AAS 26-657

MISSION DESIGN CONSIDERATIONS FOR A LOW-THRUST SPACECRAFT

Scott N. Karn*

Developing an executable low-thrust trajectory for use in a spaceflight mission requires the design and optimization of a deterministic trajectory as well as the validation that the selected architecture is robust to some set of uncertainties, execution errors, and potential contingencies. Uncertainty in the ability of the spacecraft and its launch vehicle to execute a trajectory as well as potential deviations such as in-flight anomalies combine with design-to constraints and requirements to complicate the optimization problem. The approach by which robust mission design was accomplished for the initial capability of NASA's Gateway is presented as well as associated results.

BACKGROUND

Over the past 30 years, solar electric propulsion (SEP) has grown from an experimental technology serving in handful of relatively niche applications to a workhorse widely adopted across the industry[1]. Since the early 2000s SEP thruster and spacecraft power generation technologies have matured substantially, allowing SEP to move out of the realm of bespoke, cutting-edge spacecraft such as Deep Space 1 and Dawn and into use on proliferated Low Earth Orbit (LEO) constellations, geostationary (GEO) communications satellites, and exploration platforms. In particular, the maturation of hall effect thrusters (HETs) in the last three decades[2] has led to their first use in deep space[3] as well as the development of high-power systems for use in human exploration[4]. This rise in the use of SEP-equipped spacecraft has also resulted in increased interest in the formulation and development of new missions as well as significant advances in the field of low-thrust trajectory design and optimization.

Despite increasing interest and research, low-thrust trajectory design remains an underdeveloped field when compared to its high-thrust counterpart. This is particularly true in the area of robust mission design, where a wide gap exists between methods which can be used to design trajectories that are mathematically valid and those that are executable by a physical, operational spacecraft. While the design and optimization of deterministic low-thrust trajectories is often computationally difficult and resource intensive, robust trajectory design increases this difficulty significantly as mission performance must be assessed and validated against expected execution errors, performance uncertainties, and risks. Beyond-Earth missions are particularly sensitive to these sources of error as they transit often-chaotic multibody dynamical environments, but the rarity and

*Mission Design Engineer, Mission Architecture and Analysis Branch, NASA Glenn Research Center, 21000 Brookpark Road, Cleveland, OH, 44135

bespoke nature of these missions means that the tools used to develop robust, flight-ready trajectories are often built from scratch or involve significant overhauls of heritage methods.

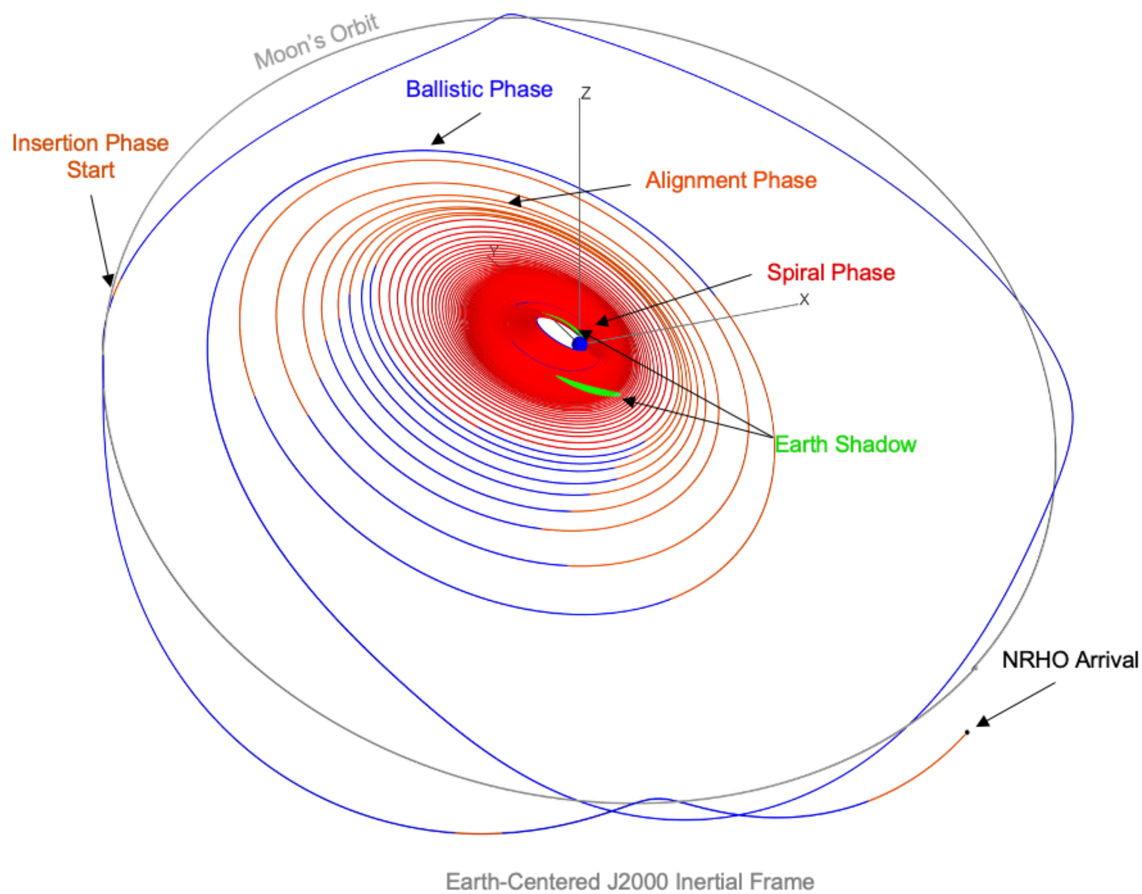

Figure 1: Lunar Transit mission diagram

The initial capability of NASA's Gateway lunar space station was envisioned to be one such mission, wherein a 50kW SEP-equipped Power and Propulsion Element (PPE) would transfer a large, human-exploration spacecraft from a Geostationary Transfer Orbit (GTO) to a 9:2 Earth-Moon Near Rectilinear Halo Orbit (NRHO)[4]. The maturation of this low-thrust Lunar Transit trajectory (Figure 1) from early, feasibility-focused, design reference missions (DRMs) towards full-fidelity, executable flight trajectories necessitated the development of novel methods and tools for robust mission design. Many of these methods leveraged heritage approaches to low-thrust mission design authored by the Deep Space 1 and Dawn mission teams[5]. The assessment of navigation and cycle-to-cycle statistical uncertainties evolved from these methods and has been described previously[6]. This publication will focus on how uncertainty in the performance of the integrated spacecraft, including launch vehicle insertion errors, SEP thruster and power system performance, as well as contingency scenarios and off-nominal events such as missed thrust events (MTEs) were assessed.

INTRODUCTION

In a conceptual or low-fidelity space, the objective of trajectory design is often simply to find a solution which successfully transits between two locations of interest (the Earth and the Moon, for example). For flight missions, this already challenging optimization problem becomes significantly more complex due to the presence of physical systems which must execute a mission as well as the programmatic context in which a mission exists. Mission objectives, requirements, constraints, and "desirements" or "nice-to-haves" must all be balanced against one another. A solution space which before may only have been bounded by physics and a handful of critical choices (spacecraft mass, propulsion system performance, etc) now becomes heavily laden with myriad inputs and constraints (many of which are often competing or contradictory). In this way, an already-challenging optimization problem becomes significantly more daunting as a mission develops and matures.

Adding every possible argument directly into the optimization of a trajectory would rapidly result in a model that is too dense and too complex to be useful (or to even converge in the first place). Many details are thus considered a-priori or a-posteriori in order to simplify computation; lower-fidelity methods can be used to filter out some set of initial conditions prior to simulation, trajectory results may be post-processed to determine compliance with a known set of requirements, and statistical methods such as Monte Carlo analyses can be used to assess trajectory performance under known uncertainties and risks. This, however, requires an extensive set of tools and methods to accomplish. And while varied and complex, these analyses must be reliable and scalable in order to support the cadence associated with mission development.

This a-posteriori approach to uncertainty analysis proved to be a substantial undertaking for the PPE and the Lunar Transit mission. Leveraging previous work to develop a robust framework by which deterministic nominal reference trajectories could be developed[7] and a similarly standardized methodology for the assessment of MTEs[8], an expanded analysis suite was developed to enable large scale assessments of integrated performance uncertainties in critical mission parameters. This package, named Aberrant and built around the existing Copernicus trajectory design software[9], enabled large scale Monte Carlo analysis of launch vehicle insertion errors, uncertainties in SEP thruster and spacecraft power system performance, and off-nominal events such as thruster failures as well as including the existing MTE framework. Results from large-scale runs could be quickly assessed for compliance with existing mission requirements and constraints including operating limits of critical limits and data provided critical insight into the necessary propellant margins and reserves to ensure successful execution of the Lunar Transit mission.

MISSION GROUND RULES AND ASSUMPTIONS

The mission design process, as with any modeling effort, starts with some known set of initial conditions with which to develop a trajectory. This set may be referred to as ground rules and assumptions (GR&As) and it contains formal mission requirements, operating and design-to constraints, performance assumptions for major subsystems, programmatic objectives and desires, and details of the computational environment in which a trajectory will be modeled. In the context of the Lunar Transit mission these GR&As can include, among others (Figure 2 presents a larger, but not exhaustive list):

1. Minimum mass delivery to the NRHO (requirement/constraint)
2. Qualified operating lifetime of SEP thrusters (requirement/constraint)
3. Maximum eclipse duration (requirement constraint)
4. Maximum total ionizing dose (requirement/constraint)

5. Launch vehicle performance (system performance estimate)
6. Thruster and power system performance (system performance estimate)
7. Launch date (programmatic direction)
8. Available computing resources (computational consideration)
9. Trajectory design software capabilities (computational considerations)

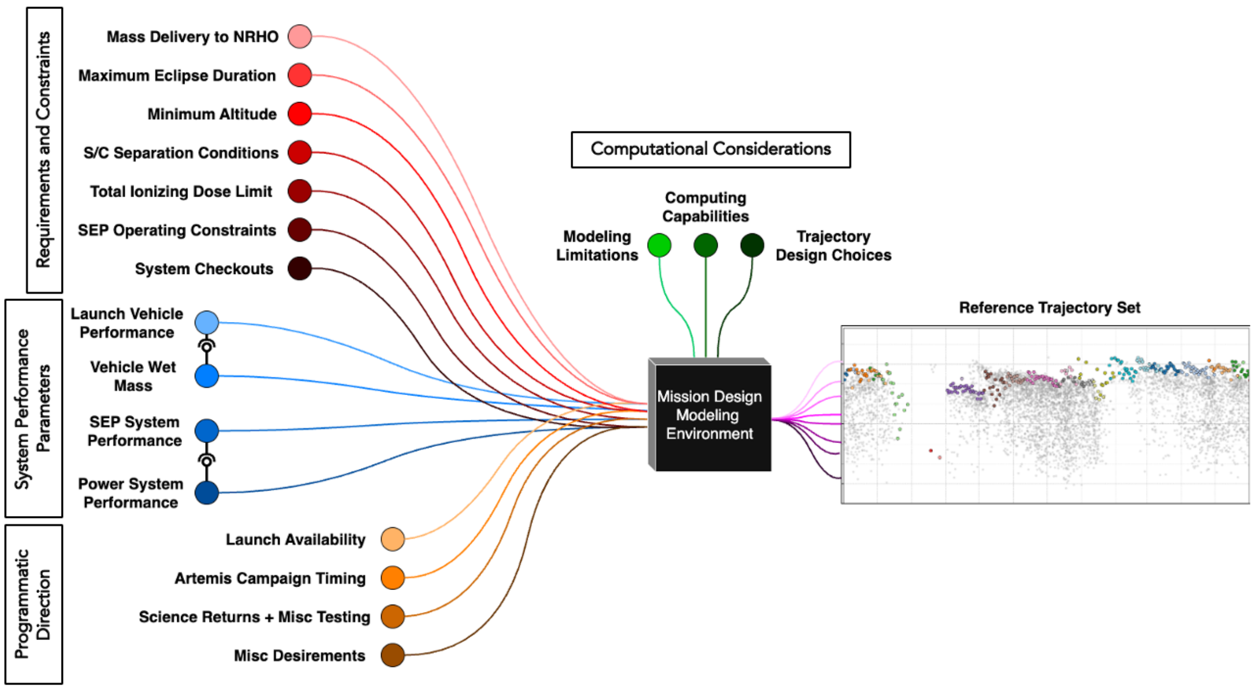

Figure 2: Dependencies of the Lunar Transit integrated trajectory design problem

The complete list of GR&As form the context within which a deterministic trajectory is designed. Some GR&As, such as the choice of trajectory simulation software and available computational resources, may be selected according to specific mission needs or supplemented at cost. Others, such as concept of operations (CONOPS) events and timelines, may be determined through engineering assessments of mission needs and can be adjusted within a defined risk posture to ease the burden on trajectory design. Many GR&As, often taking the form of requirements and constraints, are dictated to engineering teams or are natural consequences of decisions such as the selection of a specific launch vehicle, thruster type, and solar array design.

The specific hardware that is selected for a spacecraft, and the manner in which it is tested, forms an inherent limitation on the operation of the vehicle. Hardware is typically qualified to some defined test criteria, whether it be a safe envelope of operating temperatures, operating lifetime before a component may fail, or limits on operating environment before reliability is reduced. These qualification limits are themselves constraints, governing the manner in which a system can be assumed to be safe to operate. In practice, this means that a mission which accounts for a physical spacecraft finds itself far more limited than a theoretical one. Life limits, maximum rate of change of vehicle attitude[10], minimum fly-by altitudes governed by thermal limits of components, or maximum eclipse durations due to temperature and energy storage capabilities all form bounds which are often more constraining than formal mission requirements or objectives. Requirements may be waived if the situation demands it, while hardware failures that lead to loss of mission cannot be ignored or written off.

These limits and criteria rigidly bound a trajectory and must either be directly included in the optimization problem or assessed a-posteriori for compliance. Additionally, system performance assumptions may contain interdependencies which complicate modeling. The parking orbit a trajectory may begin from is dependent on the performance of the selected launch vehicle as well as the mass and separation condition requirements of the spacecraft and the power available to the SEP system is a function of the performance of the vehicle power system, including the solar array capabilities and the ability of the energy storage system to recover from periods of eclipse. These system interactions may be modeled and determined a-priori (such as park orbit energy) or they may need to be modeled within the trajectory optimization problem (such as SEP system performance over time), further complicating the process of developing a deterministic trajectory.

MISSION UNCERTAINTIES

In addition to the complexities of developing a single, nominal, deterministic trajectory a flight mission must consider a host of uncertainties and unknowns to determine whether a mission design is executable. To carry a sufficiently high confidence that a mission can be successfully flown a mission designer needs to incorporate some set of performance parameters and events which *may* be realized during flight but cannot be accurately predicted beforehand. Beginning with launch, the spacecraft will be inserted into an initial orbit with some previously unknown (but quantified) error, vehicle systems will perform differently than predicted pre-flight (these differences may be within quantified uncertainties or they may be significant if an anomaly has occurred), and so the as-flown trajectory will depart from the pre-flight reference (Figure 3). The magnitude of these uncertainties and unknowns which the mission designer must protect for is defined through the mission's risk management process.

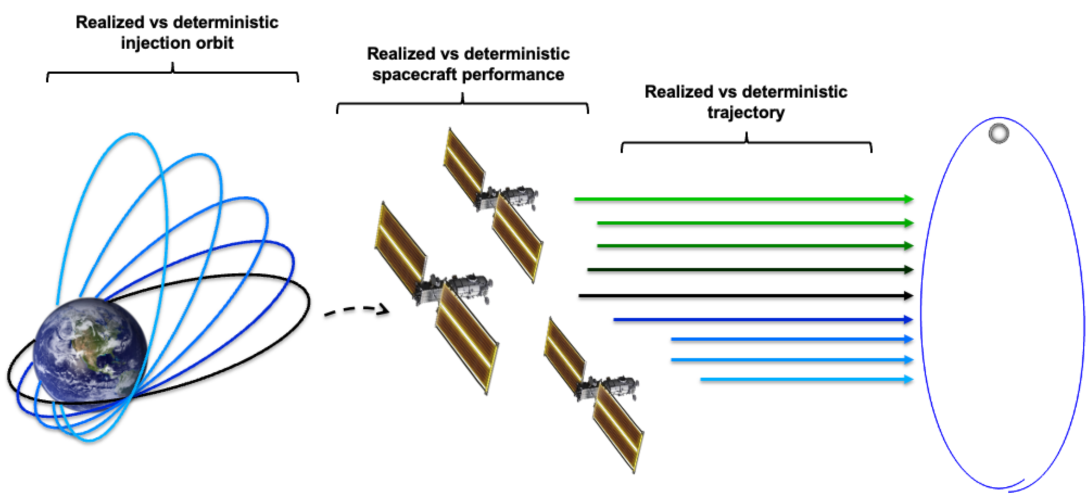

Figure 3: Simplification of differences between deterministic (black) vs realized (colored) launch vehicle injection orbits, spacecraft performances, and trajectories to reach the NRHO

General Mission Uncertainties

All space missions, regardless of size, objective, or design, must deal with the epistemic and aleatoric uncertainties inherent to remotely operating a complex system in an environment that cannot be fully recreated or simulated in a test facility. Pre-flight designs rely on the quality of test data, system-level simulations, and environmental models. This pre-flight knowledge can vary depending on the capabilities of the facilities used for testing, the precision of the instruments and methods used to collect data, and the fidelity of the models used to generate performance predictions. The uncertainty present in test measurements and integrated models can often be quantified, meaning that it is a “known unknown”. Some statistical understanding of the uncertainty provides engineers with a probability distribution of the degree to which uncertainty may be realized during mission execution, allowing its impact to be characterized through Monte Carlo methods.

The manner in which a spacecraft is flown also dictates how uncertainties may impact operations. No physical system can perfectly execute commands and so there will always be some delta between what the ground wants a spacecraft to do and what it actually does (this is typically referred to as an execution error). Determining what a spacecraft has executed, and at what time, is not necessarily a straightforward task. Unlike ground-based systems, it is not possible to physically inspect an on-orbit spacecraft and the knowledge that operators have of the vehicle in real-time is only as extensive as the onboard sensors and downlinked telemetry allow. Telemetry is also not always available depending on the frequency with which ground stations are available or scheduled, meaning that available knowledge can quickly become obsolete. This gap between the state of the spacecraft and the ground’s understanding of it must be considered pre-flight and included in uncertainty analysis to validate a trajectory architecture.

Uncertainties Inherent to Electric Propulsion

While HETs are a mature technology, there remains a significant lack of fundamental knowledge regarding the in-space performance of these systems. While HETs are some of the most numerous forms of propulsion on-orbit today, it remains difficult to characterize the output of these thrusters on the ground, there is no standard or agreed upon method by which to extrapolate ground test data (in a facility which imperfectly simulates vacuum) to on-orbit performance, and it is not often possible to recreate precise estimates of on-orbit performance from spacecraft telemetry.

The AEPS thruster, of which three were planned to be used on PPE, produces a maximum of 586 mN of thrust at a discharge power of 12 kW[11]. The thruster has a mass of <47 kg[12], meaning that AEPS has a thrust to weight ratio of roughly 0.001 at standard gravity. Measuring such a small force acting upon a much more significant mass is inherently challenging. For PPE, the AEPS thruster performed qualification and acceptance testing at NASA Glenn Research Center’s electric propulsion test facility. Measurement uncertainty at this facility has previously been characterized as +/- 6.9 mN at a 95% confidence interval for a thruster as powerful as AEPS[13]. While small, this uncertainty still represents 1.2% of the full thrust magnitude, an error that must be accounted for across roughly 300 days of thrusting in the Lunar Transit mission.

In addition to accurately measuring the performance of a thruster on the ground, this performance must be extrapolated to an on-orbit environment. Previous work has shown that HETs, including the magnetically shielded AEPS, are sensitive to facility backpressure when tested on the ground[14]. Generally, a HET will display higher specific impulse (efficiency) at higher backpressures associated with ground-based vacuum test chambers. As space forms a more perfect vacuum than any test chamber, the on-orbit background pressure is very close to zero and the realized efficiency and thrust magnitude of a HET is often lower than that observed in pre-flight testing. Unfortunately, due to the general lack of in-space data available for HET performance, no standard method exists by which to characterize this difference and account for it in pre-flight analysis.

Just as ground-test data carries a high-degree of uncertainty for HETs, in-flight estimates can be difficult to obtain and imprecise. Performance estimates of thrusters in space are often obtained by reconstructing thrust magnitude over some period of time from orbit determination (OD) solutions, meaning that these estimates are dependent on the quality of the OD being performed. Vehicle telemetry pertaining to the state of valves within the propellant management system can additionally provide insight into the mass flow rate of the system over time, allowing specific impulse and mass evolution to be reconstructed as well. These methods also carry their own uncertainties and missed communications passes may lead to onboard data being permanently lost due to data storage limits, meaning that an accurate accounting of the spacecraft's mass over time may not be available. Without an accurate mass estimate thrust estimation is further complicated, even when onboard accelerometers are available.

While precise characterization of how a thruster will perform on or-orbit when commanded is currently not available to mission designers, the Psyche mission has shown that even systems with considerable flight heritage may exhibit unexpected behaviors. The SPT-140, flown on the Psyche spacecraft, has a strong flight heritage with Intuitive Machines (formerly Maxar) having operated at least 4 SPT-140-equipped spacecraft prior to Psyche's launch[15]. Despite this extensive history of on-orbit use and characterization, the mode-hopping behavior[3] observed during Psyche's electric propulsion checkout activities had never been documented previously. This performance instability, in a thruster with a significant amount of previous operation, further exemplifies the uncertainties inherent in the use of high-powered electric propulsion in space missions.

Mission Specific Uncertainties

Gateway and the PPE were envisioned to be a first-of-its-kind application of SEP to a large, high mass, human-rated exploration vehicle. Gateway was to incorporate a number of never-before-flown technologies, including the largest roll-out solar arrays (ROSAs) and the highest power SEP system developed to date. This system included two newly developed thruster types (AEPS and the BHT-6000) neither of which has operated in-space. The result was a mission that inherently carried uncertainties in the performance of critical systems that have never previously operated on-orbit (such as power and propulsion). Further uncertainty is added by the novel nature of this mission to NASA human space flight (HSF), which has previously only operated high-thrust spacecraft and has largely been confined to low-Earth Orbit (LEO).

While all space missions account for some uncertainty in system performance about a mean/nominal specification, the number of new technologies onboard PPE/Gateway led to a statistically significant degree of variation in the mean value itself. Where heritage low-thrust verification methods utilizing Veil and the newly developed Stampede[6] have assessed cycle-to-cycle variation of performance about a mean value, the Lunar Transit required significant focus to be put on analysis of variations in mean performance values for SEP thrust magnitude and power availability. Further complicating this analysis was the need to perform this analysis across a wide range of potential launch dates due to system sensitivity to seasonal variations in solar array performance and Earth-Moon system alignment.

Missed Thrust and Contingency Scenarios

In addition to the uncertainties discussed above, the Lunar Transit mission design needed to consider a number of potential disruptive events that could occur during execution. These include MTEs (a catch-all term used to describe any time that a vehicle safe-mode or unplanned coast period disrupts the execution of a low-thrust maneuver) as well as contingency scenarios such as failed SEP thruster. MTEs are a persistent threat to any low-thrust mission, and the Lunar Transit was no different. With more than 300 days of planned thruster operation, missed thrust proved to be a significant driver of sensitivity analysis. Additionally, due to the high thruster throughput required to reach the NRHO (more than 3,000 m/s Δv), the loss of a SEP thruster and particularly the loss of one of the three AEPS would make successful execution of the remainder of the trajectory challenging.

Application of the methods described by Imken, et al[16], to the Lunar Transit reflected a high likelihood of multiple vehicle safe mode events over the duration of the mission. These empirical forecasts, developed against historical safe-mode event occurrences and durations in deep space missions, indicated that operators could expect to experience between 2 and 7 individual MTEs with a total outage duration between 3 and 10 days (Figure 4). In addition to MTEs caused by safe mode events, operational considerations such as orbital object conjunctions were accounted for under the missed thrust umbrella. Previous publications[17] have detailed the strategy PPE MD developed for collision avoidance maneuvers (CAMs) during the Lunar Transit by which an additional, unplanned, coast period would be used to generate separation between the Gateway spacecraft and a potential conjuncting object. Due to the highly eccentric nature of the Lunar Transit during the early phases of flight, a coast on the order of an hour placed at apogee passage was shown to be effective at producing miss distances on the order of tens to hundreds of kilometers. Due to the low initial perigee altitude of the Lunar Transit (200 km) and the dense object environment at low Earth altitudes below 2,000 km, multiple CAMs were anticipated to be necessary during flight. As CAM implementation was based in the addition of unplanned coasting, this was, by definition, intentional missed thrust and was included in the broader MTE context.

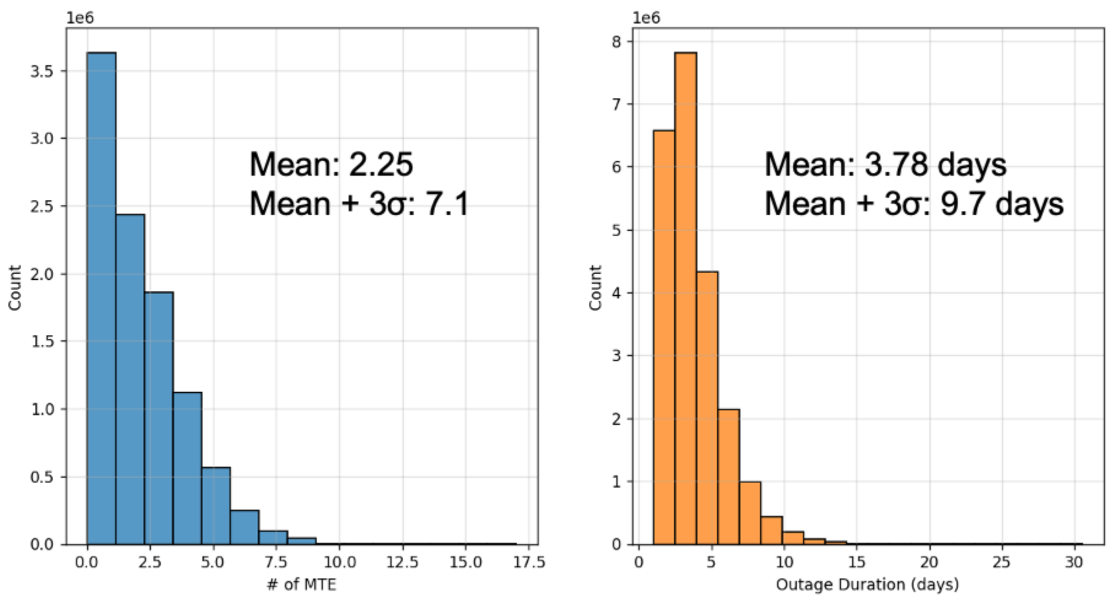

Figure 4: Empirical MTE forecasting using methods outlined by Imken, et al[16]

Failures within the primary propulsion system can have catastrophic consequences for spacecraft and lead to loss of mission. PPE's SEP system inherently carried redundancy through the integration of 7 separate thrusters, 3x AEPS and 4x BHT-6000 thrusters. During transit, however, the AEPS strings would have produced roughly 80% of the total impulse required (a natural consequence of the AEPS units being roughly twice as powerful as the BHT-6000). Thruster string failures were thus analyzed through the bounding case of the loss of a single AEPS unit on the first day of the mission, at mission elapsed time (MET) +0 days. Analysis focused on the ability of the spacecraft to reach the NRHO within existing mission constraints while utilizing the remaining 6 healthy strings in order to show fault tolerance within the propulsion system.

RISK MITIGATION

Uncertainties and contingency scenarios represent risks to the ability of the combined spacecraft and operator teams to successfully execute the mission within established constraints and requirements. To ensure that sufficient capability exists to protect against a set of accepted risks, margins and reserves were applied to critical mission and spacecraft parameters. If, when, and to what magnitude a risk would be realized during flight would prompt the consumption of part of these margins and reserves to offset the impact.

Within the context of mission design, margins and reserves were primarily applied to propellant (particularly xenon mass) and time, including thruster duty cycles included for maneuver expansion periods, explicit coast periods reserved for things like missed thrust recovery, and padding placed on required coasting events such as eclipse passage to protect against navigation state uncertainty. Broader vehicle margins such as mass growth allowance are beyond the scope of mission design and thus not discussed in this publication.

For the Lunar Transit mission margin and reserves were defined as separate quantities and applied to specific forms of uncertainty. Reserves were aimed at protecting against "known unknowns", uncertainties that were likely or guaranteed to be realized but the magnitude of the realization would not be known pre-flight. This includes epistemic and aleatoric uncertainties such as thruster performance dispersions and launch vehicle injection error. Missed thrust was also included under reserves due to the high likelihood of realization indicated by empirical forecasting. Margins, on the other hand, targeted "unknown unknowns", events or uncertainties that *may* or *may not* occur during flight. This category broadly included anything where both the occurrence and the magnitude of realization was not certain, such as contingency scenarios and significant dispersions beyond the magnitude protected for with reserves (typically 3 standard deviations). Also included in margins were any risks, uncertainties, or events that were not being tracked *at all* in preflight analysis due to incomplete or insufficient knowledge of their potential existence. While some margins, such as contingencies, could be directly quantified through analysis these "surprises" lurking beyond the limits of the engineering team's knowledge would be covered with whatever propellant or time was naturally "left over" in the development of the deterministic mission. For example, should a surplus of xenon exist between the onboard storage capacity and the nominal mission and reserve and margin needs, then this extra propellant would flow to margins. Table 1 summarizes these definitions.

Table 1: Reserves and margins overview

<table>
<tr><th>Reserves</th><th>Margins</th></tr>
<tr><td>Known Unknowns
Thruster performance uncertainty
Orbit insertion dispersions
Power availability uncertainty
Navigation uncertainty
Execution errors
Missed thrust events</td><td>Unknown Unknowns
Thruster failures
Prolonged loss of comms
Beyond 3σ dispersions
Risks not being tracked preflight
Vehicle contingencies</td></tr>
</table>

RESULTS

Monte Carlo analyses and targeted assessments of specific contingency cases became a focus of PPE MD as the Lunar Transit mission advanced from conceptual design towards execution. Validating the robustness of trajectories to expected uncertainties became a critical and necessary step towards certification of flight readiness. While many analyses covering a large array of uncertainties, execution errors, and operational implementations were undertaken, results presented here will focus on launch vehicle injection error, thruster performance uncertainty, MTEs, and thruster failures. Reserves and margins are presented in terms of the percent of the deterministic xenon cost required to account for each realization. A brief discussion of time margins as it pertains to a maximum experienced eclipse requirement of 90-minutes is also included.

Launch Vehicle Injection Error

The Lunar Transit trajectory begins with a target interface point (TIP), which serves as a patch point between the ascent trajectory modeled by the launch vehicle team and the in-space trajectory modeled by PPE MD. The spacecraft will not perfectly be delivered to the identified TIP state by the launch vehicle and so the parking orbit from which the Lunar Transit begins will initialize with some insertion error. The first week of the Lunar Transit was envisioned to consist of system checkouts and a multi-day SEP system commissioning, meaning that insertion errors would propagate freely during this time without the ability to perform correction maneuvers. Due to the low perigee altitude of this parking orbit (200 km), perturbations due to drag, Earth spherical oblateness, and solar gravity would magnify dispersions in initial state.

Lunar Transit operations planning assumed that the pre-flight reference trajectory would be redesigned during the initial 7-day checkout to incorporate real-time knowledge in park orbit state and vehicle system performance. Dispersions in TIP insertion would thus be incorporated into a newly designed reference trajectory which would mitigate their impact relative to requiring short-term retargeting of the original, pre-flight reference. Analysis completed using the Aberrant package simulated this process by sampling from a known distribution of perturbations to park orbit parameters including:

1. Apogee altitude
2. Perigee altitude
3. Inclination
4. Argument of perigee (AOP)
5. Right Ascension of the Ascending Node (RAAN) (determined by launch epoch)

The resulting trajectory was then reoptimized to reach the NRHO. Of note is that due to the manner in which TIP states were defined (in an Earth body-fixed frame), RAAN was solely a function of the time of day in which a launch occurred (assuming that ascent guidance and duration did not change from one minute to the next). Monte Carlo analysis sampled across distributions in each parameter simultaneously and covered a potential launch window of 20 minutes centered about the optimal epoch determined in reference trajectory development (+/- 10 minutes). Representative analysis covers a single, notional, month of potential launches (presented here as March of 2027) and contains roughly 20,000 samples.

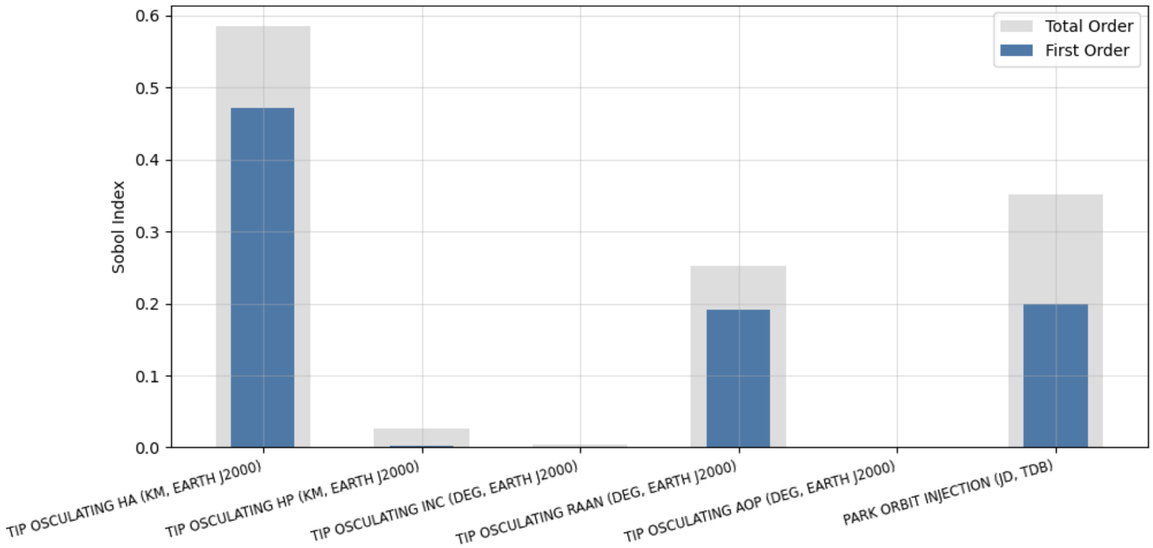

Figure 5: Sobol sensitivity indices of propellant consumption to varied TIP state parameters

As multiple parameters are being varied simultaneously (apogee, perigee, inclination, RAAN, AOP, and launch date), a Sobol sensitivity analysis[18] is conducted to determine trajectory propellant usage sensitivity to variations in each value. From Figure 5 it can be seen that variations in apogee altitude (and thus orbital energy) dominate the data set. This is to be expected, as increases in orbital energy decrease the Δv required from the spacecraft to transit from GTO to NRHO (and the converse is true as well). The second most sensitive component is launch date, represented by park orbit injection epoch, which is also to be expected due to the daily evolution of the Earth-Moon system geometry and the fixed launch inclination of 28.5˚. Finally, variations in TIP RAAN form the final significant contribution to changes in xenon use, which will be discussed below.

As indicated by the Sobol results, xenon propellant required to correct for dispersions in apogee altitude forms a clearly identifiable trend, where higher apogee altitudes (positive dispersions) require less propellant than the nominal trajectory (lower xenon reserve). Figure 6 shows that 3σ dispersions in apogee altitude result in between +/- 3% impact to available xenon reserves with hot insertions potentially saving xenon relative to the mean. The results also show two clear solution families separated by TIP RAAN, with one family forming around 45˚ and one forming near -15˚. This trend is known to PPE MD, where solutions fall into one of two RAAN orientations depending on launch date. This is a result of the orientation of a launch trajectory with a 28.5˚ inclination and the Moon's orbital plane about the Earth, where the -15˚ and 45˚ solution families produce a spacecraft trajectory that arrive in a plane conducive to a lunar intercept at the end of the Earth-centered spiral portion of the Lunar Transit. Solutions within these two families can exhibit different xenon costs and eclipse trends from one another and this behavior has been noted across three separate iterations of reference trajectory design covering launch dates spanning three years (2026, 2027, and 2028). The RAAN sensitivity seen in the Sobol results are thus a product of this solution family behavior, and not the +/- 10-minute launch window perturbation applied to the Monte Carlo analysis (and in fact no significant sensitivity between launch window and trajectory performance have been noted by PPE MD).

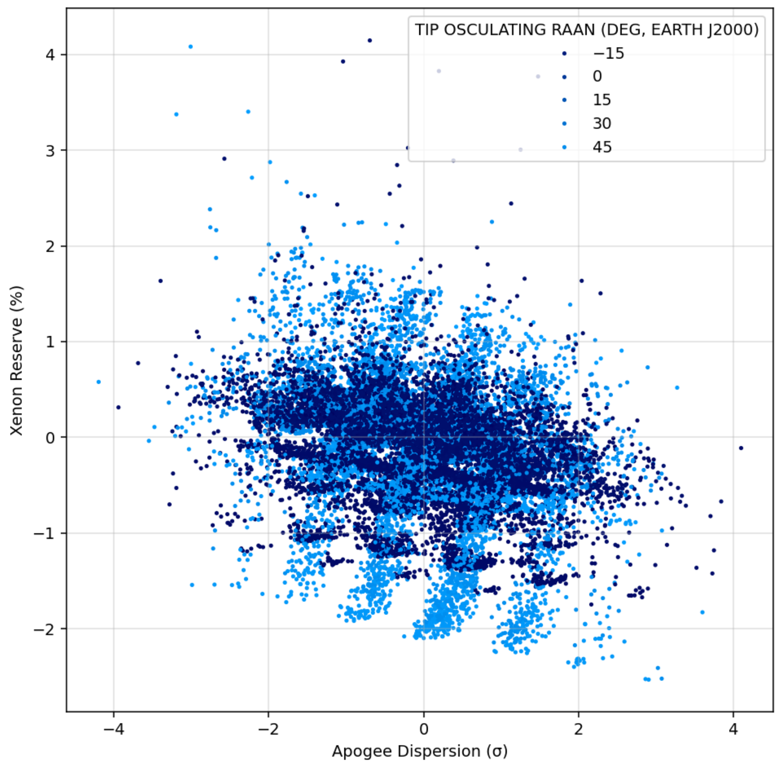

Figure 6: Xenon reserve sensitivity to apogee altitude dispersions

Thruster Performance Uncertainty

As discussed previously, PPE carried inherent uncertainty in the on-orbit performance of the SEP thrusters due to ground measurement uncertainty and performance differences in varying back-pressure environments. As with TIP insertion accuracy, a statistical distribution of thruster performance (thrust magnitude and specific impulse) as well as power system performance (in power available to the SEP system) was developed and applied to Lunar Transit reference trajectories in a Monte Carlo fashion.

Once again, a Sobol sensitivity analysis is employed to determine the sensitivity of xenon usage to variations in thrust magnitude, specific impulse, power to SEP, and launch date (this analysis is performed over the same March 2027 period as above). Sobol results (Figure 7) show a dominant sensitivity to variations in thruster efficiency with a secondary impact from thrust magnitude, a small contribution from power to SEP (which serves as an additional adjustment to thrust and specific impulse via thruster throttle level), and very little impact from launch date. This result is indicative of the comparatively larger uncertainty in SEP thruster performance than power system performance, which is better characterized due to the mature implementation of ROSAs in spaceflight missions. The second order impact of thrust magnitude to xenon usage has been determined to be a result of lower-thrust trajectories experiencing a greater number of eclipses than nominals (to be discussed below). This greater time spent in shadow increases the time required to spiral out of Earth's gravity well, resulting in worsening geometric alignment of the trajectory with the Moon's orbital plane when the spacecraft reaches lunar intercept. This alignment is corrected propulsively, resulting in increased propellant consumption at lower thrust magnitudes.

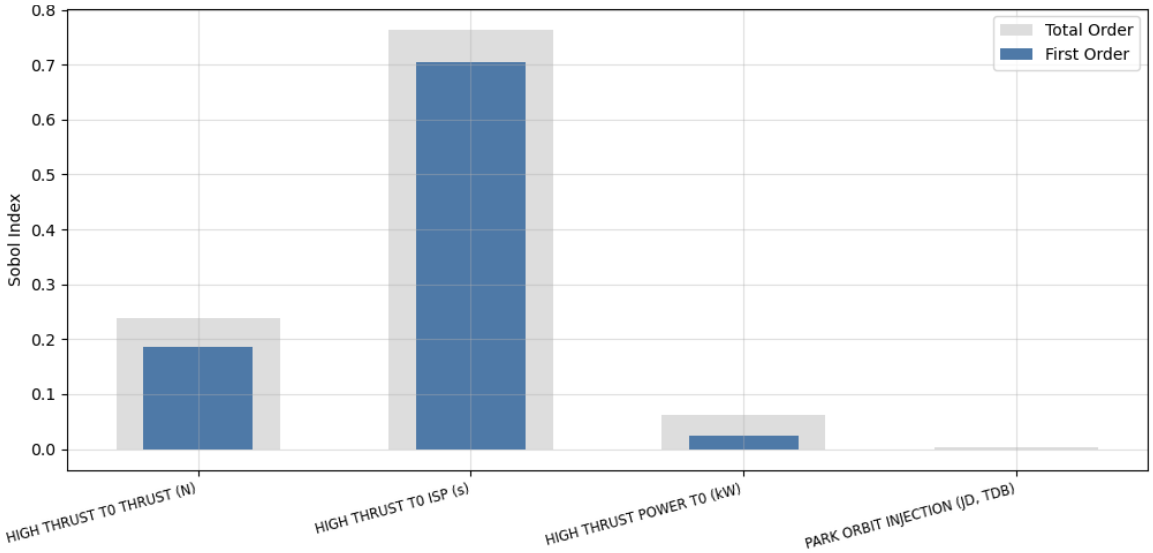

Figure 7: Sobol sensitivity of xenon usage to changes in SEP performance parameters

Results (Figure 8) show a strong correlation between specific impulse and xenon use, as expected via the rocket equation. This trend is nearly 1:1, with a 1σ reduction in specific impulse requiring, on average, 1% more xenon to recovery from. This nearly linear relationship between propellant use and specific impulse is a happy coincidence of the magnitude of performance dispersion, which, for the purposes of analysis, was assumed to be roughly 0.8% 1σ. A ~1% variation in specific impulse resulting in a ~1% variation in propellant use not surprising to mission designers, but this linearity is an example of not only the law of small numbers but also how the high baseline SEP specific impulse of over 2,600 s helps to flatten the curve of the rocket equation. These results additionally show that just above 4% xenon reserve is required to protect against 3σ dispersions in thruster specific impulse.

Unlike propellant consumption, which is primarily a function of specific impulse, eclipse profiles are primarily a function of thrust magnitude. Higher thrust magnitudes mean that when the spacecraft reaches the beginning of an eclipse season, the size of the vehicle's orbit will have grown at a faster rate than the nominal. Meaning that each revolution within the eclipse season will be larger and thus orbital velocity across that revolution will, on average, be lower. This means the spacecraft will transit the Earth's umbra more slowly, leading to increased eclipse durations with higher thrust magnitudes. The spacecraft will, however, exit an eclipse season entirely more quickly due to the higher rate of change of semi-major axis. This behavior is seen in Figure 9, which shows increasing eclipse durations and decreasing Earth eclipse counts with higher thrust dispersions across the entirety of March of 2027. Across 3σ dispersions in thrust, the duration of the longest experienced eclipse increases by ~3 minutes over baseline, while the number of eclipses decreases by roughly 10. Figure 9 also shows trajectories which violate the 90-minute requirement at higher thrust dispersions, meaning that while the baseline trajectory may have been requirement compliant, a mission design which accounts for uncertainty in performance is not. This trend reflects the need for margin to be applied to path constraints such as a maximum eclipse duration in order to account for real world uncertainty in execution.

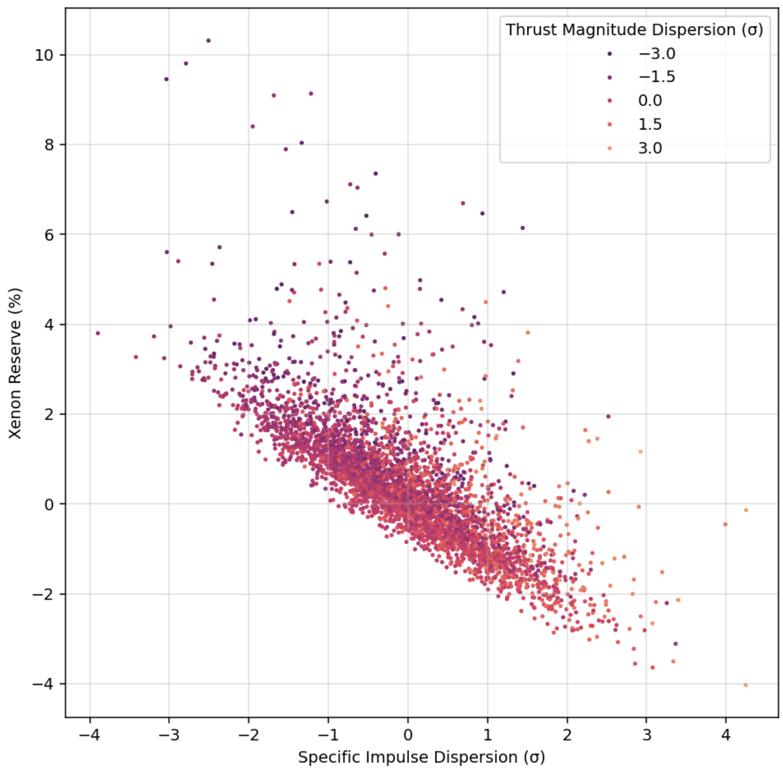

Figure 8: Xenon reserves required vs specific impulse dispersions

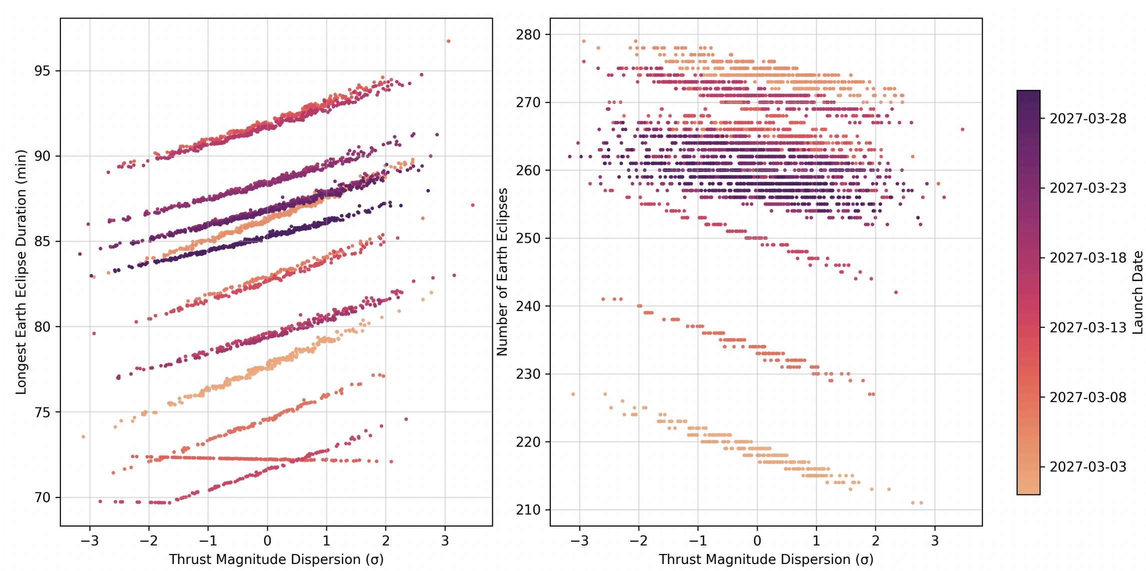

Figure 9: Eclipse duration (left) and number of total Earth eclipses (right) vs thrust magnitude dispersion, colored by launch date

Missed Thrust

Previous work[10] as well as the empirical results in Figure 4 have highlighted the probability of missed thrust across the end-to-end Lunar Transit. While the empirical results indicated a 3σ probability of encountering just under 10-days of total thrust outage during the Lunar Transit, the combined PPE engineering and operations team conducted a separate, vehicle specific, assessment to determine a required missed thrust tolerance for the Lunar Transit. Flight controllers from NASA's Mission Control Center Houston (MCC-H) along with mission design concluded that a resilience to 21-days of total thrust outage occurring at any time during transit would provide a reasonable not-to-exceed time for anomaly resolution and spacecraft recovery during operations. This maximum outage duration of 21-days became a "known unknown", with the xenon required to protect against an outage of this duration coming in the form of reserve.

Using the framework contained in Aberrant, 21-day MTEs were analyzed using the Rolling Outage approach[10], beginning at MET +0 and continuing through the end of the Spiral Phase (near MET +275). Further analysis assessed MTEs in the Alignment and Insertion Phases but have not been included here for conciseness. In each case, the remaining spacecraft trajectory was reoptimized post-MTE to target the NRHO. To limit the propellant impact of an outage, 21-day MTEs prompted an arrival delay of 4 NRHO revolutions, roughly corresponding to one lunar month, or 29 days. This delay prevented the need for excessive thrusting during later phases of the mission and preserved optimal coasting in the Alignment Phase while maintaining a favorable geometry by which the Gateway would execute a lunar intercept.

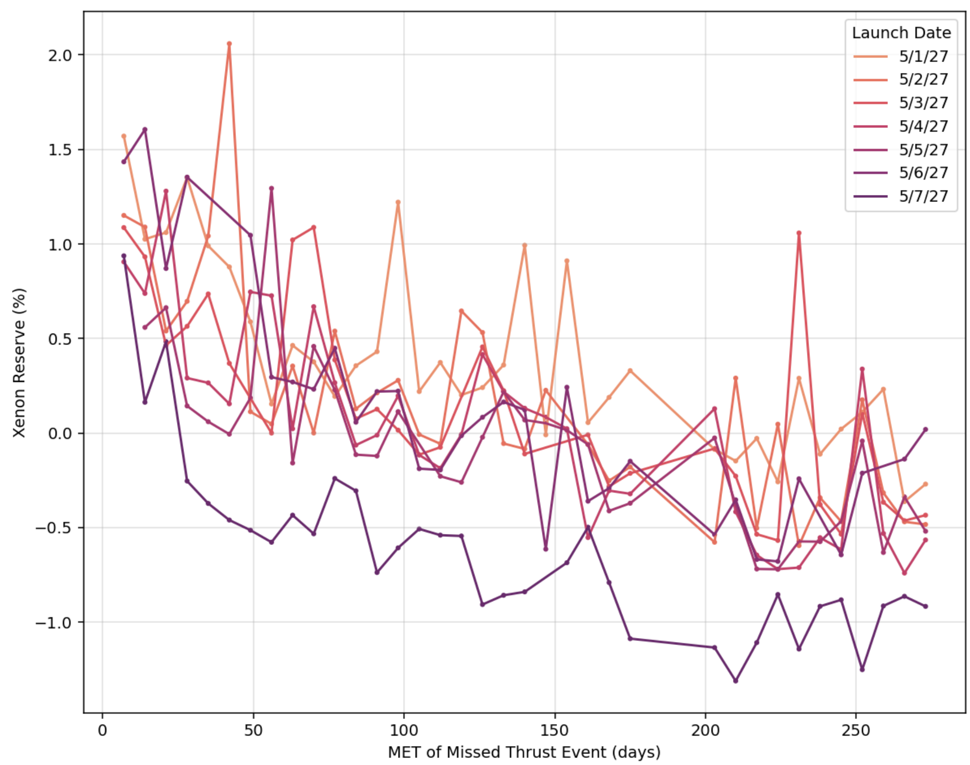

Figure 10: Xenon reserve as a function of the MET occurrence of a 21-day missed thrust event

Missed thrust results for 21-day outages show that a xenon reserve of roughly 2% is sufficient to recover from MTEs at any point in the Spiral Phase, provided the spacecraft can arrive at the NRHO 29 days later than planned (Figure 10). Sensitivity is not constant over the duration of the Spiral Phase, however, with xenon costs decreasing as the MET at which an outage occurs increases. This is due to the increasing perigee altitude of the trajectory over time, meaning that atmospheric drag and Earth spherical oblateness induce fewer perturbations to trajectories with outages later in the mission. The results also show a launch date sensitivity, with certain launch dates exhibiting lower costs than others (the May 7th solution is noticeably more robust to MTEs than solutions on May 1st through the 6th).

Thruster Failures

Given the strategic value of the Gateway to the Artemis architecture as envisioned in the beginning of 2026, it was desirable to maintain fault tolerance and anomaly resilience where possible. This included resilience to thruster failures, with the MET +0 loss of a single AEPS thruster determined to be the most stressing case for the Lunar Transit. To assess this contingency, the same set of trajectories analyzed in the missed thrust assessment were modified to represent the loss of a single AEPS. While the PPE carried 3x AEPS, it also would have carried 4x BHT-6000s, and so following the loss of a 12 kW string the Lunar Transit would be completed on 2x AEPS and 4x BHTs. Complicating this analysis is the need to complete a recovery within the qualified lifetime of the BHT-6000 strings which, due to their lack of magnetic shielding, are expected to degrade over their life and eventually fail due to this wear mechanism.

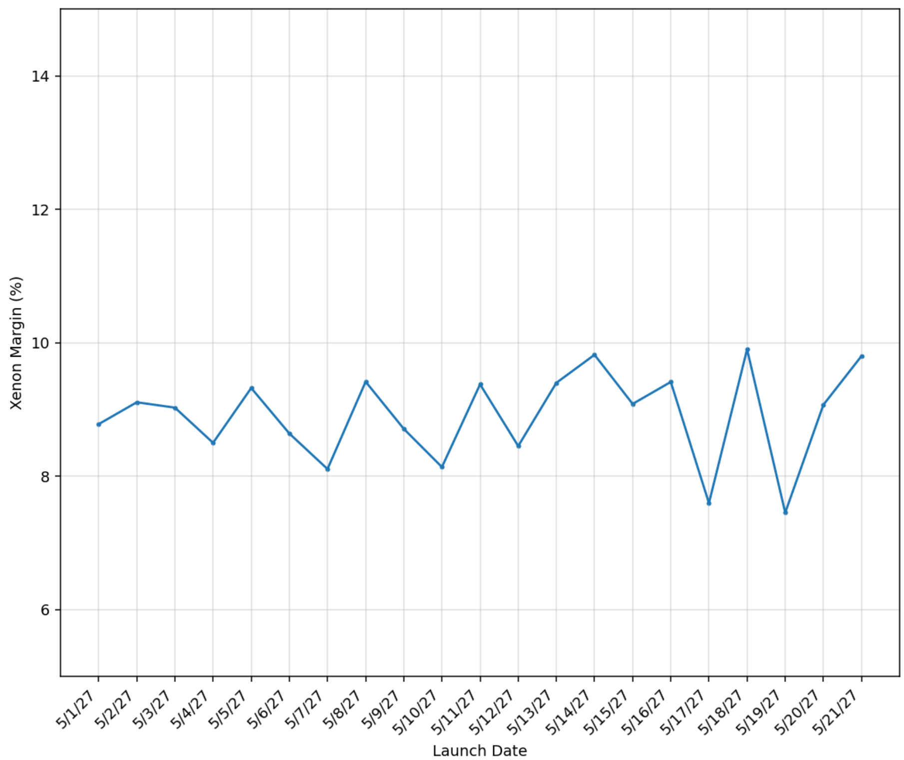

Figure 11: Single AEPS out recovery results for May 2027

Figure 11 shows the xenon margin (not reserve, as contingencies are defined as unknown-unknowns) required to recover from an MET +0 AEPS failure across launch dates in May of 2027. These recovery costs are significantly higher than those for the three types of stochastic uncertainties described above, with a maximum value of ~10%. Due to the significant shift in impulse handling away from the AEPS and towards the remaining BHT-6000s, total system efficiency drops in these recovery cases and propellant consumption increases. While the AEPS thrusters produce nearly 80% of the total required mission impulse in nominal scenarios, that share drops to just 60%

in a failure case. Margin, like reserves, is not made of just one value but is meant to cover multiple potentialities. One individual scenario requiring 10% margin represents a significant driver of propellant needs and posed program leadership with a challenging risk posture to incorporate into mission-level accounting.

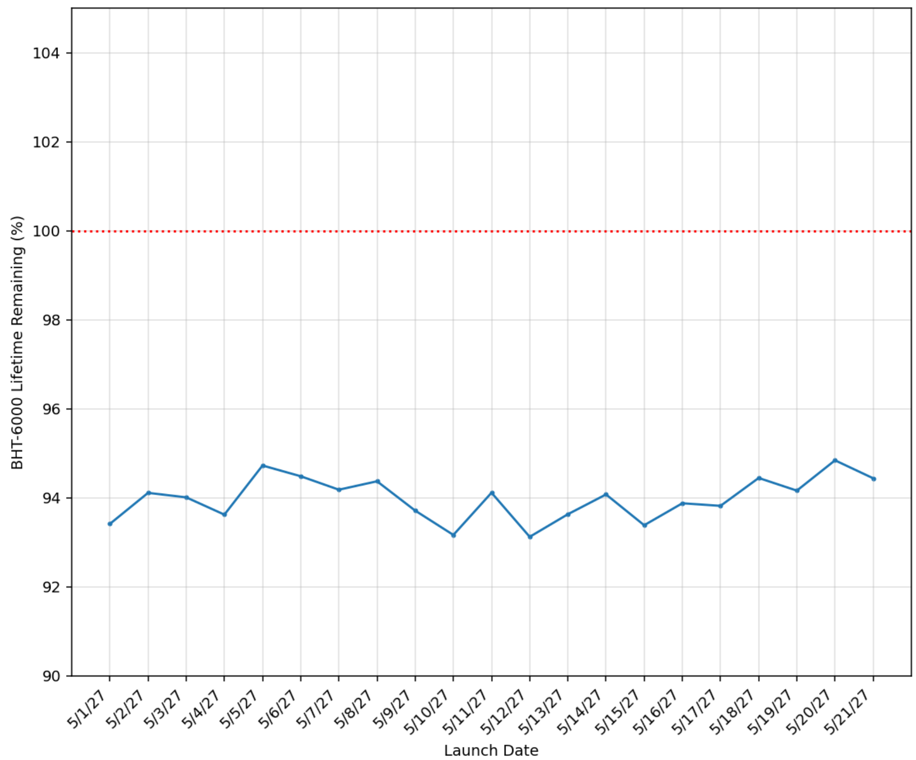

Figure 12: BHT-6000 lifetime utilization

Nominal Lunar Transit trajectories required roughly 50% of the available lifetime of each BHT-6000 thruster to complete (a more comprehensive discussion of thruster utilization has been presented previously[19]). As Figure 12 reflects, AEPS failure recoveries nearly double that usage to between 93% and 95%. In practice, mission designers were forced to throttle the BHT thrusters back in these recovery scenarios in order to prevent lifetime exceedances, further reducing the overall system efficiency and increasing propellant consumption further. This qualification limit became a difficult design constraint to comply with and one that drove thruster implementation decisions during contingency recoveries.

Impact to Margins and Reserves

These four analyses each are the result of the application of some set of accepted GR&As to reference trajectories representing deterministic missions. The results reflect a potential distribution of risk postures, with mean values for stochastic uncertainty sensitivity near zero for all three reserve-based analyses up to 4.6% 3σ in the case of thruster performance errors. Likewise, AEPS out cases represent a choice to protect against single fault tolerance in the SEP system, driving a high propellant margin of nearly 10%. When added together, these analyses represent 9% (3σ) of xenon reserve and 9.9% of xenon margin (Table 2). It is, of course, possible to reduce these margins and reserves by choosing to accept a higher degree of risk on mission execution. Mission managers can choose to account for 2σ uncertainties in values such as thruster performance or reliability assessments may indicate that an AEPS failure is not likely enough to warrant carrying an explicit margin. These decisions are themselves part of the GR&As that guide analysis and the risk posture that has been accepted at a project or program level drive not only trajectory design and analysis decisions but also hardware decisions on the spacecraft (such as propellant tank size to accommodate additional xenon).

Table 2: Reserves and Margins Impacts

	Mean + 3σ Xenon Reserve	Xenon Margin
TIP Injection Error	1.7%	0%
SEP Uncertainty	4.6%	0%
Missed Thrust (21-days)	2.7%	0%
AEPS Failure	0%	9.9%
Total	**9%**	**9.9%**

CONCLUSIONS AND DISCUSSION

Hardware, as a mission designer may come to learn, is a synonym for constraint. Real, physical systems inherently are imperfect as are the humans that operate them. Ground based facilities provide imperfect or incomplete recreations of the space environment, models imperfectly simulate the operations of systems, and spacecraft and the supporting ground-based apparatuses imperfectly execute missions within the context of incomplete knowledge.

When hardware is qualified for use on a mission, the parameters of those tests set the limits to which systems are accepted to safely operate. On-orbit, systems may safely operate well beyond those bounds but they do so at risk without existing data. In the case of the BHT-6000 thruster (as an example), this takes the form of a qualified lifetime. This number is not a true limit after which the thruster catastrophically fails, it is a limit in knowledge that becomes a design-to constraint in pre-flight analysis. During flight, mission managers may choose to exceed this limit particularly in the event of an anomaly or contingency. Such a capability cannot, however, be depended upon to ensure mission robustness in development. These limits in knowledge become the limits to which engineers can design and analyze, and the final design of a mission will reflect the known capabilities of a spacecraft, not the true maximums.

Any low-thrust spacecraft may face the same uncertainties and contingencies presented in this work. Launch vehicle injection errors apply to all space missions, the thruster performance uncertainties described here are inherent to all electric propulsion technologies, missed thrust is a persistent risk to low-thrust missions, and thruster failures may strike any spacecraft. Likewise, the magnitude of margins and reserves necessary to protect against not just these, but all forms of errors, uncertainties, and risks that face a spacecraft must be accounted for in the development of any flight mission.

To offset just four unknowns in the Lunar Transit required 18.9% of the deterministic xenon utilization (9% for reserves and 9.9% for thruster failures). This is a valuable lesson about the nature of maturing a mission from formulation to flight. Early in mission development (typically following the preliminary design review), the overall architecture and structure of a spacecraft trajectory is likely to be frozen. At this stage, if one were to assume that the spacecraft mass was 10,000 kg, and transit required consuming 1,000 kg of xenon, systems engineers may allocate the remaining 9,000 kg to the remainder of the vehicle. Avionics, power systems, structure, and all other manner of hardware will then be developed, fabricated, and integrated onto a physical spacecraft such that it weighs in at allocation (or often slightly above allocation) prior to launch. If this work is completed without considering margins required to ensure robust mission design, then systems engineers may be shocked to learn during later phases of development that the analysts are not just asking for 1,000 kg of xenon, they are asking for 20%, or 200 kg more! If the now 10,200 kg spacecraft cannot complete the mission, managers must undertake the painful process of either accepting increased risk to burn down margins and reserves or maintain their risk posture and offload mass from the spacecraft. Should this happen after systems are integrated, this may mean directing a technician to physically cut hardware off a nearly complete spacecraft.

Propellant is not the only margin or reserve that must be accounted for, as seen in the case of the 90-minute maximum eclipse requirement on the Lunar Transit. Dispersions in thruster performance indicated that eclipse durations may vary by up to +/- 5 minutes due to uncertainties in mission execution. This indicates to the mission designer that the design-to constraint for deterministic trajectories is not a maximum of 90 minutes but instead is some lower value such as 85 minutes to incorporate sufficient robustness to unknowns.

This same philosophy can be extended to myriad other unknowns; some conservatism must be added to the deterministic trajectory in order to ensure that the mission can be executed under real world realizations. A flyable trajectory is therefore less performant than a conceptual optimal solution designed under less stringent assumptions. Engineers may then choose to incorporate this conservatism early into the formulation process to ensure that a mission is not designed so close to the theoretical maximum performance as to become brittle or impossible to execute when uncertainties are later realized. Likewise, in the development of testing and qualification campaigns it may be valuable to consider the situations in which a system *may* be used and not just the case of a nominal mission. Whether or not a mission may successfully advance from formulation to detailed design to execution is a function of not only available funding and resources, but also to how well the limits of knowledge were considered during development.

REFERENCES

[1] J.R. Brophy, "Perspectives on the Success of Electric Propulsion", *Journal of Electric Propulsion*, 2022
[2] M. J. Patterson and J. S. Sovey, "History of Electric Propulsion at NASA Glenn Research Center: 1956 to Present", *American Society of Civil Engineers*, 2013
[3] S. Snyder, "Initial Checkout of the Psyche Electric Propulsion System", *Journal of Electric Propulsion*, 2025
[4] M. McGuire, S. McCarty, D. Grebow, T. Pavlak, S. Karn, K. Ponnapalli, D. Davis, K. Hack, "Overview of the Lunar Transfer Trajectory of the Co-Manifested First Elements of NASA's Gateway", *AAS/AIAA Astrodynamics Specialist Conference*, Aug 2021
[5] D. Han, "Orbit Transfers for Dawn's Vesta Operations: Navigation and Mission Design Experience", *International Symposium on Space Flight Dynamics*, 2012
[6] C. Venigalla and S. McCarty, "Monte Carlo Analysis of Low-Thrust Many-Revolution Spiral Out Trajectories with Statistical Errors for NASA's Gateway", *AAS/AIAA Space Flight Mechanics Meeting*, 2025
[7] S. McCarty, "Mission Scan Approach for the Gateway Low Thrust Spiral to a Near Rectilinear Halo Orbit", *AAS/AIAA Spaceflight Mechanics Meeting*, Jan 2025.
[8] S. Karn and M. Wentzel-Long, "A Missed Thrust Framework for Low-Thrust Spiral Trajectories to the NRHO", *National Aeronautics and Space Administration*, 2026
[9] Copernicus Trajectory Design and Optimization System, Software Package, Ver. 5.4.0, *NASA Johnson Space Center, Houston*, TX, 2025
[10] M. Martini and B. Humphreys, "The Effects of Solar Beta Angle on Low Thrust Spirals for Solar Electric Propulsion Vehicles", *AAS/AIAA Space Flight Mechanics Meetings*, Jan 2025
[11] R. Shastry, H. Kamhawi, and J. Frieman, "12-kW Advanced Electric Propulsion System Hall Thruster Current Qualification and Production Status", *International Electric Propulsion Conference*, 2024.
[12] "Advanced Electric Propulsion System (AEPS)" *L3 Harris Technologies*, 2024
[13] J. Mackey, et al, "Uncertainty in Inverted Pendulum Thrust Measurements", *54th AIAA/SAE/AIAA Joint Propulsion Conference*, Jul 2018
[14] H. Kamhawi, et al, "Performance, Facility Pressure Effects, and Stability Characterization Tests of NASA's Hall Effect Rocket with Magnetic Shielding Thruster", *52nd AIAA/SAE/AIAA Joint Propulsion Conference*, Jul 2016
[15] T. Kerl, et al "Maxar Electric Propulsion Development for Deep Space", *AIAA Propulsion and Energy Forum*, 2020
[16] T. Imken, R. Randolph, M. DiNicola, and A. Nicholas, "Modeling Spacecraft Safe Mode Events", *IEEE Aerospace Conference*, 2018
[17] S. Karn, et al, "Ballistic Leverage for Collision Avoidance During the Lunar Transit Trajectory of NASA's Co-Manifested Vehicle", *AAS/AIAA Astrodynamics Specialist Conference*, 2023
[18] I. Sobol, "On sensitivity estimation for nonlinear mathematical models", *Matematicheskoe Modelirovanie*, 1990
[19] M. McGuire, S. McCarty, K. Hack, S. Karn, "Application of Solar Electric Propulsion to the Low Thrust Lunar Transit of the Gateway Power and Propulsion Element", *38th International Electric Propulsion Conference*, 2024