NASA/TM-20260001370

A Missed Thrust Framework for Low-Thrust Spiral Trajectories to the NRHO

Scott N. Karn
Glenn Research Center, Cleveland, Ohio

Matthew R. Wentzel-Long
HX5, LLC, Brook Park, Ohio

March 2026

NASA STI Program Report Series

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NTRS Registered and its public interface, the NASA Technical Reports Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or cosponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at http://www.sti.nasa.gov

NASA/TM-20260001370

A Missed Thrust Framework for Low-Thrust Spiral Trajectories to the NRHO

Scott N. Karn
Glenn Research Center, Cleveland, Ohio

Matthew R. Wentzel-Long
HX5, LLC, Brook Park, Ohio

National Aeronautics and
Space Administration

Glenn Research Center
Cleveland, Ohio 44135

March 2026

Level of Review: This material has been technically reviewed by technical management.

This report is available in electronic form at https://www.sti.nasa.gov/ and https://ntrs.nasa.gov/

NASA STI Program/Mail Stop 050
NASA Langley Research Center
Hampton, VA 23681-2199

A Missed Thrust Framework for Low-Thrust Spiral Trajectories to the NRHO

Scott N. Karn
National Aeronautics and Space Administration
Glenn Research Center
Cleveland, Ohio 44135

Matthew R. Wentzel-Long
HX5, LLC
Brook Park, Ohio 44142

Abstract

A framework is developed by which end-to-end optimization of many-revolution low-thrust spiral trajectories can be completed in the presence of missed thrust events. This framework is applied to the Lunar Transit trajectory by which the initial capability of NASA's Gateway lunar space station will be delivered to a Near Rectilinear Halo Orbit. This low-thrust mission consists of three subphases, each designed according to the specific objectives and dynamical regimes encountered as the mission progresses from a medium Earth insertion orbit to cislunar space. The presented framework accounts for the unique considerations demanded by each mission phase and incorporates appropriate capabilities into a novel mission analysis tool. This methodology enables large scale and reliable analyses of missed thrust events across the end-to-end Lunar Transit to verify the robustness of flight trajectories across the full range of considered launch dates.

Nomenclature

Δv	delta-v
t	time
m	mass

1.0 Introduction

Recent interest in using high-powered solar electric propulsion (SEP) to deliver large spacecraft to cislunar space has generated significant advancements in the field of low-thrust trajectory design and optimization (Ref. 1). A focus of this work has been the development of a framework by which to design and optimize end-to-end trajectories to transit the initial capability of the Gateway lunar space station from a geostationary transfer orbit (GTO) to a 9:2 synodic Near Rectilinear Halo Orbit (NRHO) in the Earth-Moon L2 family (Ref. 2). Development of this framework highlighted the computational difficulties involved in the optimization of these trajectories due to the sensitivities inherent in the dynamical regimes, mission requirements (such as eclipse path constraints), and propulsive capabilities of low-thrust spacecraft. An example of this form of trajectory is shown in Figure 1.

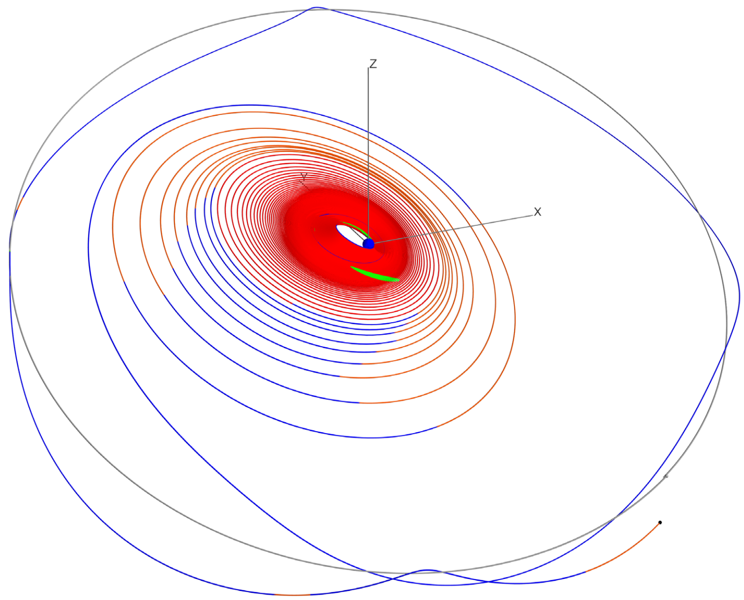

Figure 1.—Example reference trajectory of the Gateway Lunar Transit.

To execute a mission, it is critical to ensure that the trajectories developed under this framework are robust and resilient to variations in expected performance during flight. These variations can take the form of expected execution errors due to the inability of physical hardware to perfectly execute a commanded maneuver as well as wholesale departures from the reference trajectories due to unforeseen contingency or off-nominal scenarios. Many scenarios manifest themselves as missed thrust events (MTEs) which can be broadly defined as any time the spacecraft fails to execute a planned low-thrust maneuver. MTEs include situations where the execution of the maneuver is completely missed and no portion of the thrust arc is executed as well as partial executions where the maneuver is aborted or terminated earlier than planned. Regardless, any MTE of any duration causes the spacecraft to depart from the designed reference trajectory. Once the anomaly or cause of the MTE is resolved and the spacecraft can resume nominal operations, the vehicle must either target back to the existing reference or the reference trajectory must be redesigned. In the context of this analysis, MTE recovery may result in Gateway arriving at the target NRHO (Ref. 3) with some additional propellant consumption and possibly a delay to arrival time relative to the baseline mission.

Just as the development of the optimized reference trajectories for the Lunar Transit mission is computationally difficult and expensive to perform, so too is the assessment of these references with respect to MTEs. Built upon the optimization framework described in previous publications (Ref. 2), a methodology is developed and tested to determine the sensitivity of precomputed reference trajectories to MTEs at any point in the end-to-end mission. The results of a large-scale test of reference trajectories are summarized and used to adjust the nominal trajectory design to increase robustness to MTEs.

2.0 Background

Spacecraft represent highly complex integrated systems. To successfully operate, a wide range of hardware and software must work together seamlessly with an array of ground-based systems. The harsh conditions of space, imperfect nature of physical hardware, sheer complexity of supporting infrastructure, and the presence of human operators in-the-loop all conspire to make the perfect execution of any space mission extremely unlikely. When one then considers a cutting-edge exploration mission such as Gateway, which relies on many newly developed and previously un-flown technologies, the likelihood of completing a mission with zero missteps nears zero.

Safe-mode events, in which a spacecraft attempts to minimize the damaging impacts of an anomaly or out-of-spec operating condition by terminating nonessential functions, are not uncommon across legacy space missions (Ref. 4). During a safe-mode event the spacecraft will autonomously attempt to isolate any faults that may have occurred onboard by terminating all functions other than those necessary to keep the vehicle alive and communicating with Earth. This includes ceasing any thruster operations that are not necessary to maintain a power-positive and Earth-communications-conducive attitude, such as the execution of a low-thrust maneuver.

When a safe-mode event overlaps with a planned low-thrust maneuver an MTE occurs. The duration of the MTE is determined by the time required for the spacecraft to return to nominal execution of the mission. Given that a typical low-thrust mission is under thrust for days to weeks at a time it is likely that any safe-mode event will result in an MTE. For example, the Dawn mission operated thrusters over 90% of the time during regular cruise (Ref. 5). This probabilistic reality is reflected in recent experiences with JPL's Psyche mission, which experienced a lengthy MTE driven by an anomaly in the propellant feed system (Ref. 6), and the ESA's Bepi Columbo, which experienced a significant MTE that resulted in an 11-month delay to the spacecraft's arrival at Mercury (Ref. 7). Robustness to any deviation from planned thrusting is crucial for successful execution of a low-thrust mission.

Previous work by Imken, et al. (Ref. 4) describes the historical occurrence rates and durations of MTEs as a function of time-between-events. These rates take the form of a Weibull probability distribution describing the likelihood of encountering a safe-mode event as a mission progresses and a statistical forecasting of how long recovery from the event will require. This generalized empirical study includes data from 21 NASA/APL/JPL deep space missions. While Imken et al.'s work includes a variety of missions only two of the spacecrafts utilized SEP (Deep Space 1 and Dawn) and only one represents a lunar mission (the Lunar Reconnaissance Orbiter). Further, Imken et al.'s timeline of a safe mode event and return to nominal operations (see Figure 1 in Ref. 4) describes a science-focused mission, which Gateway is not. Thus, while the work of Imken et al. provides a reference point from which to start an analysis of safe-mode events, it is clear that a study of MTEs for low-thrust SEP trajectories to cislunar space requires significant development of a framework from the ground up.

3.0 Considerations

The design of the baseline Lunar Transit trajectory involves the development of many reference trajectories, spanning multiple days or even many weeks of potential launch dates. Each of these references is assessed against expected execution errors, performance dispersions, and MTEs to ensure that the identified launch opportunities are robust and executable in the face of likely perturbations. MTEs must be considered across the end-to-end trajectory as they may occur at any point during flight from the commencing of low-thrust maneuvers to arrival at the NRHO. To efficiently analyze the impact of MTEs of varying durations across the full range of time-of-flight for tens or even hundreds of reference trajectories, it is desirable to develop an automated and reliable analysis framework.

Achieving computational efficiency and solution reliability is often an objective of large scale optimization efforts but can be elusive when working with complex, high-fidelity problems. As described in (Ref. 2), the design of the Lunar Transit reference trajectories is computationally challenging and so the following details must be considered in the development of an MTE framework:

1. Objectives of the analysis
2. Design and architecture of the underlying reference trajectory and mission phases
3. Limitations of the selected trajectory design and optimization software

3.1 Objectives

In addition to the development of the framework to conduct the analysis, there are two desired outcomes of this missed thrust analysis. First, the Lunar Transit trajectory shown in Figure 1 will be capable of supporting some natural maximum duration MTE before it is no longer possible to reach the NRHO prior to the nominal arrival epoch. In this event, the underlying reference trajectory would need to be redesigned to support a later-than-planned NRHO arrival. This is largely contingent on whether or not the spacecraft is capable of recovering from an MTE and reaching the baseline low-energy ballistic NRHO approach trajectory given the acceleration capability of the low-thrust propulsion system. The maximum feasible outage duration is a function of when an MTE occurs in the mission, with some phases of the trajectory expected to be more sensitive to outages than others. Identifying this location, the maximum feasible duration, and changes that can be made to the design of the underlying references to increase this duration is important to increasing the robustness of the nominal mission.

The second objective is to identify the sensitivity of the Lunar Transit to MTEs in terms of the propellant mass (or Δv) required to recover from an outage of a given duration. To protect against a maximum-expected outage duration identified through engineering analysis, the propellant mass required to recover from an outage of that duration must be understood. Similar to the maximum feasible outage duration, the propellant required to protect against this maximum expected outage duration will be a function of outage location. This propellant sensitivity informs the sensitivity to MTEs at various points in the mission.

3.2 Design of the Underlying References

The Lunar Transit trajectory is designed in three primary mission phases, which have each been described previously Reference 1. The first is the Spiral Phase, beginning with Gateway's insertion into a highly elliptic Earth parking orbit similar to a geostationary transfer orbit (GTO). From there, Gateway's 50 kW class SEP system is used to climb out of the Earth's gravity well over a period of hundreds of days. This mission phase is characterized by near-continuous thrusting over hundreds of revolutions around the Earth during which a simple Earth-relative velocity-aligned thrust guidance method is used to maximize the rate at which orbital energy is increased. The Spiral Phase terminates once the vehicle reaches a semimajor axis of 130,000 km. With hundreds of days of thrusting in this mission phase alone, an MTE of modest duration (less than 30-days for example) is likely to represent a small fraction of the overall Spiral Phase. Additionally, due to the fixed control law and semimajor axis-based termination conditions, the impact an MTE will have on the Spiral Phase is primarily limited to a 1:1 slip in required time under thrust (a 1-day outage will require roughly 1-day of additional time to reach 130,000 km). An exception to this may be significant outages in the early periods of the Spiral Phase, where the right ascension of the ascending node (RAAN) and argument of perigee (AOP) of the spacecraft orbit can change relatively rapidly due to Earth J2 perturbations. Any errors introduced by this outage are rolled forwards into the Alignment Phase, where steered thrusting is used to correct back to the reference trajectory.

The Alignment Phase patches together the end of the Spiral Phase with the designated target interface necessary to approach the NRHO. This phase, typically consisting of 12 to 14 revolutions and less than 100 days time of flight, patches together the end of the Spiral Phase with the designated target interface necessary to enter onto the desired low-energy Ballistic Phase. Unlike the fixed velocity-vector guidance of the Spiral Phase, the thrust vector is allowed to align in any direction and vary as a function of time across the Alignment Phase. Additionally, unlike the near-constant thrusting of the first portion of the trajectory, coast arcs are included in the Alignment Phase and the fraction of the phase spent under thrust is typically 50% or less. The direction of the thrust vector and the location and durations of the thrust and coast arcs in this phase are allowed to optimize in order to maximize spacecraft mass at NRHO arrival. During this phase, lunar gravity becomes significant as the spacecraft enters the multibody dynamical regime.

The terminal Ballistic/Insertion Phase consists of a low-energy trajectory beginning at the final perigee passage of the Alignment Phase which naturally drifts towards the NRHO. Short, deterministic low-thrust maneuvers are applied on the final three of five revolutions about the Moon to wind the spacecraft onto the target orbit. These three maneuvers are typically each 1 to 2 days in duration with a total Δv expenditure of between 5 and 100 m/s. The vast majority of the Ballistic Phase, which is roughly 60 days in duration, is deterministically spent under coast although statistical trajectory correction maneuvers are expected to be necessary. A depiction of all three mission phases, highlighting the relative amounts of thrusting (red) and coasting (blue), is shown in Figure 2.

In the design of the reference trajectory, the terminal approach to the NRHO is the first mission phase to be designed. For a given launch date, a specific revolution of the target NRHO is selected and a low-energy approach trajectory is designed using a multiple-shooting method to define the Earth-centered target interface for the Alignment Phase to be constrained to. This preliminary Ballistic Phase solution is, in effect, just one trajectory within the reachable set stemming from the target NRHO and defined by the chosen architecture of one low-thrust maneuver during each of the final three revolutions about the Moon. When the final optimization problem is solved, this preliminary solution may vary slightly in shape and Δv cost and the precise location and timing of the ballistic interface (the final perigee passage of the Alignment Phase). These variations tend to be small, the epoch at which the ballistic interface may occur is typically constrained to variations of 1 or 2 days and the reachable set at this epoch may only encompass a handful of degrees of variation in RAAN, AOP, inclination, and true anomaly and a likewise restricted region of semimajor axis and eccentricity.

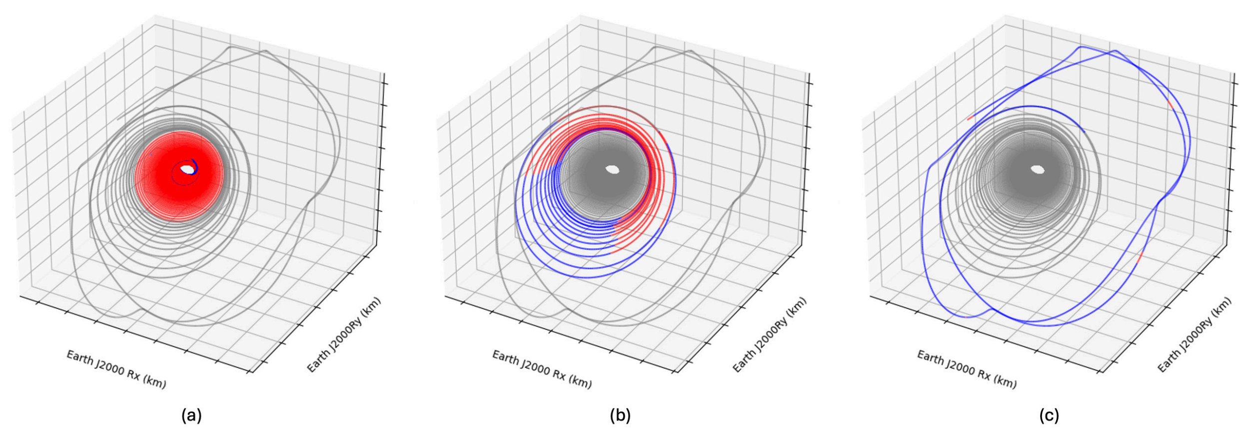

Figure 2.—Lunar Transit trajectory showing thrust arcs in red and coast arcs in blue, with the Spiral Phase highlighted in (a), the Alignment Phase in (b), and the Ballistic/Insertion Phases in (c)

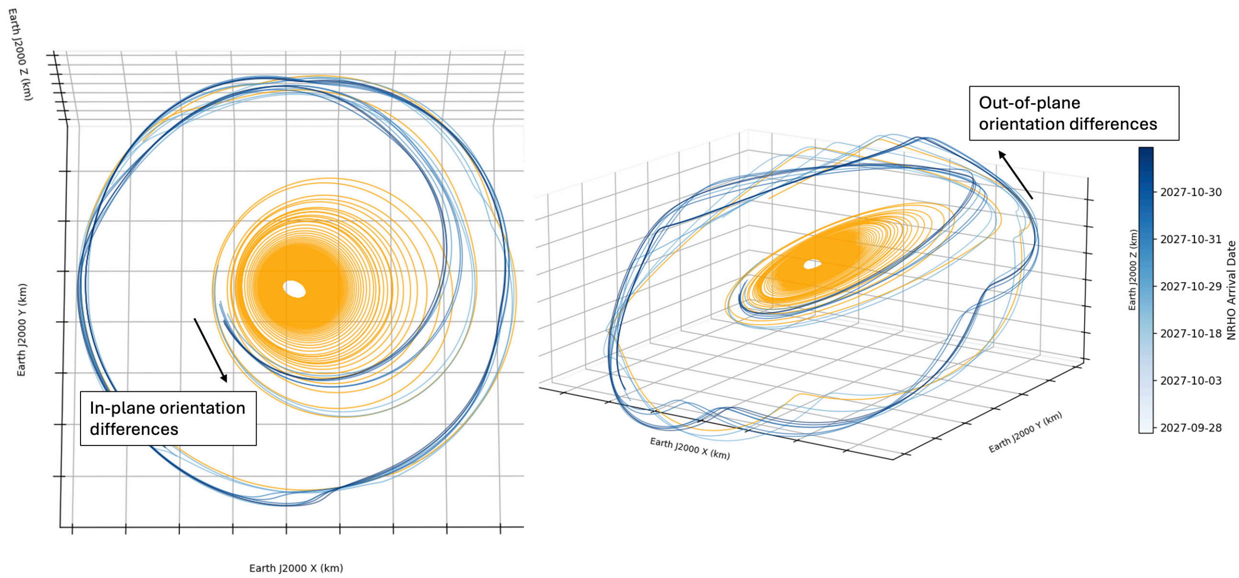

Figure 3.—In-plane and out-of-plane orientation differences between one Lunar Transit trajectory (orange) and trajectories arriving on four subsequent revolutions of the NRHO spanning 1-month (blues)

Should the spacecraft be unable to enter the reachable set of the NRHO a backup Ballistic Phase will be necessary to complete the mission. As seen in Figure 3, this is not as straightforward as one may assume due to the natural dynamics of the Earth-Moon system and the NRHO's place in it. Approach trajectories for one revolution of the NRHO may differ from those for subsequent revolutions significantly in the orientation of the interface (Earth-centered RAAN and AOP). While backup transfer opportunities do exist, it is preferable from both a propellant cost and operational complexity perspective to maintain the reference target. Each reference trajectory has been designed to marry a particular NRHO insertion with a particular launch from the Earth's surface (targeting a particular elliptical Earth park orbit) and so adjusting a trajectory to target a later NRHO insertion opportunity may result in increased propellant usage, sometimes significantly so. Additionally, in order to determine the robustness of the as-built trajectory, it is desirable to maintain the baseline Ballistic Phase target and not rearchitect the terminal phase of the mission. To this end, the maximum feasible outage duration mentioned above is defined as the MTE duration which causes the spacecraft to no longer be capable of reaching the reference Ballistic Phase entry target. This approach determines if the trajectory is robust enough to the maximum expected outage duration to maintain the ability to reach this interface without a wholesale redesign of the NRHO insertion (which would be undertaken in the event of a significant enough anomaly to exceed this MTE length).

3.3 Limitations of the Selected Design Software

All trajectory design and optimization activities for the Lunar Transit are conducted in Copernicus (Ref. 8), developed by NASA's Johnson Space Center. Copernicus is an exceptionally powerful high fidelity trajectory design tool used across the Orion, Human Landing System (HLS), and Gateway programs. Leveraging existing optimization algorithms such as SNOPT (Ref. 9), Copernicus enables the design of flight-fidelity spacecraft trajectories in complex multibody dynamical environments. This capability is leveraged extensively to design and optimize the Lunar Transit to minimize propellant consumption.

For flight-fidelity trajectories for SEP spacecraft such as Gateway, the trajectory optimization problem becomes quite complex and typically include:

1. Time-varying magnitude of the thrust vector (as power available varies according to range to the Sun)
2. Time-varying direction of the thrust vector (solved for as part of the optimization problem)
3. Maximum rate-of-change of the thrust direction (bounded by spacecraft capabilities)
4. Location and placement of optimal coasting (solved for as part of the optimization problem)
5. Required coasting for events such as eclipse passage (a constraint for solar electric spacecraft)
6. Path constraints such as minimum planetary fly-by altitudes (often determined by thermal conditions)
7. Mission requirements and constraints such as maximum time-of-flight

The above list is not exhaustive but drives much of the computational complexity involved in optimizing many-revolution spiral trajectories. For the Lunar Transit, the Spiral Phase consists of hundreds of revolutions most of which contain at least one cycling of the SEP system (e.g., due to an eclipse passage). Copernicus, unlike other trajectory design tools, is not an event-based framework where trajectories are propagated forwards in time until some termination condition is reached. Instead, the entire trajectory is both propagated and iterated upon simultaneously to achieve the desired feasibility and optimality conditions. This leads to a high incidence of convergence failures and, in cases where the trajectory does converge, the global minima is rarely found. For this reason, the fixed control law of the Spiral Phase eliminates the need to solve an optimization problem as only the duration of the mission phase must be solved. This duration is typically consistent across launch dates. So, a high-quality initial guess enables a simple differential correction scheme to converge upon the desired SMA termination within a handful of iterations.

4.0 Methodology and Approach

Given the considerations outlined above, any methodology for MTE analysis must be flexible enough to adapt to the demands and design philosophy of each mission phase. Maintaining multiple approaches to the missed thrust problem within a single framework enables this flexibility.

4.1 Framework Approach

The first approach uses a traditional Monte Carlo method to determine the duration and mission elapsed time (MET) at which the MTE occurs. Probabilistic distributions, such as those developed by Imken et al. (Ref. 4), are sampled from to determine the duration of the inoperability period of an MTE. The total duration of the outage from safing to the return to nominal thrusting is then determined by adding mission-specific assumptions for the length of the discovery delay and design cycle build phases of an MTE. Likewise, the MET at which an outage occurs is sampled from some probability distribution. Many samples of each distribution can then be run against one or more reference trajectories to develop a solution space covering some span of MTE durations across the full range of METs.

A second approach applies a more rigid case-matrix approach to MTEs similar to that used in the Dawn mission design (Ref. 10). Referred to as the Rolling Outage approach, this method is implemented by a-priori determining a fixed MTE duration corresponding to some outage length of interest (the maximum expected outage, for example). This outage is first applied at the start of the thrusting, then marched forwards in time in fixed time steps until the end of the trajectory or the end of a given mission phase. For instance, a 10-day outage can be inserted at MET +0 days, then stepped forward to MET +10, +20, and so on.

The final approach, referred to as Maximized Outage, leverages the optimization capabilities of Copernicus (similar to a method in Venigalla et al. (Ref. 11)). For a given MET location, a zero-duration outage is inserted into the trajectory. The optimizer, namely SNOPT, is then used to maximize the duration of the outage at that location. This approach can be applied with no other constraints, enabling the optimizer to find the maximum feasible outage duration at a given MET, or it can be combined with a mass constraint at NRHO arrival, setting an upper limit to the amount of propellant the optimizer has available to reach the NRHO. This is particularly useful to analyze the maximum recovery capabilities of a spacecraft that must operate within established propellant margins and reserves (due to the inconveniently finite nature of onboard propellant tanks). The location of the MTE can be determined either a-priori, sampled from some preconstructed matrix of cases, or sampled from some probability distribution.

4.2 Benefits and Limitations

Each of the described approaches carries its own unique set of advantages and disadvantages (summarized along with the methods themselves in Table 1). These considerations make each method more suited for one or more mission phases, but not all three.

The Rolling Outage approach is the simplest of the three, allowing the mission designer to directly analyze an MTE duration of interest at particular locations of interest, offering a targeted analysis capability not available with stochastic or probabilistic methods. By tailoring the input case matrix, the size of an analysis scan can be contained to only the specific set points needed to inform a specific data product such as the propellant required to protect against the maximum expected outage. This is critical to the deployability of the framework given the computational expense of computing a single end-to-end solution to the Lunar Transit (which may take in excess of 60 min to solve and optimize). Large scan sizes, while valuable for the quantity of data they generate, are often wasteful in the clock time required to complete when compared to the results gained from smaller, more targeted analyses.

TABLE 1.—OVERVIEW OF FRAMEWORK METHODS

	Rolling outage	Maximized outage	Monte Carlo
MTE duration	Fixed	Optimized	Sampled from distribution
MTE location	Predetermined matrix	Predetermined matrix or sampled from distribution	Sampled from distribution
Advantages	• Smooth results • Small scan sizes • Directly assess durations and locations of interest	• Smooth results • Small scan sizes • Directly determines upper bound	• Extensive results space • Resolution includes full range of potential MTE durations and locations • Can be seeded with vehicle-specific probability distributions
Disadvantages	• Low resolution • No guarantee of feasibility/requirement compliance for lower MTE durations • Does not determine maximum feasible outage duration	• Low resolution • No guarantee of feasibility/requirement compliance for lower MTE durations • A-priori decide the propellant mass to allow in recovery • Relies on an optimizer	• Large scan sizes • Noisy results due to numerical reliability of software • Quality of results depend on quality of input distributions • Significant fraction of results may not be valuable

Similarly, the Maximized Outage approach can limit the size of scans by directly analyzing the upper bound of possible outage durations. By default, the Maximized Outage method will use the optimizer to determine the maximum feasible outage at a given epoch, but it can also determine the maximum executable outage given a propellant mass constraint as discussed above. Finding these upper bounds via a Monte Carlo method would be computationally expensive and potentially take one or more orders of magnitude more samples to determine, significantly increasing the time required. Directly determining the upper bound additionally informs a key data product, the maximum feasible outage, and does so precisely as the optimizer will converge within predefined tolerances as opposed to a probabilistic method which may get close to the exact value of the upper bound but never meet it exactly.

Both the Maximized Outage and Rolling Outage approaches suffer from similar disadvantages, mostly stemming from the low resolution of the resulting data set. The Rolling Outage approach, by definition, only analyzes a single outage duration per run while the Maximized Outage only analyzes to the upper bound. This means that both methods lack insight into the impact of, and sensitivity to, outages of lesser durations. For the Lunar Transit mission, this means that additional analysis must be completed to ensure that the designed reference trajectories remain requirement compliant for outages of lower durations than what is analyzed. In general, the compliance of longer duration outage recovery solutions can be extended to shorter duration outages occurring within the same trajectory. This assumption is not necessarily valid, however, when working within a chaotic dynamical environment such as cislunar space, requiring higher resolution results to validate solutions.

The Maximized Outage approach implicitly relies on an optimizer. This makes it uniquely poorly suited for use in the Spiral Phase, where the use of an optimizer is purposefully avoided due to the numerical limitations of working with a many-revolution spiral trajectory. The Rolling Outage approach avoids this issue and by using a course grid size for MTE location, it can step through the hundreds-of-days-long Spiral Phase with a modest number of samples. This is particularly advantageous for use in the Spiral Phase, as the phase is so long that the sensitivity to MTEs changes very slowly with advancing epoch (discussed further in results). The reliable nature of the Spiral Phase alleviates the resolution issue, making verification of shorter outages relatively straightforward to accomplish.

Unlike the other two, the Monte Carlo approach relies on a supplied probability distribution to seed both the duration and the location of an MTE. This comes with an immediate benefit; any suitably large number of samples will cover the full range of expected outage durations as well as the full range of possible occurrence epochs. The resulting data set will represent a fine resolution and directly analyze the feasibility and compliance of short and long outages. The underlying probability distributions can also be directly developed with spacecraft-specific reliability and probabilistic risk assessment (PRA) data such as mean time between failures or expected recovery procedures. This offers an opportunity to conduct representative analysis of how the vehicle is expected to perform. Conversely, given the inherent difficulties in meaningfully predicting the performance of a brand-new spacecraft flying brand new systems, these inputs are more likely to be poor representations of flight.

Many of the advantages of the Monte Carlo approach are double edged swords and can quickly become pitfalls without careful application. While a suitably large sample size will generate an extensive and complete data set, the computational requirements of solving those samples may be prohibitively large. Additionally, some natural level of noise is typically present in Lunar Transit results due to the numerical limitations of applying Copernicus to low-thrust spiral trajectories. This noise can cloud Monte Carlo results, a disadvantage when compared to the other two methods which typically produce more smooth result sets due to the smaller scan sizes as well as the performance of the optimizer in the Maximized Outage approach. As discussed above, sensitivity to MTEs of a given duration does not typically vary significantly with small changes in occurrence epoch. For a fixed or similar epoch, results are also relatively insensitive to small variations in the duration of an MTE (a 10- and 11-day outage require similar amounts of recovery propellant). This means that much of the result space generated with a Monte Carlo approach will not be valuable from an analysis perspective. In practice, a large scan is no more valuable than a smaller one if it does not result in additional understanding of the underlying problem. And the additional time spent executing a large scan is time which cannot be spent on other analysis.

Given these relative advantages and disadvantages, each approach naturally lends itself to application to a particular mission phase. The simple and broad-search nature of the Rolling Outage approach makes it well suited for analysis of the Spiral Phase. The Maximized Outage method, with its unique ability to directly determine upper bounds for MTE feasibility, is applied primarily to the Alignment Phase, where sensitivity to outages is expected to be highest due to the criticality of the ballistic interface. Finally, the terminal NRHO insertion is left to the Monte Carlo method, where the relatively limited scope (thrust arcs of 1 to 2 days duration across ~30 days time of flight) enables large numbers of samples to be run without the use of prohibitive computational resources.

5.0 Framework Implementation

The methodologies described above are implemented in the form of a suite of python-based tools. This code library includes a variety of scripts and functions used to extract information from reference trajectories and their time histories, manipulate and solve trajectory models, and collect resulting data into output products. In total, more than 3,000 lines of code make-up this analysis package which has been integrated into a global mission analysis library which forms the common tool package for the Gateway Lunar Transit mission designers and includes the reference trajectory development framework described in Reference 2. This global framework allows both reference development and subsequent dispersion and perturbation analysis efforts to be conducted using the same primary functions and methodologies, increasing reliability of results across the board.

The base capability of the missed thrust framework is the ability to apply an outage, a forced coast, of a given Δt at a particular MET within the transit. For each reference trajectory, tabular time history data is generated in the form of a csv file. This time history is queried at the desired epoch for spacecraft state and thruster performance information. This information is used to create a control point representing a state on the reference trajectory. A forced coast whose duration corresponds with the desired outage is then applied at this control point. For Maximized Outage solutions, the Δt of the forced coast is initially set to zero and maximized through use of the optimizer. Following this Δt, the remainder of the Lunar Transit is reoptimized to reach the NRHO. The reference trajectory forms the initial guess for this solution so that the resulting dispersed trajectory maintains the same structure and mission phase construction as the baseline. Should an MTE occur in the middle of a thrust segment (such that the thrust arc will continue once the outage ends), the control history of that segment is mapped forwards in time so that the initial steering law of the reference is maintained. All trajectory segments prior to the control point are frozen, meaning that the resulting optimization process cannot retroactively optimize the entire trajectory to account for the presence of an MTE. The only portions of the trajectory that are allowed to vary are those occurring after the end of the forced coast.

The benefits of initializing in this manner are two-fold; the resulting post-MTE trajectory can be converged more easily and reliably as the initial guess is, itself, a converged high-quality trajectory and the resulting post-MTE trajectory is likely to remain in the same solution family as the reference. As the optimizer can often find local minima in neighboring solution families (instead of the same family as the baseline), generating an entirely new initial guess for the post-recovery trajectory can easily result in a large difference in propellant usage between the recovery and baseline trajectories (Δm). Using the reference as an initial guess has proven effective at limiting the occurrence of these neighboring minima and minimizing the Δm of the solution.

Once the process, outlined in Figure 4, is complete, it can be repeated for the next epoch and coast duration of interest. For each case, a new ideck is created (each a copy of the reference trajectory file) and a new control point is populated from the associated time history file. This entire process can be paralyzed and deployed on advanced computing resources, allowing many trajectories to be modified and solved simultaneously. By leveraging these resources, thousands of missed thrust cases covering tens or hundreds of reference trajectories (each corresponding to a different potential launch date) can be analyzed within several days of clock time. Usage in the summer of 2025 resulted in more than 16,000 missed thrust samples, requiring more than one calendar years' worth of CPU hours, yet completed in less than 2 weeks of runtime.

Once trajectories have been solved, output information is collected for the converged results. These results take the form of time history data, eclipse durations and occurrences, and Δm and Δt information describing the additional propellant and time required to recover from a given MTE. This entire process is automated such that the mission designer need only provide a directory which contains the reference trajectories to be analyzed, the MTE duration and interval step to be used for Rolling Outage analysis, and the parameters for Maximized Outage occurrence epoch sampling. The automated analysis process (described in Figure 5) then steps through these reference trajectories, applying the specified number of MTE samples according to the provided parameters, then collecting results data into a series of csv output files.

This automated approach has revolutionized the missed thrust analysis capability of the Gateway Mission Design team. Previous efforts to assess the impact of MTEs has focused on a single reference trajectory and has required significant amount of analyst time and effort to complete. While the development of the framework described here required the better part of 6-months to complete, the result is the ability to rapidly and reliably analyze thousands of MTEs in a matter of days. This is a critical capability and one that will be central to the verification of flight trajectories. Although thousands of samples across hundreds of references will not be required for flight, the ability to rapidly analyze dozens of samples across a handful of potential launch dates will be necessary prior to launch. This framework enables that capability in a streamlined, rapid, and robust package.

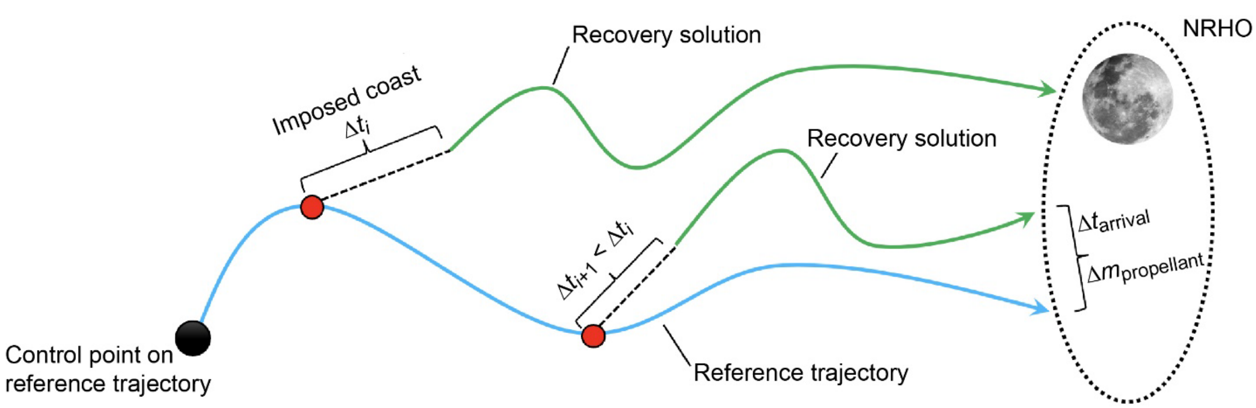

Figure 4.—Missed thrust analysis methodology.

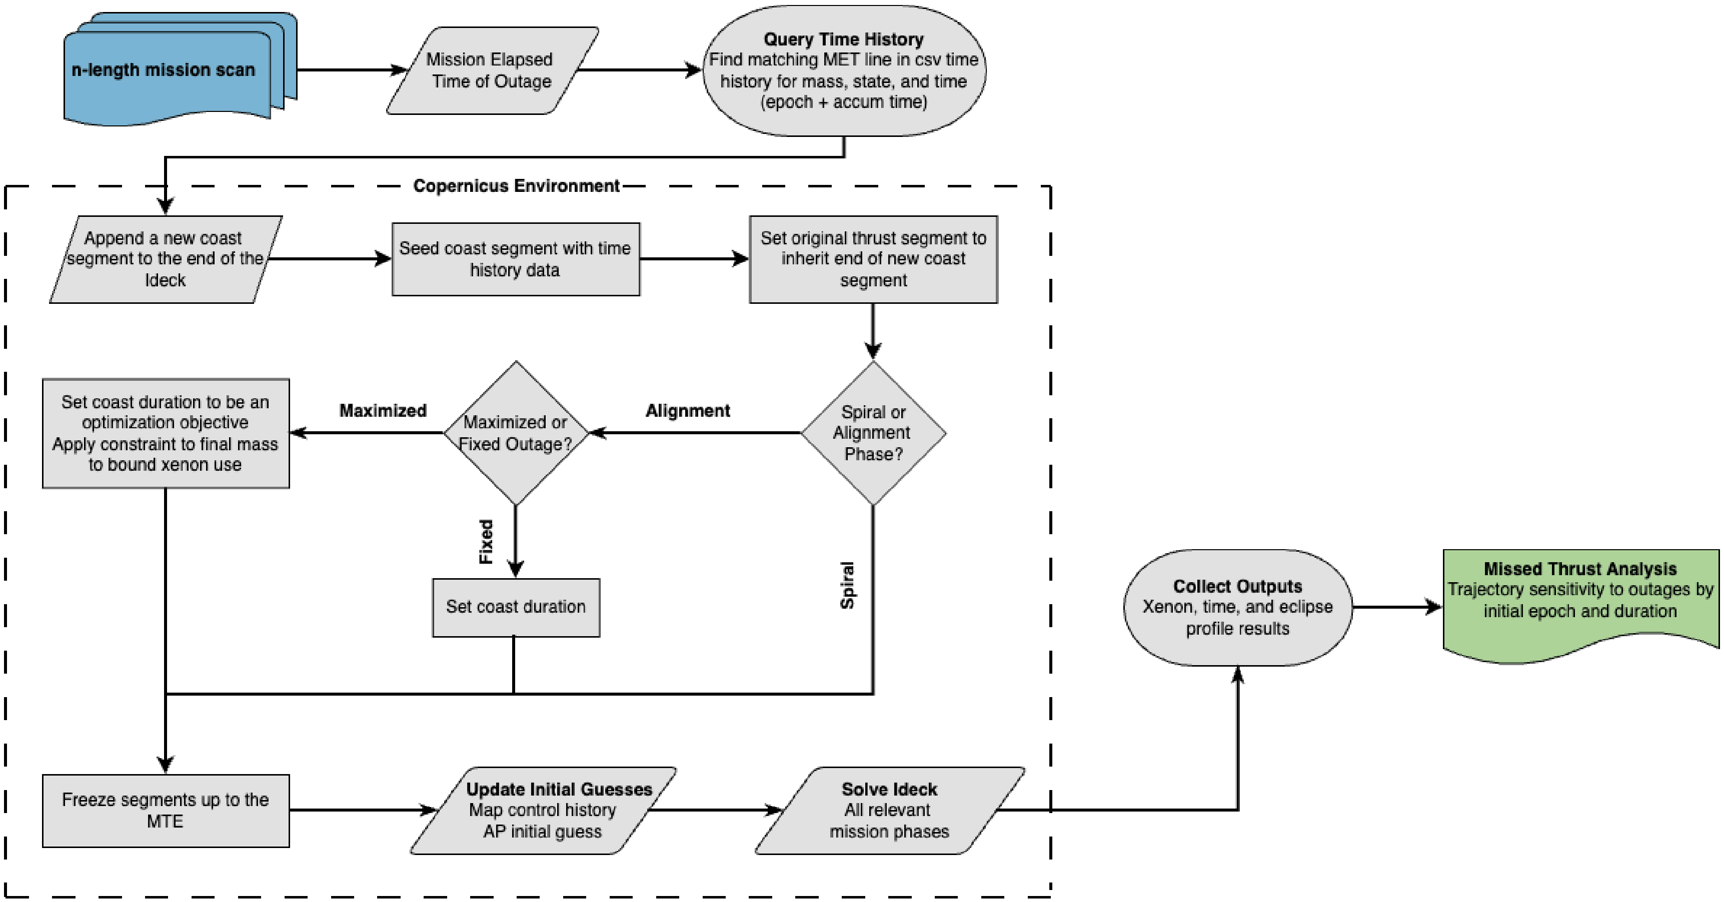

Figure 5.—Missed thrust analysis process.

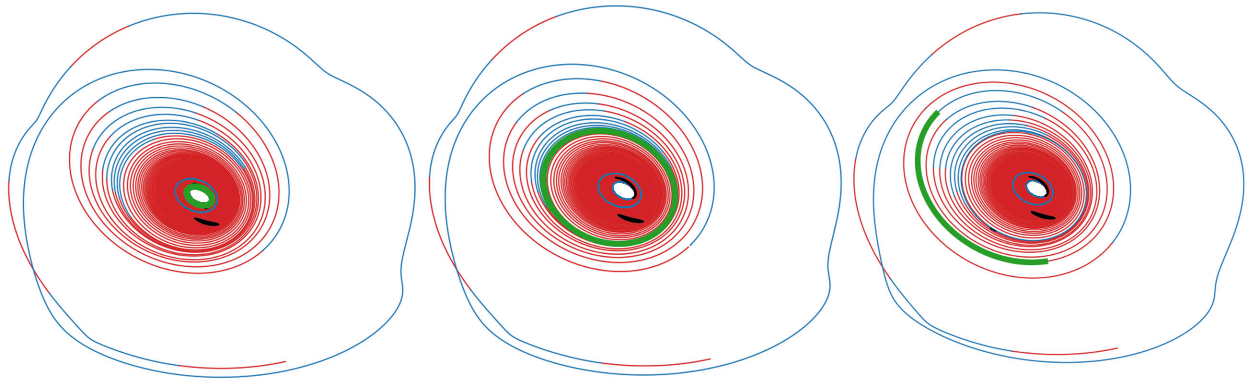

Figure 6.—A 7-day missed thrust event (green) applied to three different epochs in a reference trajectory, MET +49 (a), MET +329 (b), and MET +362 (c) with thrusts shown in red and coasts shown in blue.

6.0 Analysis

Results are split into three sections according to phase and methodology. Two samples of trajectories were used for testing: A) the most xenon-efficient trajectory per day across roughly a year (343 total trajectories) and B) a smaller subset used of 25 trajectories used for varying MTE length in the Spiral Phase and recovery xenon in the Alignment Phase (Group B is contained within Group A). Group A was tested with the goal of examining 7-day MTEs spread across each trajectory to determine sensitivities due to varying initial orbit orientation as well as seasonal differences. Group B was tested against three MTE lengths in the Spiral Phase (7-, 14-, and 21-day MTEs) and six maximum recovery xenon limits in the Alignment Phase.

For MTEs occurring in the Spiral or Alignment Phases, finding a recovery solution relies largely on the ability of the Alignment Phase to adjust such that its rendezvous with the ballistic lunar capture orbit, to the destined apolune, remains feasible. Since the Alignment Phase consists of alternating coasting and thrusting arcs, adjustments may come in the form of new suboptimal steering along with variations in the durations of the thrusting and coasting segments.

Sensitivity to MTEs can be visualized directly in the trajectory itself, as shown in Figure 6. At three different locations in a single reference trajectory (MET +49, +329, and +362 days) a 7-day MTE (green) is applied as a forced coast, and the resulting dispersed trajectory is reconverged to the same target NRHO apolune. The overall structure of the underlying trajectory is maintained in all three examples.

6.1 MTEs in the Spiral Phase

The Spiral Phase utilized the rolling outage method applying forced coasts evenly throughout this phase. Figure 7 shows the nondimensionalized additional xenon required to recover from 7-day MTE tests as a function MTE start across Group A trajectories, representing over 16,000 converged trajectories. The most striking feature of these results is the near constant mean cost of the recovery solutions, showing an insensitivity to when the MTE occurs. Regardless of when the 7-day MTE began, the average recovery trajectory used about 30% of additional xenon to reach the original target apolune. This behavior indicates that the Rolling Outage approach is well suited for use in the Spiral Phase, and that a course step size can be used in analysis without reducing the fidelity of the results

This constancy of the average recovery xenon cost is likely due to the underlying trajectory design. Since each Spiral Phase is solved to reach the same semimajor axis, MTEs will primarily perturb the right ascension of the ascending node (RAAN) and argument of perigee (AOP) and time at which the phase ends. These perturbations all move in single directions. J2 precesses RAAN and AOP in the same directions regardless of MTE, and MTEs will always push the end of the new Spiral Phase later in time. Therefore, recovery from MTEs in the Spiral Phase should be expected to hover around a similar cost for additional xenon used for most trajectories.

A negative value in Figure 7 means that the recovery solution used less overall xenon than the nominal trajectory. While the nominal trajectories of this dataset are meant to represent the optimal trajectory each day, these results highlight the difficulties in solving long low-thrust transfers in Copernicus. Primarily, the difficulty of finding solutions which represent the global minimum not simply a local minimum in the overall optimization problem. This is also reflected in the spread of the recovery trajectories additional xenon costs. The wings of the 2-σ distribution show wide variation in the costs of recovery solutions. One cause of this noise is the variation, caused by MTEs, in the ending state of the Spiral Phase.

Extended Spiral Phase MTEs of 14 and 21 days were performed on a smaller subset of 25 trajectories due to the long computation time. MTEs longer than roughly 21 days were found to require modification of the underlying ideck structure. Figure 8 shows the additional xenon used cost to recover these trajectories as a function of MTE start. Like the larger Group A results, a similar trend of roughly constant additional xenon costs can be seen across MTE starts. For each additional 7 days of missed thrust the additional xenon cost rises by roughly 0.25%.

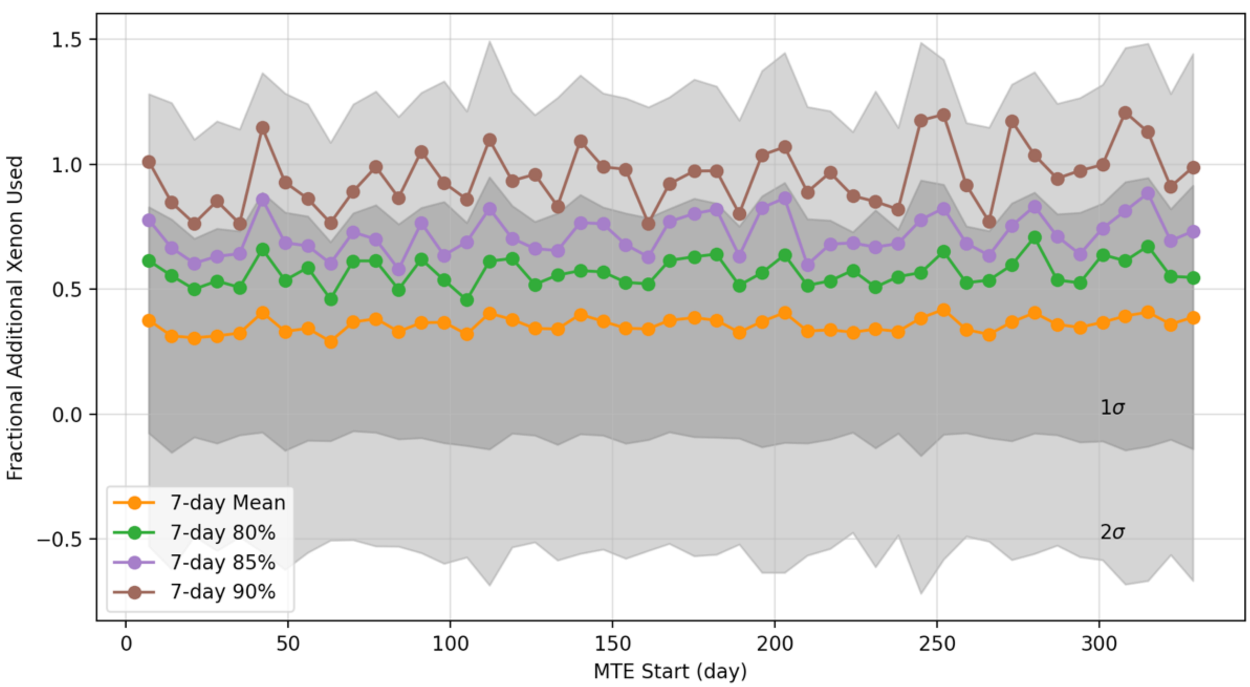

Figure 7.—Additional xenon used for recovery from 7-day forced coast rolling outages in the Spiral Phase as a function of the start of the MTE in mission elapsed time. These results used 343 nominal trajectories to generate over 16,000 MTE solution trajectories. Colored curves show the mean, 80th, 85th, and 90th percentiles of results that are binned every 7 days. Shaded regions overlay 1- and 2-σ bounding regions.

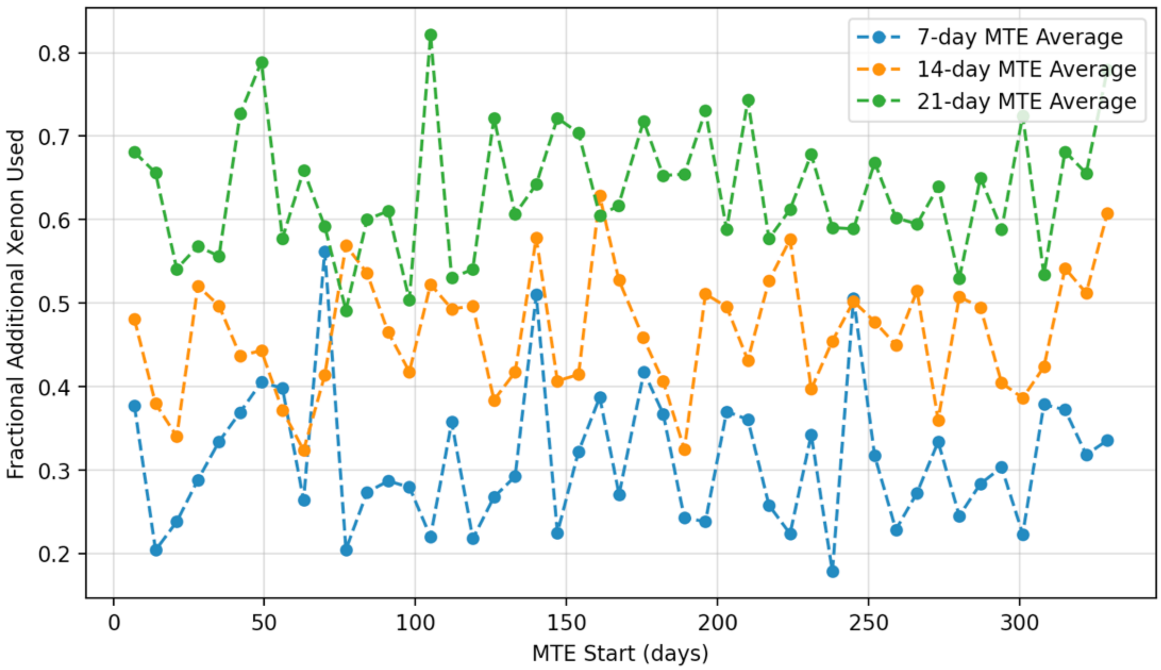

Figure 8.—Additional xenon used to recover from MTEs across the Spiral Phase as a function of MTE start in mission elapsed time. Each curve represents 7-day binned average values across a subset of 25 nominal trajectories colored by the length of the MTE. In each case, the indicated MTE length was applied every 7 days across each Spiral Phase, for consistency.

6.2 MTEs in the Alignment Phase

MTE tests in the Alignment Phase use the maximized outage method. Like the Spiral Phase, MTE starting times correspond to a rolling frequency that picks out MTE starts across both phases. The end of this phase marks a critical stage in the overall trajectory when the last apogee allows for a ballistic transfer into a lunar orbit. Robustness to MTEs in this phase is critical for the success of the mission. While MTEs in the Spiral Phase have an abundance of time, comparatively speaking, to be "corrected" and Insertion Phase MTEs always results in targeting a downstream apolune arrival, that is not the case in the Alignment Phase. Each target apolune arrival should be thought of as extending from insertion into an NRHO to the transition point between Alignment and Ballistic Phases. If entry onto a specific Ballistic Phase is missed, there is no guarantee that the orbit can be re-aligned to rendezvous with a later apolune.

Before examining the aggregated results, Figure 9 shows MTEs across an example Alignment Phase of a single reference trajectory that shows typical behavior that is otherwise difficult to see in the combined results. Each solution family is colored by the reference trajectory thrust arc in which the MTE occurs. From these results it can be seen that as the time remaining in the Alignment Phase decreases (from beginning near MET +335 days to the end of the AP near MET +420), the maximum outage duration also decreases. No missed thrust recovery margin (MTRM) exists near the end of the Alignment Phase, and results tend to trend upwards within the bounds of a single thrust arc. This indicates that partial maneuver executions are more beneficial than complete failures to execute any portion of a thrust arc.

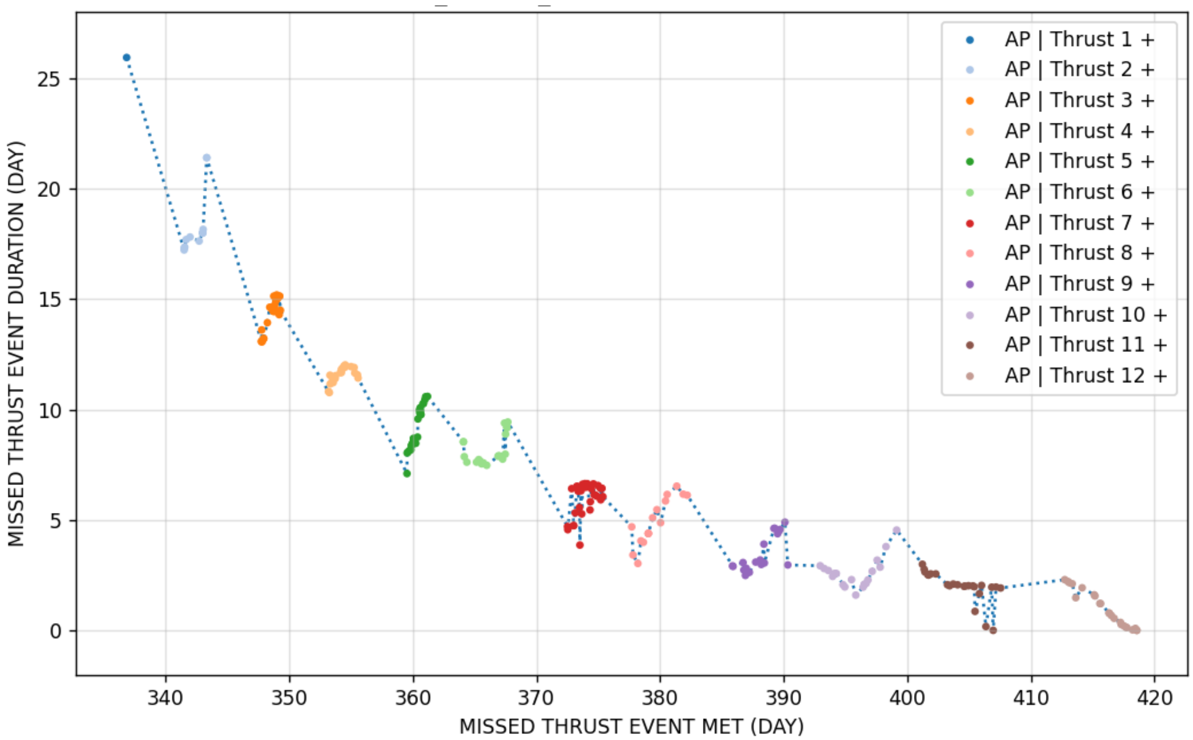

Figure 9.—Example Alignment Phase results using the maximized outage method. For each epoch (presented in mission elapsed time), an optimizer is used to determine the maximum possible outage for which a viable solution still exists within a given mass of propellant used. Results are colored by the thrust arc in which the MTE occurs (there are 12 thrust arcs in the Alignment Phase of this trajectory).

Results from the larger Group A maximized outage scan show that, like the Spiral Phase solutions, AP solutions can vary significantly. The plot of MTE length versus MTE start in Figure 10 shows that solutions overall follow a stair step pattern of decreasing MTE length as MTEs start later in this phase, like the single trajectory example. However, it also clear that some trajectories can accommodate MTEs better than others. The rough groupings of points in Figure 10 correspond to thrusting arcs of the Alignment Phase. Near the 7th AP thrusting arc, an MTE Start of ~380 days, the average MTE length dropped below the 7-day target threshold and continued to drop until the end of the phase. This is in part due to the fraction of time spent coasting in the nominal AP, which gives an MTE solution more flexibility to expand maneuvers, and in part due to lunar gravitational interactions. Resonances between the orbital periods of the nominal trajectory and lunar ephemeris result in very efficient Alignment Phases as the gravitational interaction naturally raises the semimajor axis (e.g., see (Ref. 12)).

In the Alignment Phase, trajectories in Group B were solved six times where the upper limit of usable xenon to recover was increased in increments of 10 kg. The mean MTE lengths are plotted as a function of Alignment Phase thrust arc in Figure 11. Similar to the previous results, the same decline in MTE length is seen in each family of solutions. Increasing the xenon limit provides the largest gains in MTE length early in the Alignment Phase and steadily converges to zero by the end. Even from this small subset, it is clear that if the temporary goal of maintaining robustness to MTEs of 7 days in duration is to be met, changes in the nominal trajectory design are needed. Ideally, these changes would not only raise the MTE length above the 7-day target but also allow for small increases in the xenon limit to have a more significant change across more of the Alignment Phase.

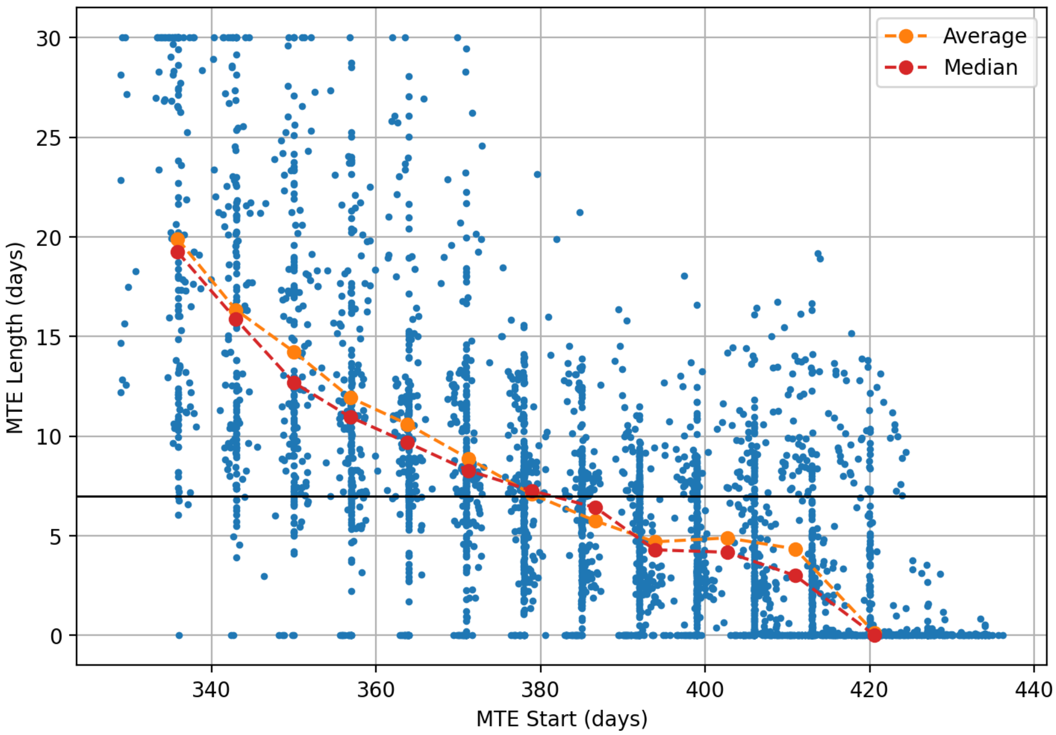

Figure 10.—Group A Alignment Phase results showing MTE length as a function of MTE starting time. Two colored curves show the 7-day binned averages and medians across each thrusting arc. The solid black line shows the target MTE length threshold of 7 days. Each MTE above could use a maximum of 60 kg of additional xenon to maximize the MTE outage duration.

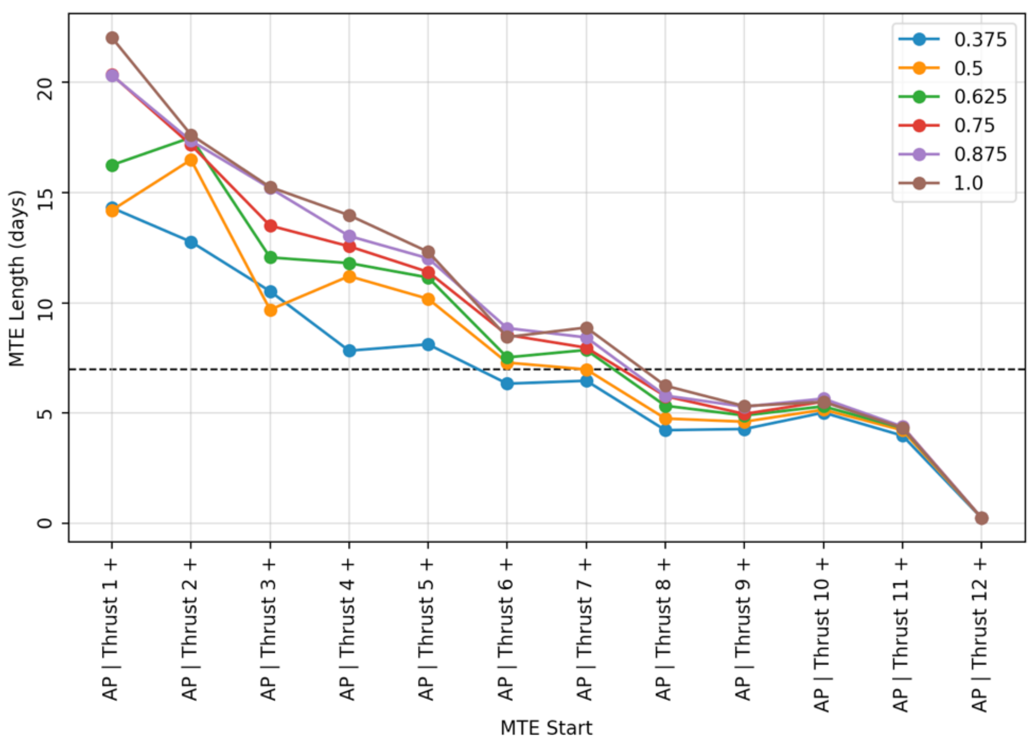

Figure 11.—Group B Alignment Phase results showing MTE length as a function of MTE starting time. Each color represents an upper limit for fractional additional xenon that could be used for recovery in the maximized outage method.

6.3 MTEs in the Insertion Phase

MTEs in the Insertion Phase are assessed according to the Monte Carlo method outlined above and follow the methodology described previously in References 13 and 14. For a given trajectory, the starting time of an MTE is sampled from one of the three deterministic thrust arcs in the Insertion Phase. The trajectory is modified to include an MTE of the sampled duration, then reconverged to insert into the NRHO three revolutions following the end of the MTE. In this manner, three recovery thrust arcs are used to correct back to the NRHO, with insertion occurring at or near apolune (insertion is constrained to occur at a lunar orbital radius of greater than 60,000 km).

By definition, the selected MTE recovery strategy results in delayed insertions into the NRHO relative to the reference trajectory (Figure 12). With an orbital period of roughly 6.5 days, this means that MTEs in the Insertion Phase will result in the spacecraft arriving to the NRHO ~6.5, ~13, or ~19 days later than the reference depending on the maneuver in which the MTE occurs. For MTEs longer than 6.5 days, an entire revolution (or more) is spent under coast, meaning that recovery, and thus NRHO insertion, is further delayed by one orbital period. As seen in Figure 12, this means that NRHO insertion can be delayed by over 1 month for MTEs in excess of 12 days duration.

As shown in Figure 13, the cost of a given MTE is primarily a function of the length of the outage. Longer MTE durations result in greater deviations from the reference trajectory, with MTEs in excess of 12 days causing the spacecraft to miss not only the maneuver in which the MTE initially occurs but also the subsequent maneuver on the next revolution. This results in a greater amount of thrusting required to correct the trajectory back towards an NRHO rendezvous, and thus higher propellant consumption. While propellant costs are directly correlated to the length of the experienced MTE, they are largely agnostic to the epoch at which the MTE occurs. The results shown in Figure 13 present fractional additional xenon usage as a function of the start time of the MTE, covering the entirety of the Insertion Phase across a subset of 21 Lunar Transit missions. Propellant costs are similar across the entire range of occurrence epochs, with no strong correlation visible.

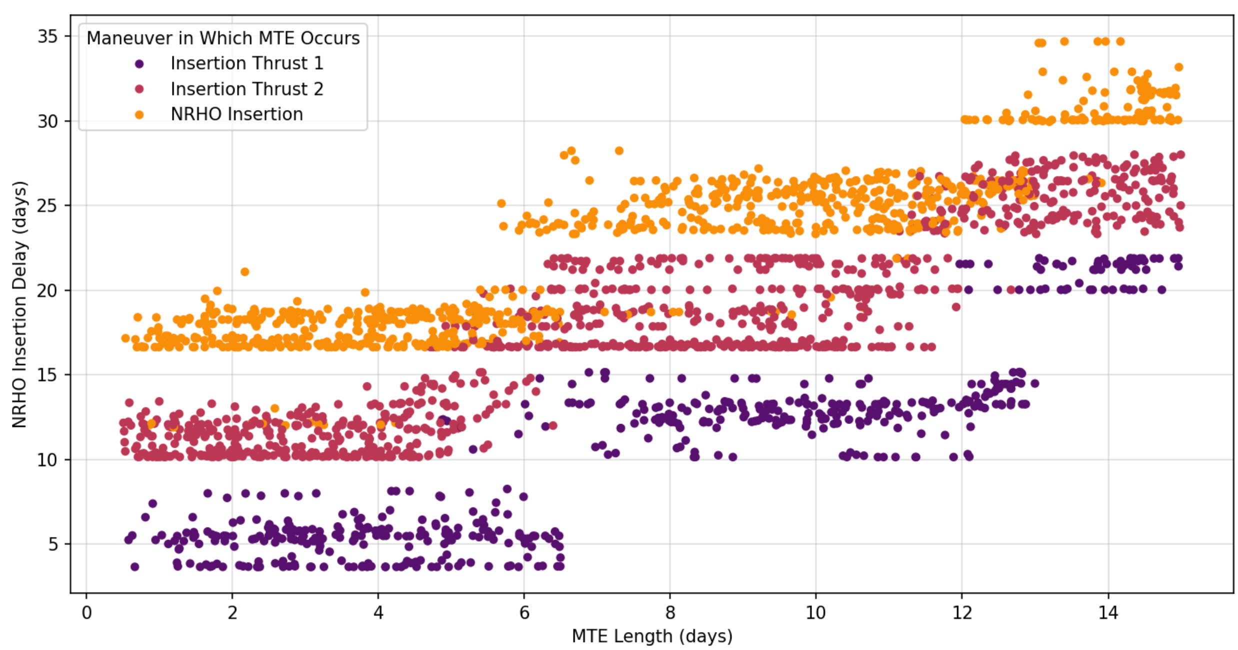

Figure 12.—NRHO insertion delay as a function of Insertion Phase MTE duration, colored by the maneuver in which an MTE occurs

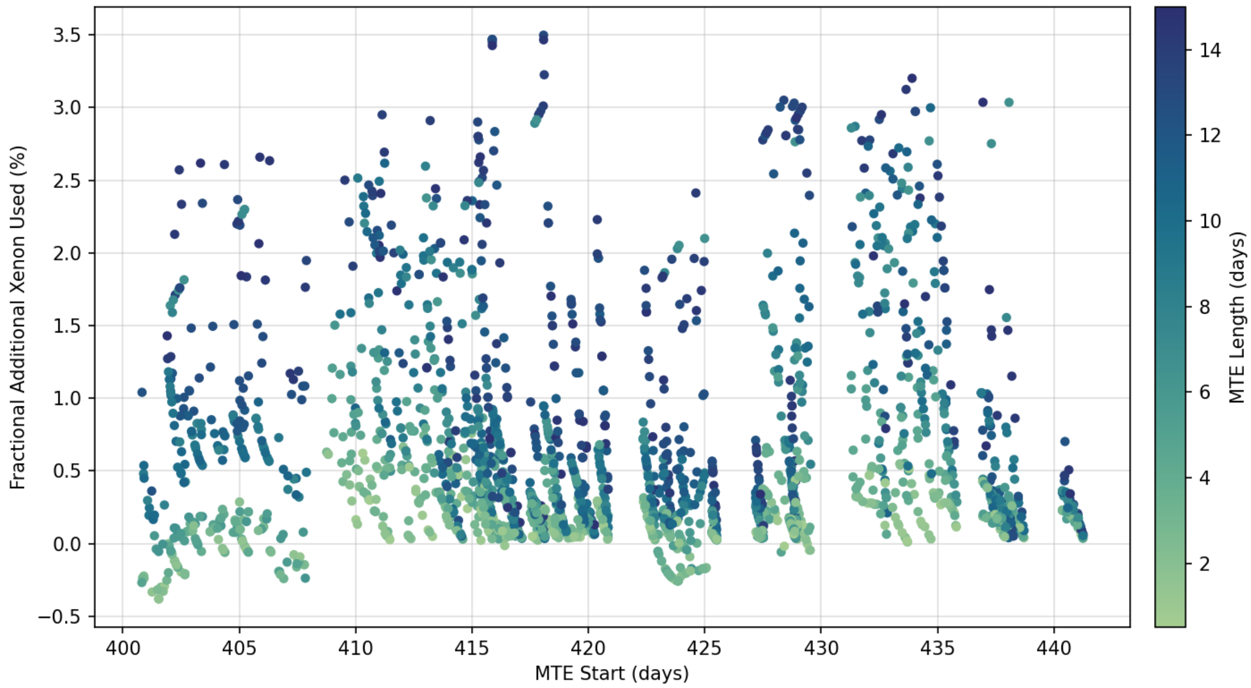

Figure 13.—Nondimensionalized additional xenon used as a function of MTE start time, colored according to the length of the MTE

6.4 MTEs Impacts to Eclipses

MTEs allow for orbital perturbations (such as those due to Earth J2 effects) to precess the RAAN and AOP of the orbit for an additional period of time. Changes in both of these orbital elements shift the angle between the orbit plane and Earth-Sun vector. This results in changes to when and how the Earth's shadow crosses the trajectory. Early in the Spiral Phase when perigee altitude is low, these perturbations have a greater impact on a trajectory and cause larger downstream shifts in the eclipse profile when MTEs occur early in this phase.

For the Group A dataset, Figure 14 shows the residual of the maximum Earth eclipse between nominal and corresponding MTE solutions. A negative value means that the maximum eclipse decreased in duration, and vice versa. The distribution shows that the maximum eclipse duration varied symmetrically in direction with most changes very near zero. An average of the absolute value of residuals showed that 7-day MTEs typically varied the maximum Earth eclipse duration by 13 min. The outliers on this distribution are the direct result of new Earth eclipse seasons moving into the late Alignment Phase. These eclipses tend to be prolonged in duration due to the long orbit period near the Alignment and Ballistic transition.

In general, MTEs lead to an increased number of Earth eclipses relative to nominal trajectories. When an MTE in the Spiral Phase occurs, the downstream eclipse profile is shifted to lower semimajor axes where the orbit period is shorter. MTEs that start earlier in the Spiral Phase led to the highest increases in eclipse counts while later MTEs lead to smaller increases since most eclipses have already occurred. This is shown in Figure 15 where the color axis refers to MTE start time. Figure 15 also shows an interesting relationship between the residual eclipse counts and the initial RAAN of the trajectory. Transfers to the NRHO fall into two RAAN groups based on the dynamics at the ballistic lunar capture stage that trace back to the initial RAAN of the Spiral Phase. Trajectories in the group centered near 55° typically have higher eclipse counts in their MTE recovery solutions than trajectories in the group centered near –10°. This is due to the natural precession of the orbital RAAN across the Spiral Phase which rotates the orbit RAAN closer to that of the ecliptic plane (and thus, Earth's shadow).

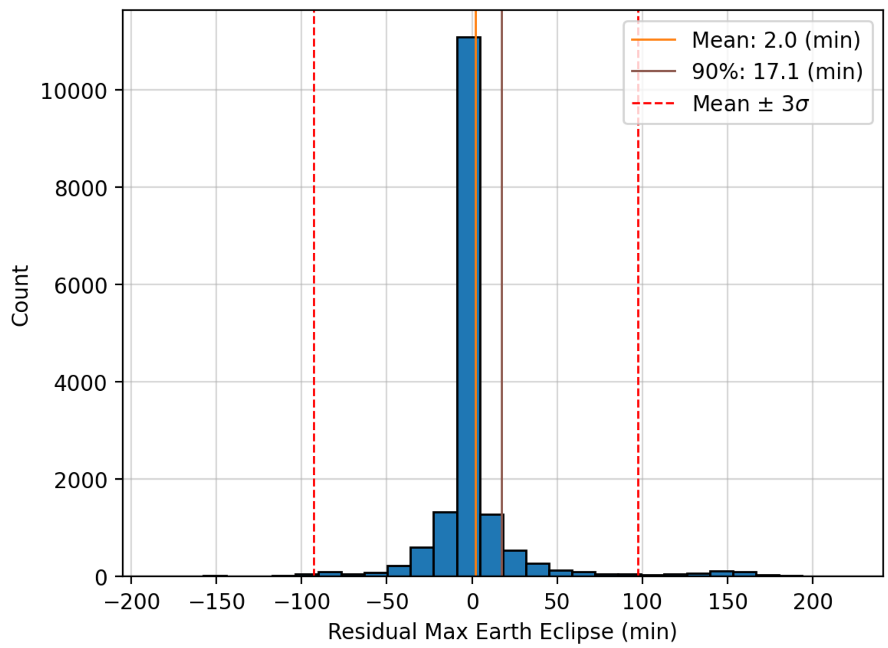

Figure 14.—Histogram of Group A residual maximum Earth eclipses (difference between maximum Earth eclipse duration encountered in the recovery solution and maximum Earth eclipse in the nominal trajectory). Vertical lines represent the mean, 90th percentile, and mean ± 3-σ points.

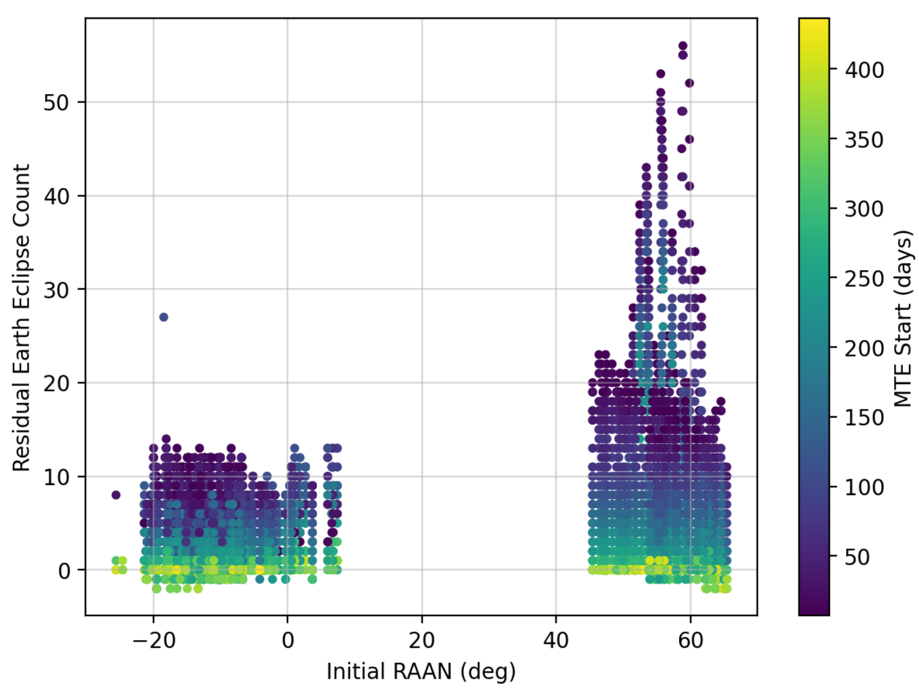

Figure 15.—Plot of the residual Earth eclipse counts (difference between MTE solutions and corresponding nominal trajectories) as a function of the initial RAAN of the trajectory for Group A. Points are colored by the MTE starting time.

7.0 Forward Work

The presented results show the capability of the developed framework, with more than 16,000 samples assessed across 340 reference trajectories in a single scan. Much of the planned forward work involves broad application of this capability to existing and future reference trajectories. While a full year's worth of references, as presented here, is not indicative of analysis needed to validate a single multiweek launch period, the ability to assess thousands of samples will be critical to thorough verification of planned flight trajectories.

These results also highlight the limitations of the current architecture. Existing reference trajectories have not been designed to be robust to missed thrust across the end-to-end mission (as evidenced by Figure 9). Additionally, existing references have not been designed to facilitate analysis of expected dispersions and perturbation events. This means that models are often computationally brittle, and even small perturbations to the baseline can prevent the optimizer from reconverging into the same solution family. The result is the noise shown in Figure 7, and reducing this noise is a central focus of future work. The behavior identified through initial analysis has highlighted the need to increase the robustness of the trajectories themselves as well as the numerical robustness of the models. Several design and architecture changes have been identified as a result and incorporated into more recent reference trajectory development. Some of these design changes, and the improvements covered in References 2, which were not incorporated in the references analyzed here, also hope to reduce the variability of perturbed cases.

One proposed change to increase downstream Alignment Phase robustness, is to require the termination of the Spiral Phase to occur at the same true anomaly. This change will allow for the initial setup of each individual Alignment Phase to be better tailored to the trajectory itself. Another change to be implement in the design of the trajectory is the addition of a forced coast arc at the end of the Alignment Phase. As seen in Figure 8, recovery solutions to MTEs in the Alignment Phase cannot adjust to MTE lengths longer than 7 days past the seventh or eighth thrusting arc. For MTEs that target the same nominal apolune arrival, there is simply not enough time left in the current Alignment Phase for recovery maneuvers to be performed. The addition of a forced coast arc in the nominal design will allow for recovery trajectories of these late-occurring MTEs to become not only feasible but also more robust to longer duration MTEs.

8.0 Conclusion

As a solar electric propulsion mission, the Gateway Lunar Transit will be susceptible to unplanned loss-of-thrust events, more commonly referred to as missed thrust. As MTEs are a probable, expected eventuality during flight, the ability to analyze these events is a critical part of validating that preflight reference trajectories can be executed in the real-world. Building on a continuously expanding low-thrust mission design framework, a broad methodology to approach the problem of MTEs during Lunar Transit has been developed. This methodology consists of three approaches to missed thrust, applying MTEs according to fixed matrices of analysis cases as well as probability distributions of expected event durations and occurrence epochs. This work enables large-scale MTE analyses to be conducted against complex low-thrust trajectories. By leveraging advanced computing resources, thousands of samples can be analyzed in a matter of days. This forms a critical capability to advance the Gateway Lunar Transit towards flight and provides a basis for future low-thrust missions to approach missed thrust in a comprehensive and reliable fashion.

References

1. McGuire, M., "Overview of the Lunar Transfer Trajectory of the Co-Manifested First Elements of NASA's Gateway," AAS/AIAA Astrodynamics Specialist Conference, Aug. 2021.
2. McCarty, S., "Mission Scan Approach for the Gateway Low Thrust Spiral to a Near Rectilinear Halo Orbit," AAS/AIAA Spaceflight Mechanics Meeting, Jan. 2025.
3. Lee, D. E., "White Paper: Gateway Destination Orbit Model: A Continuous 15 Year NRHO Reference Trajectory," JSC-E-DAA-TN72594, NASA Johnson Space Center, Houston, TX, 2019.
4. Imken, T. Randolph, R., DiNicola, M., and Nicholas, A., "Modeling Spacecraft Safe Mode Events," IEEE Aerospace Conference, 2018. https://doi.org/10.1109/AERO.2018.8396383
5. Hart, William, et al. "A Margin Management Strategy for Low Thrust Trajectories on the Psyche Project," 2020 IEEE Aerospace Conference. IEEE, 2020. https://doi.org/10.1109/AERO47225.2020.9172250
6. "NASA's Psyche Spacecraft Using Backup Fuel Line," NASA Jet Propulsion Laboratory, URL: https://science.nasa.gov/blogs/psyche/2025/05/28/nasas-psyche-spacecraft-using-backup-fuel-line/, 2025.
7. "Glitch on BepiColumbo, Work Ongoing to Restore Spacecraft to Full Thrust," European Space Agency, URL: https://www.esa.int/Enabling_Support/Operations/Glitch_on_BepiColombo_work_ongoing_to_restore_spacecraft_to_full_thrust, 2024.
8. Copernicus Trajectory Design and Optimization System, Software Package, Ver. 5.4.0, NASA Johnson Space Center, Houston, TX, 2025.
9. Gill, P., Wong, E., et al., "User's Guide for SNOPT Version 7.7: Software for Large-Scale Nonlinear Programming," University of California San Diego, Mar 2021.
10. Oh, David, et al. "Analysis of System Margins on Deep Space Missions Utilizing Solar Electric Propulsion," 44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit. 2008. https://doi.org/10.2514/6.2008-5286
11. C. Venigalla, "Multi-Objective Low-Thrust Trajectory Optimization With Robustness to Missed Thrust Events," NASA Goddard Space Flight Center, 2022.
12. Lee, Jinsung, and Ahn, Jaemyung. "Low-Thrust Resonance Gravity-Assist Trajectory Design for Lunar Missions," Journal of Spacecraft and Rockets 61.4 (2024): 1074-1083. https://doi.org/10.2514/1.A35825
13. Karn, S., McCarty, M., and McGuire, M., "Recovery from Missed Thrust During Low Thrust Insertion of NASA's Gateway into a Near Rectilinear Halo Orbit," AAS/AIAA Astrodynamics Specialist Conference, Sep. 2024.
14. Karn, S., McCarty, M., and McGuire, M., "Recovery from Missed Thrust During Gateway's NRHO Insertion Using Auxiliary Propulsion," AAS/AIAA Spaceflight Mechanics Meeting, Jan. 2025.