Automated Optimization of Weak Stability Boundary Lunar Transfers for the LunaH-Map Mission

Jeremy Knittel [1,2], Derek Nelson[1], Andrew Levine[1], Carly VeNard[1], Daniel Wibben[1], Craig Hardgrove[3]

[1]*KinetX, Inc., Space Navigation and Flight Dynamics (SNAFD)*
Simi Valley, CA, USA
Corresponding Author Email: jeremy.m.knittel@gmail.com

[2]*Laboratory for Atmospheric and Space Physics, University of Colorado*
Boulder, CO, USA

[3]*School of Earth and Space Exploration, Arizona State University*
Tempe, AZ, USA

The LunaH-Map (LMAP) mission was manifested and launched as a secondary payload on the Artemis 1 mission in November 2022. With a scientific objective of mapping hydrogen concentration at the Lunar South Pole, the mission required a low polar orbit. As a secondary payload, the LMAP mission design had to overcome many unique constraints to generate Lunar transfers to reach this desired orbit. This paper will discuss the methods used to overcome all these challenges and incorporate two sets of three-body weak-stability dynamics into one low-thrust trajectory optimization problem. Emphasis will be given to the many operational constraints and how automation was employed to meet mission requirements. Instructive trends and results will be presented for the thousands of trajectories generated across the over 400 potential launch opportunities.

I. Introduction

The LunaH-Map (LMAP) spacecraft was manifested and launched as a secondary payload on the Artemis 1 mission in November 2022. With a scientific objective of mapping hydrogen concentration at the Lunar South Pole (see Ref [1]), the mission required a low polar orbit. As a secondary payload, the LMAP mission design had to overcome many unique constraints to generate Lunar transfers to reach this desired orbit. Most notably, Artemis 1 deployed the Orion spacecraft and nine other CubeSats along with LMAP, creating unprecedented constraints on NASA's Deep Space Network (DSN). In turn, this created unprecedented challenges for mission design teams needing to adjust thrusting profiles around communication periods that were provided sometimes only days before attempted launches.

Orion, the primary Artemis payload, desired an injection vector that subsequently put LMAP on a trajectory that would impact the Moon or complete a low leading edge Lunar flyby. Immediately capturing into orbit about the Moon during the first Lunar encounter would have required an extreme propellant load with a chemical propulsion system. Returning to the Moon and inserting into orbit was most feasible with the efficiency of a low-thrust solar electric propulsion (SEP) system. Embarking on such a trajectory required thrusting to begin days before the Lunar encounter, itself only five to six days post-launch.

LMAP's on-board ion engine's predicted performance was mission enabling, however the required burn durations presented many challenges. The LMAP spacecraft (see Fig. 1) could not modify the Lunar flyby sufficiently to avoid being flung to or past the edge of Earth's gravitational sphere of influence before returning to the Moon. This weak stability boundary region has intricate dynamics requiring higher fidelity modeling than typically used for initial trajectory design.

Upon returning to the Moon, long finite burn durations also would prohibit a large orbit insertion. While it requires longer than optimal transfer durations, a separate weak stability boundary, between the Earth and the Moon, allows for a Lunar capture that is nearly ballistic and feasible for a low-thrust spacecraft.

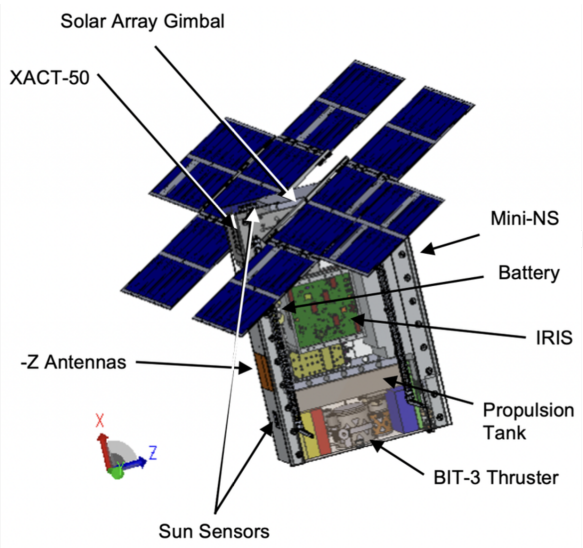

Fig. 1. Diagram of the LMAP spacecraft

Once captured at the Moon, the LMAP spacecraft would have spent months to years spiraling down to a low, elliptical Science orbit. The mission design for that phase of the LMAP mission is discussed in [2]. The navigation strategy for LMAP has been presented previously as well in [3].

This paper will discuss the methods used to incorporate two sets of three-body weak-stability dynamics into one low-thrust trajectory optimization problem. Emphasis will be given to the many operational constraints and how automation was employed to meet mission requirements. Instructive trends and results will be presented for the thousands of trajectories generated across the over 400 potential launch opportunities. While the LMAP spacecraft did not ultimately prove capable of executing the presented mission designs, the tools and capability will hopefully enable future missions.

II. Background

A. Previous Work

The foundational work that informed LMAP's initial design process is presented in [4] and [5]. These early studies helped inform the general structure of viable trajectories that were eventually worked into the presently discussed operational tools. Important lessons include:

1. Early thrusting was required to raise the altitude of the Lunar gravity assist (LGA). Without this thrusting, the energy boost from the LGA would be so strong that LMAP would be ejected from the Earth-Moon system.
2. For practical amounts of thrusting and LGA altitude raise, LMAP would make a near 180 degree turn at the moon and return towards Earth.
3. Less than one week following the LGA, the spacecraft would pass through perigee. Operationally it was required that this occur above 1e5 km to remain above radiation belts. Further perigee raise would require infeasible LGA altitude raise and pre-LGA thrusting.
4. Apogee would occur at roughly 1e6 km.
5. Once at apogee, many solution families exist with varying transfer time to a ballistic capture, sometimes involving additional perigee and apogee passages.
6. The final boundary condition at the Moon can be defined as a near-circular, ballistically captured polar orbit.

These lessons directly translated into constraints and problem formulation in the resulting mission design tools.

B. Software

LMAP mission design work was completed using NASA Goddard Space Flight Center's Evolutionary Mission Trajectory Generator (EMTG) [6] wrapped with a suite of scripts named the Python EMTG Automated Trade Study Application (PEATSA) [7] as well as MIRAGE [8], KinetX Inc.'s licensed version of JPL's ODP/TRAJ [9].

i. EMTG

EMTG employs a non-linear programming (NLP), gradient-based solver to yield locally optimal solutions. The Sparse Nonlinear OPTimizer (SNOPT) [10] is used as the NLP solver, with EMTG supplying the algorithm with decision variables, constraints and an objective function that define the LMAP astrodynamic journey.

To achieve some level of global optimization, upon completion of each SNOPT optimization run, EMTG applies pseudorandom perturbations to decision variables in a method known as monotonic basin hopping (MBH) and then re-runs the gradient based NLP solver. This process is repeated until a time-based stopping criteria is reached. It is therefore a stochastic solver and does not deterministically reach any feasible solution even though one might exist, let alone the globally optimal solution.

Finally, all EMTG tasks are facilitated through use of PEATSA. This script suite provides a simple to use command line and text file interface to accomplish a wide variety of trajectory optimization tasks. This involves creating EMTG input scripts for all scenarios, running those cases in parallel through the EMTG software, post-processing EMTG results, and finally using those results to share initial guesses from case to case setting up subsequent iterations. All of these tasks are automated and require only manual generation of input files.

Through the combination of MBH and PEATSA's initial guess sharing, EMTG is extremely effective and efficient at reaching globally optimal solutions.

ii. MIRAGE

The MIRAGE software package has been used to successfully navigate many deep-space missions including New Horizons, OSIRIS-REx, EMM, and most recently, Lucy. It is flight fidelity operational orbit determination (OD) software, with some limited mission design capabilities.

In order to extend the mission design capabilities of MIRAGE, particularly in the way of low-thrust trajectories, the PVDrive Interface and Robust

Astrodynamic Targeting Engine (PIRATE) was written [11]. This allows modern scripting and complex mission design optimization using SNOPT to be connected to the fully validated and verified propagation and force modeling capabilities within MIRAGE.

III. Low-Fidelity Initial Design

LunaH-Map mission design was performed in stages of increasing modeling accuracy, allowing rapid intuition-free mapping of the design space. This section and the following will walk through each stage of design including discussion of the software used, trajectory modeling and characteristic results.

High fidelity solvers require longer run-times and often are unable to converge without approximate initial guesses for long duration finite burn thrusting. The purpose of the first stage of LMAP mission design was to create trajectory solutions that could seed higher-fidelity solvers starting with nothing but trajectory injection states and epochs for each launch opportunity.

A. Modeling

Initial LMAP mission design employed EMTG's two-point shooting algorithm (based on the Sims-Flanagan method [12]) named MGALT (Multiple-Gravity Assist with Low-Thrust). Each thrust arc is modelled as a series of n impulsive maneuvers with magnitude equal to the net delta-v capable of the finite burn system in equivalent time. Between each impulsive maneuver, trajectory propagation is performed analytically using Keplerian, 2-body modeling. As n increases, the trajectory fidelity increases, but so does computation time.

Given the approximations with the thrusting model, explicit navigation coasts are not modeled. Instead, when calculating delta-v available over a given timespan, thrust and mass flow rate are scaled down by an assumed duty cycle to mimic thruster off time. By scaling both thrust and mass flow, Isp is held constant. The principle behind this strategy is that for a duty cycle of x, thrusting at x% for 100% of the time is roughly equivalent to thrusting at 100% for x% of the time.

Using a lower than 100% duty cycle also provides the trajectory with margin to recover from any incidents which prevent the propulsion system from operating when it would otherwise be preferable. These are called missed thrust events and less than 100% duty cycling provides missed thrust margin. The majority of the LMAP trajectory is relatively insensitive to missed thrust events, as there is significant deterministic coasting in the transfer, however the thrusting prior to the LGA is extremely sensitive, and thus a lower duty cycle of 90% was used in this portion of the trajectory beyond what would be needed for navigation coasts (~95%). An even lower duty cycle would have been preferred; however, this would significantly limit the feasibility of the trajectory.

Generating weak stability boundary trajectories, particularly those involving low-thrust maneuvering is relatively challenging to accomplish in low-fidelity, which typically involves simplifying assumptions like 2-body gravitational acceleration. As mentioned, between each impulsive maneuver, trajectory propagation here is modelled as 2-body for rapid propagation. However, to capture perturbation effects, an impulse is also applied with magnitude proportional to the averaged effect of the perturbation. This method is discussed in detail in [13]. Once again, as the number of timesteps, n, increases, computation time increases, but so does accuracy. Any time varying perturbation accelerations can be modelled in this way, but most commonly this is used to model solar radiation pressure (SRP) and 3rd body gravitation.

In this work, Earth is primarily used as the integration center, with the Sun and Moon modelled as perturbing 3rd bodies. During the Lunar Gravity Assist and the final ballistic capture at the Moon, the integration center is swapped to the Moon, with a handover at the Moon's sphere of influence (SOI). SRP effects are also included throughout.

B. Deployment

The only required input for an LMAP trajectory study was injection vectors and epochs from the Artemis launch provider team. They provided a 6-degree of freedom (DOF) state at the deployment epoch, to which the LMAP team had to then apply the delta-v imparted from the deployment mechanism.

See Fig. 2 for a schematic of the deployment geometry. The Artemis vehicle was spin-stabilized about a direction 55 degrees from the Sun direction towards the velocity vector. The rotational state at the exact moment of deployment would be unknown. The ten secondary CubeSat deployers were oriented at a 34 degree angle relative to this spin axis.

For mission design purposes, only the deterministic component (parallel to spin axis) of that deployment velocity was applied as delta-v to the initial state. Navigation would need to correct for the deployment velocity in the direction perpendicular to the spin axis, along with all other unknowns.

C. Optimization

Initially, the design process was performed for all provided launch injection states across launch windows

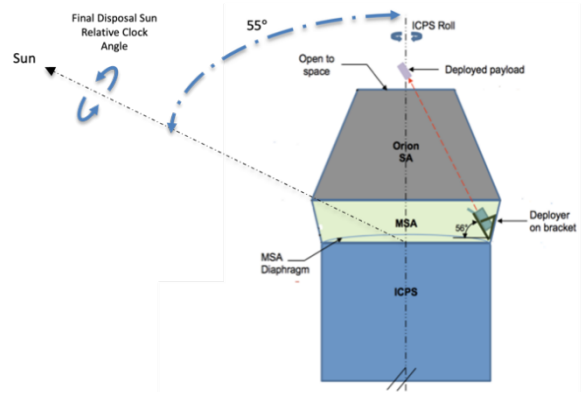

Fig. 2. Artemis CubeSat Deployment Geometry. Original (unannotated) image courtesy of Artemis program, NASA Johnson Space Center.

and launch periods. After building the medium-high and high fidelity tools discussed in later sections, it became clear that optimizing across the launch windows did not provide sufficient benefit to justify this effort. A trajectory solution available for the launch period open would easily converge for the launch period close. I.E., the one minute to two-hour delay in launch conditions did not provide sufficient alterations to require a separate trajectory study.

A total of 114 separate optimizations were run for each launch day. This was a balance between computational resources available and the computational time required. Each optimization represented a different combination of constraints, shown in Table 1. Note that these are all maximum constraints, and for each value the lower bound was either the next lower constraint, or 0.

Table 1. Optimization Grid

Parameter	Min	Max	Step
Time of Flight [days]	50	230	10
Capture Eccentricity	0.1	0.20	0.05
Capture Semi-major Axis [km]	32500	35000	2500

The primary constraint of relevance defining the transfer is the time-of-flight. The other constraints were included as a proxy for the ballistic or near-ballistic weak capture at the Moon. As presented in [2], most combinations of these orbital parameters resulted in an orbit that was sufficiently circular and captured, that a low-thrust spiral could begin before the spacecraft was ejected from the Lunar SOI. Further, including a grid of capture scenarios provided initial guess seeding to work across consistent time-of-flight constraints as transfers would be very similar, and increased the probability that a solution would be found if one existed.

D. Results

The results presented here represent Artemis 1 Launch Periods (LPs) 22 through 29 covering dates from June 6th 2022 through December 23, 2022 excluding LP26 whose data was lost due to a file corruption issue. (Note: Artemis actually launched on November 16th, 2022, the 4th day of LP28). This comprises over 3800 converged trajectories.

Fig. 3 shows the required propellant to reach a lunar capture for all converged cases in the study. For reference, LMAP had 1.5 kg of Iodine propellant onboard, and a spiral down to the desired orbit would likely require at least half of that and likely closer to 1 kg. So only solutions requiring less than .5 kg for the Lunar transfer were considered. The histogram shows that launch periods 27-29 were found to be generally the most favorable for LMAP, as they largely required less propellant. This is due to the Artemis injection vectors providing a more favorable ballistic LGA encounter. With a higher nominal LGA encounter, LMAP would have been more likely to be able to maneuver such that the gravity assist put the spacecraft on a near ballistic Lunar return trajectory.

Fig. 4 shows that across all launch periods, more solutions were available with longer times-of-flight to return to the Moon. The explanation for this can be explained by considering the time-varying availability of certain trajectory types. Low-thrust Lunar transfers can be categorized by the geometry in the rotating Sun-Earth frame as well as the number of revolutions around Earth completed. The relevant metrics for geometric evaluation are the location of apsides, the location of the Moon and the presence or lack of a "figure-8" pattern. The authors have identified 16 different trajectory families resulting in different combinations of the above metrics. Launch directions and the matching Moon locations in this frame will be discussed as clock times or equivalently clock angles.

The color mapping for the remaining results in this section is presented in Fig. 5. For scaling and reference,

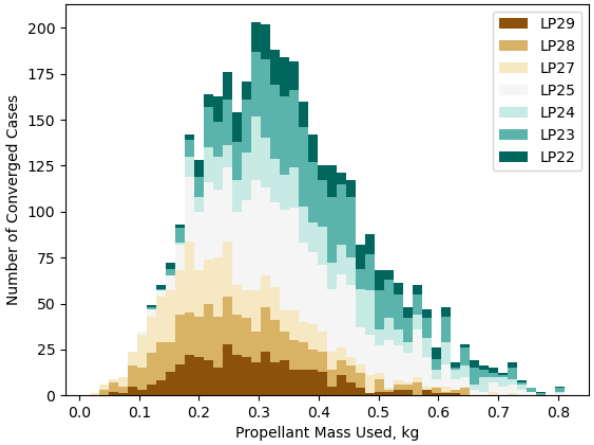

Fig. 3. Propellant Usage for the Launch Periods

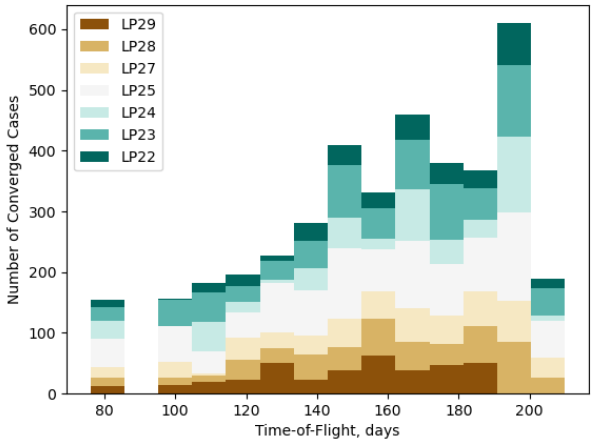

Fig 4. Time-of-Flight to Lunar Capture for the Launch Periods

all plots include locations of Sun-Earth L1 and L2 (at ~±1.5e6 km). The Sun is in the -x direction, so L1 is also in the -x direction, while L2 is in the +x direction. The authors will use the nomenclature presented in [5] such that the geometry is broken up into four quadrants. Quadrant-I, II, III, and IV refer respectively to clock angles of 12 to 3, 9 to 12, 6 to 9 and 3 to 6.

Figs. 6 and 7 show similar orbit families that include a "figure-8" and either a completed loop around Sun-Earth L1/L2 or one that comes close to doing so. The Lagrange point involved lies on the opposite side of Earth from the launch direction. I.E., a launch on the Sunward side comprising clock angles of ~8 to 12 resulted in trajectories towards or around L2. Conversely, launch clock angles of ~3 to 7 resulted in trajectories towards or around L1. This is due to the Artemis launch providing a leading-edge Lunar gravity assist, resulting in a near 180-degree turn of the spacecraft back towards Earth. LMAP did not have sufficient thrusting capability to convert the Artemis injection into a trailing edge flyby. Halo orbits tend to be relatively long duration. As time-of-flight shortens, near-Halo orbits are more common than true Halo orbits.

Fig. 8 depicts the shortest transfer family found in this study. These orbits are only possible with launch into Quad-I or Quad-III, with apoapsis reached in a near perpendicular direction to the Sun-Earth line. The spacecraft simply transfers up to apoapsis, remains there sufficiently long such that in the rotating frame, a "figure-8" is completed, and then falls back into capture at the Moon.

Lunar transfers that represent transition from so-called perpendicular "figure-8" orbits to near-Halo orbits are pictured in Fig. 9. As transfer duration increases, apoapse moves from the perpendicular direction through Quad-II or Quad-IV towards L1 or L2. Not surprisingly, launch directions have some overlap but also represent the region between those depicted In Figs. 7 and 8.

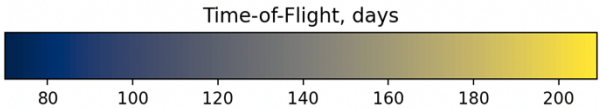

Fig. 5. Color scale used for subsequent trajectory plots

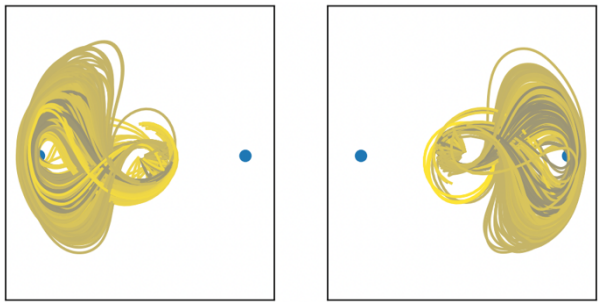

Fig. 6. L1 (left) and L2 (right) Halo-orbit transfers

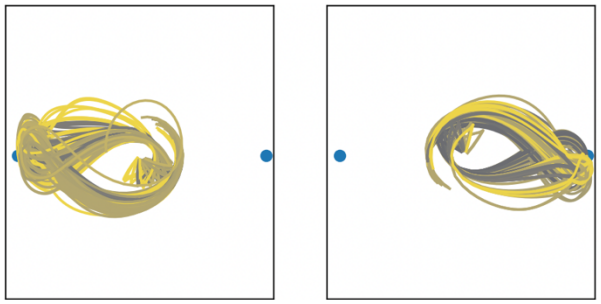

Fig. 7. Near-Halo "Figure-8" orbits

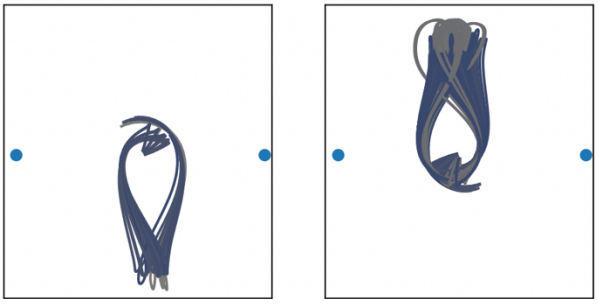

Fig. 8. Perpendicular "Figure-8" orbits

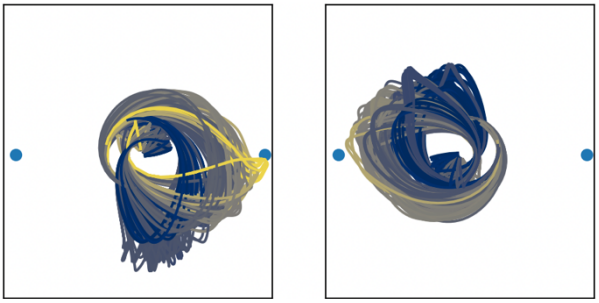

Fig. 9. Single Revolution Quad-IV (left) and Quad-II (right) orbits

Two additional single revolution orbit families found are shown in Fig. 10 and 11. These orbits are unique from those shown previously in that they reach apoapse twice during their single revolution. This also distinguishes those in Fig. 10 from those in Fig. 11. In Fig.

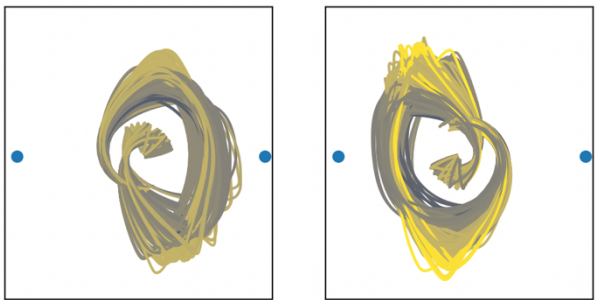

Fig. 10. Single Revolution, Quad-IV/Quad I (left) and Quad-II/Quad-III (right) orbits

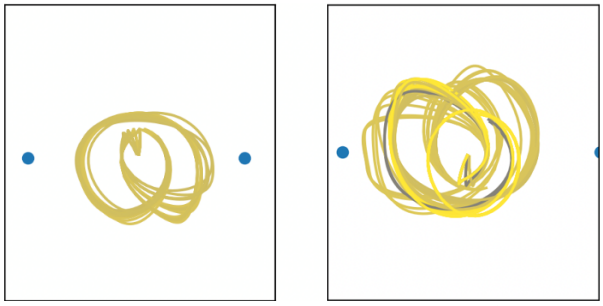

Fig. 11. Single Revolution, Quad-III/Quad-IV (left) and Quad-II/Quad-1 (right) orbits

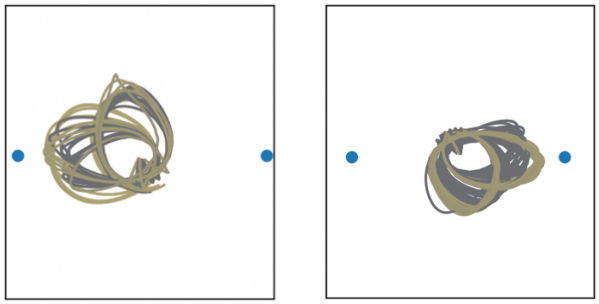

Fig. 12. Dual Revolution orbits. L1/Perpendicular (left) and L2/Perpendicular (right)

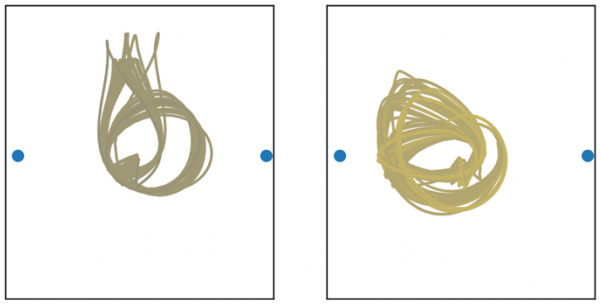

Fig. 13. Double revolution transfers: Perpendicular/Quad-I (left) and L1/Quad-II (right) orbits

10, apogee is reached once above the Sun-Earth line, and once below, whereas in Fig. 11, both apoapses are found on the same side of the Sun-Earth line. The direction of launch determines the quadrant in which these apogees occur and the order. One as-yet unexplained result is that orbits launching below the Sun-Earth line tended to have a longer time-of-flight to complete this orbit type than those launching above the Sun-Earth line. This is found in the families shown in both Figs. 9 and 10.

Figs. 12 and 13 depict the orbit families found with two-revolutions before capture at the Moon. Significantly fewer double-revolution trajectory solutions were found. Due to the stochastic nature of the solution method, the number of revolutions was not prescribed and this could be a result that these are more dynamically challenging for a two-point shooting transcription to converge. However, it is believed that the primary cause is that the range specified in time-of-flight is simply less permitting of two revolution transfers.

Fig. 12 trajectories reach an apoapsis directly towards either L1 or L2, then fall back into a periapse inside of the Moon's orbit, prior to their second apoapse in a perpendicular direction to the Sun-Earth line. Launch direction again specifies the direction of the apsides in this symmetric orbit family.

Fig. 13 depicts two somewhat similar, but unique, orbit families, for whom a symmetric pair was not found. It is the authors' hypothesis that the symmetric pairs exist, but simply weren't found due to unequal distribution of the launch directions in the particular launch periods of study, however this would need further analysis to confirm.

IV. Medium-High Fidelity Detailed Analysis

With plentiful low-fidelity solutions available, work transitioned into higher accuracy modeling, and a more reliable representation of the design space available for each launch opportunity.

A. Improved Modeling

The first major modeling improvement needed was in the propagation model. This stage employed EMTG's Finite-Burn Low-Thrust (FBLT) transcription. As the name implies, rather than using discrete impulsive maneuvers to approximate long-duration SEP thrusting, continuous accelerations were integrated using a fixed-step Runge-Kutta-8 method. The time-step used in each mission phase was reduced until the match point constraints remained constant to reasonably tight (~1e-7 km) tolerances.

With the implementation of finite burn modeling, there is benefit to increased fidelity in the handling of navigation coasts by modeling them explicitly with durations and timing matching a conservative estimate of a tracking schedule (not the tracking schedule that the DSN would actually assign though). In this case, conservatism meant assuming more tracking early when thrusting is more efficient in moving the LGA flyby

parameters. During these assumed coasting periods, the optimizer was constrained to a 0% duty cycle. Outside of these periods, a 90% duty cycle was used for the purposes of missed thrust recoverability.

Because the trajectory was numerically integrated, perturbation approximations were no longer needed, and were also continuously applied. Those considered were again 3rd body effects (Sun, and Moon or Earth depending on the integration center) and SRP, while Moon and Earth J2 gravitational effects were added at this step.

B. Optimization

Every converged low-fidelity trajectory was moved to medium-high fidelity. The optimization grid was therefore the same as presented in Table 1, however the grid was now sparse due to cases that had not converged in low-fidelity being excluded. The level of sparsity of course varied from launch period to launch period (see Table 2). As discussed, it was not worthwhile to attempt to converge cases at this fidelity without an initial guess.

C. Results

This section will summarize the general results found when moving to higher-fidelity modeling as well as provide a more detailed analysis of launch period 28, in which the Artemis vehicle departed Earth in November 2022.

While further modeling fidelity was possible and will be discussed in the following section, if a solution was found with the present modeling, the accuracy is high enough that it is reasonable to conclude that it is physically valid. The same is not true for the low fidelity modeling, and some cases found there "cheated" the optimizer by finding numerically valid, but physically invalid ways around certain constraints, such as passing below the surface of the Earth between control points. For that reason, not all low fidelity cases converged when moved to the medium-high fidelity modeling. However, Table 2 shows that over 70% of cases were found to be usable. Note that LP23 and LP24 were not moved to medium-high fidelity due to launch slips occurring before the analysis was started, while LP29 was not revisited as the launch had already occurred in LP28.

Fig. 14 shows how well the low-fidelity modeling captured the actual required propellant usage. Thankfully, in most all cases, the higher fidelity modeling was able to find a solution with lower propellant usage (above the solid line, which represents a perfect match). Longer trajectories were slightly more likely to demonstrate a mass usage mismatch. The same number of control points was used regardless of the time-of-flight, therefore, longer trajectories travelled further between perturbations and impulsive maneuvers being applied, reducing accuracy.

Table 2. Success Rate of Transferring Trajectories from Low Fidelity to Medium/High Fidelity

LP	Converged Low Fidelity Cases	Converged Medium/High Fidelity Cases	Percent Convergence
28	487	389	79.9%
27	492	411	83.5%
25	1015	897	88.4%
22	336	246	73.2%

Additional focused results pertaining to LP28 are shown in Figs. 15 and 16. Fig. 15 displays the trajectory results as a function of time-of-flight for each launch day. For brevity, only the first four launch days are shown. Additional thrusting is consistently shown to be required to reach a more circular orbit upon capture. In many cases, a more circular orbit was not even possible (ecc <= .1).

Given the nature of the trajectory families presented in the previous section, it should also not be surprising that there is not a monotonic trend as time-of-flight increases. Certain transfer types are simply more efficient for given launch conditions than others.

The amalgamated results for LP28 are shown in Fig. 16. This plot shows the best delivered mass case, regardless of time-of-flight for each launch day. As expected, due to the changing availability of certain transfer types, and the efficiency of that transfer type for this specific timing, some launch dates are more favorable than others. Further, each transfer type requires a differing level of pre-LGA thrusting, which may or may not be feasible given the specific Artemis ballistic LGA target

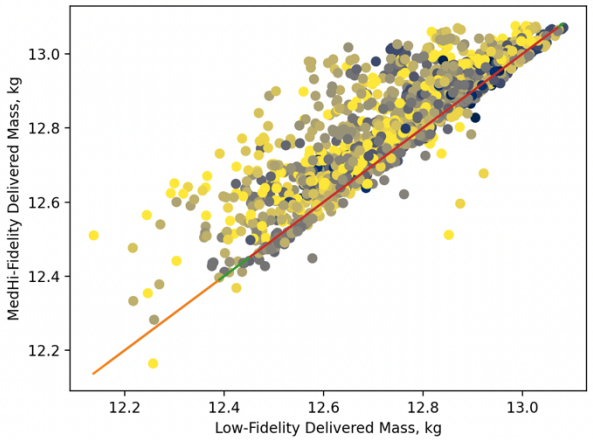

Fig. 14. Comparison between Final Trajectory Mass with Low and Medium/High Fidelity Modelling. See Fig. 5 for color scale

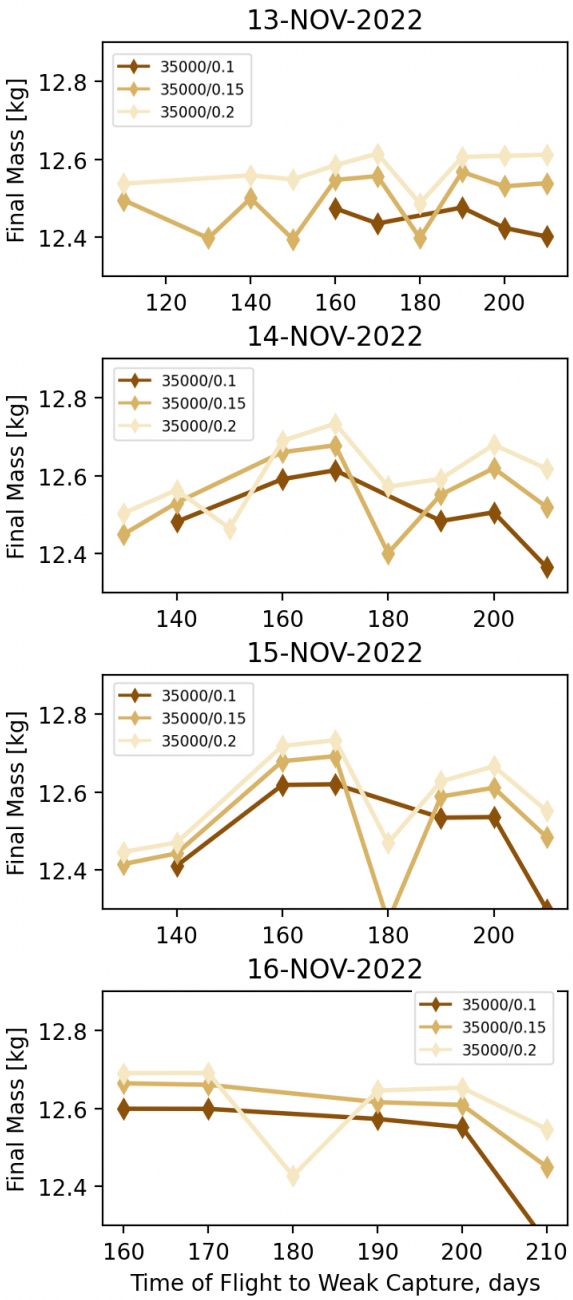

Fig. 15. Least propellant required trajectories found for each launch day sorted by Lunar orbit capture conditions [semi-major axis (km)/eccentricity] and time-of-flight (LMAP launched at 12.865 kg)

for that launch. If the optimal LGA parameters could not be achieved with pre-LGA thrusting, additional thrusting was required after the encounter to compensate.

V. High Fidelity Flight Products

The final step of the mission design process was to convert the trajectories into maneuver designs that could be built into spacecraft command sequences uploaded to the vehicle.

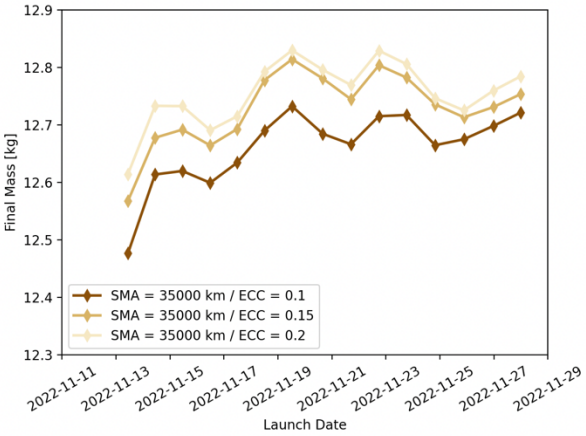

Fig. 16. Least propellant required trajectories found for each launch day removing effect of time-of-flight (LMAP launched at 12.865 kg)

A. *Improved Modeling*

While EMTG has made major improvements in recent years and has been validated against MIRAGE [14], for the creation of actual flight products, a trajectory generated with full OD level modeling including an adaptive step integrator, full *m* x *m* harmonic fields, solid tides, and a fully continuous forward integration is preferred.

Having the final high-fidelity products generated in MIRAGE also provided operational simplicity; as the OD team generated solutions, any estimated parameters in the modeling were immediately and directly fed into future maneuver designs.

B. *Duty Cycle Considerations*

As discussed above, a less than 100% duty cycle is assumed on the SEP thruster to model both navigation coasts and provide operational missed thrust margin. However, when designing flight products, it is beneficial to allow an optimizer to use the full available thrusting timespan for flight products, to increase missed thrust margin and to have the ability to overcome maneuver execution and OD error in later thrusting periods. I.E., if designing flight products for Jan 1-8, the allowable duty cycle usable in the optimization process for Jan 1-8 should be increased up to 100% with thrust modeling after Jan 8 kept at a lower (typically 90%) duty cycle. This almost always results in less deterministic thrusting after Jan 8th, thereby improving missed thrust margin. If the spacecraft suffers a missed thrust event in the Jan 1-8 timespan, then the maneuvering-to-go would need to be redesigned anyway, and nothing is lost in terms of missed thrust recoverability.

When releasing this duty cycle margin for LMAP, it was done in an iterative process to always prefer thrusting sooner rather than later. The first maneuver in a given

sequence was allowed to go up to 100% duty cycle, while all others were constrained to a 90% duty cycle. Once the trajectory design process was complete, it was repeated, with the first maneuver held constant with the design provided above, and the second maneuver allowed to go up to 100%. This was repeated for all maneuvers requiring flight products in a given design cycle. As this process occurred, the expected result was that downstream maneuvers became shorter and shorter, indicating less required downstream thrusting if no missed thrust events occurred.

While this likely would not always yield truly mass optimal solutions, the duty cycle was never forced to 100%, so the optimizer could always lower a given thrusting period's duty cycle if it was truly suboptimal to thrust at 100%. Therefore, any mass inefficiency was certainly very small, and deemed worth the benefits in improved downstream missed thrust recoverability and ability to correct for navigation errors.

For this margin strategy release process to work (i.e. using 100% duty cycle), then the timing of navigation coasts must be known. This introduced a challenge, as the DSN was unable to negotiate and deliver tracking schedules for all available launch opportunities due to antenna congestion. As a result, this margin release strategy was only implemented once the DSN provided tracking schedules, and much work was put into automating the process outlined above to dramatically reduce the required execution time needed to meet mission and operational scheduling constraints. Tracking schedules were sometimes only provided 3 days before a given launch opportunity, leaving very little time to generate and verify initial flight products.

C. Optimization

The goal of optimization was to yield a trajectory in MIRAGE matching flight-fidelity modeling. Although PIRATE provides MIRAGE with significant low-thrust optimization capability, it is forward-shooting modeling only and not able to globally re-optimize a full trajectory. Instead, chunks of 2 to 3 weeks worth of maneuvering were re-targeted to the fixed states from the EMTG globally optimized trajectory. This process is known as piecewise targeting.

There were four steps involved in the highest-fidelity conversion process:

1. Adjust the thrusting to fit around the negotiated DSN schedule
2. Re-optimize that trajectory in EMTG
3. Piecewise re-target that trajectory in PIRATE to generate a reference trajectory
4. Apply the duty cycle release methodology discussed above.

Scripts were generated so that each of the first 3 steps was executed in parallel for all launch window cases.

The 4th step was simple and quick enough that it was only done manually when flight products for a given launch opportunity were required. The flight products for the launch period open were generated before the launch day. If a launch slip happened within the window, then the LMAP mission design team would repeat step 4 for the new opportunity. There was enough time to generate and verify flight products post-Launch and prior to LMAP coming into first communication acquisition and being available for the sequence upload.

D. Results

The thrusting profile for the LMAP mission design for Nov 16th 2022 for the three levels of fidelity is shown in Fig. 17. Due to the low-acceleration produced by the thruster, the timing of maneuvers can move noticeably while still flying a qualitatively very similar trajectory. Even though LMAP would not have utilized all thrusting opportunities (the somewhat transparent regions in Fig. 17), it was still important to isolate thrusting to periods when the spacecraft would not be in communication with the DSN. These unused thrusting periods were still available operationally for maneuvers to overcome navigation uncertainty.

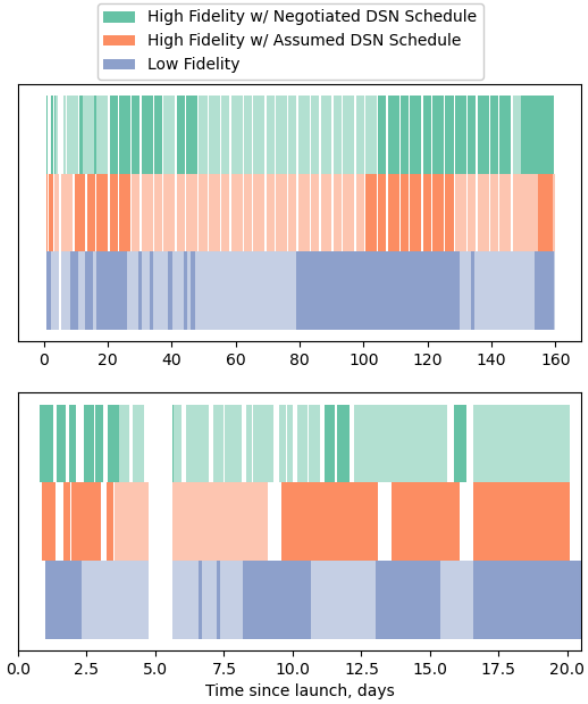

Fig. 17. Thrusting Profiles for the full trajectory (top) for the realized Nov 16 2022 launch, and a zoom in of the first 21 days with the realized negotiated DSN schedule (bottom). Opaque regions show when thrusting was required and transparent regions show thrusting opportunities not used by the optimizer.

The next two diagrams, Figs. 18 and 19, reveal the optimization performance of moving solutions from EMTG to MIRAGE. Fig. 18 shows that because an initial guess trajectory and targets optimized for the launch open were used, propellant requirements increased as the launch window progressed. In future work, the LMAP mission design team concluded that if only one trajectory was used for global target optimization, then a better strategy would have been to use the middle of each day's launch window. However, the overall propellant increases were very small.

Fig. 19 shows that the modeling in the medium/high fidelity case was quite accurate at predicting overall performance. And favorably again, the highest fidelity modeling presented opportunities to better take advantage of perturbations and reduce overall propellant in flying a weak stability boundary transfer. This shows that PIRATE does provide some optimization capability beyond simple retargeting convergence.

Finally, to visualize the movement required in the B-plane from pre-LGA thrusting, Fig. 20 shows the targets before and after the designed maneuvers. Certain

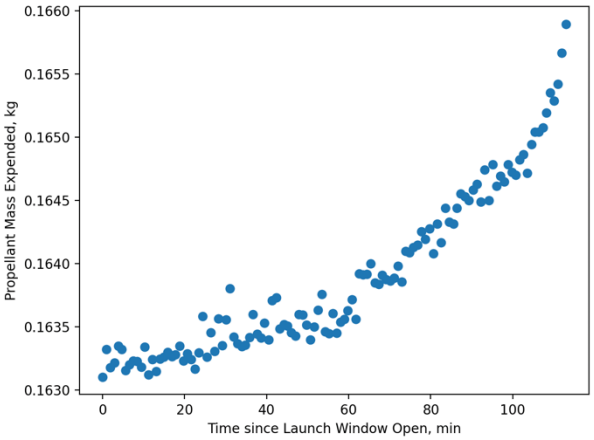

Fig. 18. Propellant requirements across the launch window

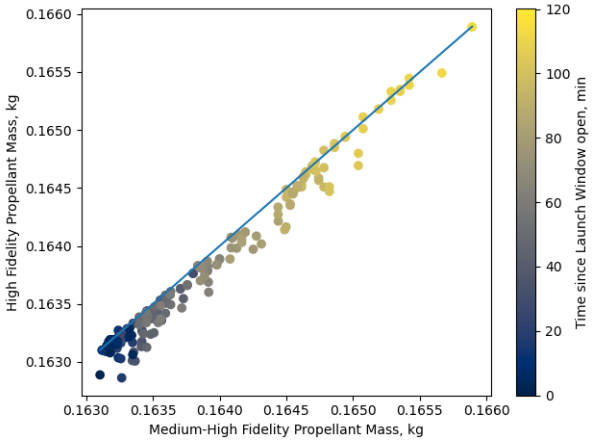

Fig. 19. Comparison between Medium/High Fidelity trajectories and High Fidelity

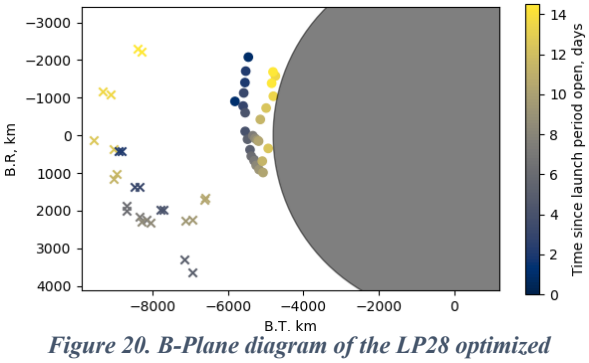

Figure 20. B-Plane diagram of the LP28 optimized trajectories (x's) and ballistic flybys (dots)

trajectory families required a different target in the B-plane, so as the launch period progressed the targets varied widely, both in location on the B-plane and the perilune raise required. Of course, in all cases, as discussed, the perilune was raised significantly.

VI. Conclusions

This work has laid out a detailed framework for performing intuition-free and initial-guess-free mission design of weak-stability boundary Lunar transfers. In much previous work, the overall trajectory structure needed to have been known and/or complicated circular restricted 3-body algorithms were required to find initial guess low-thrust transfers. This work used a direct solver and n-body inertial modeling throughout, and understanding of trajectory families was an output of the mission design and optimizer rather than an input.

Much work went into automating every step in the process so that analyst time was minimized and that the constraints of the process (i.e. receiving DSN schedules with a very short turnaround time before launch, and analyzing multiple launch periods due to launch slips) were still viable for operations. Unfortunately, an onboard propulsion failure ultimately caused a loss of the LMAP mission; however, the tools and processes developed will no doubt enable future low-thrust missions.

VII. References

[1] Hardgrove, C. *et al*. "The Lunar Polar Hydrogen Mapper CubeSat Mission," *IEEE Aerospace and Electronic Systems Magazine,* Vol. 35, 2020.

[2] Wibben, D. *et al*. "Low-thrust, many-revolution, orbit transfer design for the LunaH-Map Lunar Cubesat mission," *Astrodynamics Specialist Conference, Big Sky, MT, USA,* August 2023.

[3] Nelson, D. *et al*. "Navigation Strategy, Preparations, and Results for the LunaH-Map Lunar CubeSat Mission," *Astrodynamics*

Specialist Conference, Big Sky, MT, USA, August 2023.

[4] Genova, A. and D. Dunham. "Trajectory Design for the Lunar Polar Hydrogen Mapper Mission," *Space Flight Mechanics Meeting, San Antonio, TX, USA,* Feb 2017.

[5] Genova, A. *Existence of Ballistic Capture Transfer Trajectories from Trans-Lunar Injection to Lunar Polar Orbit for all Lunar Phases at Time of Departure.* Master's Thesis, Florida Institute of Technology, December 2020.

[6] Englander, J. and B. Conway. "An Automated Solution of the Low-Thrust Interplanetary Trajectory Problem," *Journal of Guidance, Control, and Dynamics.* Vol 40, No 1, pp 15-27, 2017.

[7] Knittel, J. *et al*. "Automated Sensitive Analysis of Interplanetary Trajectories for Optimal Mission Design," *Internationa Symposium on Space Flight Dynamics, Matsuyama, Japan,* June 2017.

[8] Guinn, J. *et al*. "TOPEX/POSEIDON Operational Orbit Determination Results Using Global Positioning Satellites," *Astrodynamics Specialist Conference, Victoria, British Columbia, Canada,* August 1993.

[9] Moyer, T. "Mathematical Formulation of the Double-Precision Orbit Determination Program (DPODP)," *TR 32-1527, Jet Propulsion Laboratory, Pasadena, CA, USA,* 1971.

[10] Gill, P. *et al.* "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," *SIAM Journal of Optimization.* Vol 12, pp 979-1006, Jan 2002.

[11] Knittel, J. *et al*. "Automated Navigation Analysis for the Lucy Mission," *Astrodynamics Specialist Conference, Portland, ME, USA,* August 2019.

[12] Sims, J. and S. Flanagan. "Preliminary Design of Low-Thrust Interplanetary Missions," *Astrodynamics Specialist Conference, Girdwood, Alaska,* August 1999.

[13] Englander, J. *et al*. "Optimization of the Lunar IceCube Trajectory using Stochastic Global Search and Multi-Point Shooting," *Astrodynamics Specialist Conference, Virtual,* August 2020.

[14] Englander, J, *et al.* "Validation of a Low-Thrust Mission Design Tool Using Operational Navigation Software," *Space Flight Mechanics Meeting, San Antonio, TX, USA,* Feb 2017.