardware. This should be done in cases where the subcontractor does not have their own independent quality assurance team.

RISK POSTURE, RATING RISKS, AND DESCOPES

Being managed as a R&T mission allowed for a higher risk tolerance than flight missions. The acceptable level of risk needs to be defined in the proposal and agreed to with both the program office and the implementing institution before project start. The implementing organization may have a different risk posture than the program office.

NASA conducted a kick-off meeting (KOM) with the Lunar Vertex team. At the KOM, NASA presented their risk posture consistent with the requirements and statement of work. At our Lunar Vertex team KOM, this risk approach was communicated to the team. However, APL institutional requirements had not been factored into the Lunar Vertex project requirements, plans, cost or schedule. Implementing these requirements after project start had significant impact. Institutional requirements need to be addressed at the proposal stage.

Program specific risk rating definitions need to accompany any review of the risks. They need to be clear, concise, unambiguous, and consistent throughout the project. Rationale provided for the consequence and likelihood ratings should justify the rating based on the rating definition.

NASA's program office developed a specific set of risk rating definitions for the PRISM projects to reflect acceptance of higher risk. Risks often fell into lower categories, High (Red) Moderate (Yellow) Low (Green) than they would have if we used risk ratings typical for flight projects. This led to internal and external reviewers recommending higher ratings to align with their flight experience of what is a High/Red risk or Moderate/Yellow risk.

Once a risk is realized, or an issue is identified, the descope list can provide a means to get back to plan. PRISM projects should be prepared to take descopes that may impact achieving baseline performance to avoid being de-manifested, so an abundant descope list is essential. It should be broad reaching and cover not only hardware but tasks, tests, deliverables, and travel. For each descope the impacts to science, technical resources (mass, power, data), schedule, and cost need to be quantified. Descopes that impact the Threshold mission should be segregated and only proposed as a last resort. In support of descope options, an incompressible test list (ITL) should be defined early in the project. The descope list was reviewed and revised as-necessary monthly. Lunar Vertex took numerous descopes throughout the project. We conveyed the recovery options and need for a particular descope to NASA. They supported our efforts to manage to the available resources by making difficult decisions.

CONCLUSION

As the first PRISM mission, Lunar Vertex was the pathfinder for this new NASA program. The Lunar Vertex project and NASA PRISM program office worked well together in developing the initial requirements and in incorporating changes as-needed. Partnering an experienced space flight hardware provider with commercial suppliers forged a team that collaboratively developed a lunar mission for a fraction of the typical flight project cost. The lessons learned will be reviewed and selected recommendations will be incorporated in future solicitations. Lunar Vertex and the NASA PRISM program have established a path to perform meaningful science investigations for a fraction of the typical costs by engaging commercial partners, implementing a "good enough" approach, and accepting more risk.