AAS 23-363

TRAJECTORY OPERATIONS OF THE ARTEMIS I MISSION

Randy A. Eckman[*], **Charles Barrett**[†], **Amelia L. Batcha**[‡], **and Brian J. Killeen**[§]

This paper describes the operational design and execution of the Artemis I trajectory. It was an operationally complex trajectory with powered lunar flybys and insertion into a Distant Retrograde Orbit (DRO). A joint team of trajectory analysts at the NASA Johnson Space Center (JSC) was responsible for the design and operation of nominal and off-nominal in-space trajectories. A process was developed to convert optimized reference trajectories into Orion burn plans that could be uplinked to the vehicle. During the mission, the joint flight controller and engineering team continuously evaluated upcoming translational burns using actual vehicle conditions, monitored the trajectory for opportunities to re-optimize the trajectory in order to reduce propellant usage, and prepared for potential off-nominal scenarios. Overall Orion in-space trajectory performance is compared to mission designs to demonstrate the success of the design and operations workflows.

INTRODUCTION

The Artemis I mission was NASA's first integrated test flight of the Space Launch System (SLS) rocket, the Orion Multi-Purpose Crew Vehicle (MPCV), and the upgraded Exploration Ground Systems (EGS) at the Kennedy Space Center (KSC). The mission's highest-priority objective was to test the Orion heat shield during re-entry at lunar return speeds. In addition to this objective, the mission sought to adequately demonstrate Orion systems' readiness to support human crews. A Distant Retrograde Orbit (DRO), chosen as the Artemis I destination for historical reasons,[1] provided an opportunity to complete these stated mission objectives, but also presented the most challenging trajectory design yet attempted by a human-rated spacecraft. A Perigee Raise Maneuver (PRM) and a Trans-Lunar Injection (TLI) burn were performed by the Interim Cryogenic Propulsion Stage (ICPS) before five planned major Orion burns. A main engine checkout at Outbound Trajectory Correction-1 (OTC-1) preceded the lunar Outbound Powered Flyby (OPF), sending Orion towards its target DRO. A DRO Insertion (DRI) burn initiated a minimum six-day stay in the orbit prior to a DRO Departure (DRD) burn that began the return to Earth. A final lunar Return Powered Flyby (RPF) sent Orion towards a targeted splashdown in daylight off the coast of San Diego. Much of the conceptual design for this mission profile is discussed in Ref. 2. An example of the Artemis I mission design is illustrated in Figure 1.

During Orion on-orbit operations, the Orion Flight Software (FSW) maintains onboard knowledge of the reference trajectory in a table of thousands of parameters that comprise the set of burn targeting and execution data, referred to as the *burn plan*.[3,4] A completed burn plan generated by the ground teams is then divided into a series of commands and uplinked to the vehicle, which

[*]Orion Mission Design Engineer, NASA Johnson Space Center, Houston, TX 77058, randy.eckman@nasa.gov
[†]Exploration Flight Dynamics Officer, NASA Johnson Space Center, Houston, TX 77058, charles.p.barrett@nasa.gov
[‡]Orion Mission Design Engineer, NASA Johnson Space Center, Houston, TX 77058, amelia.l.batcha@nasa.gov
[§]Orion Mission Design Engineer, NASA Johnson Space Center, Houston, TX 77058, brian.j.killeen@nasa.gov

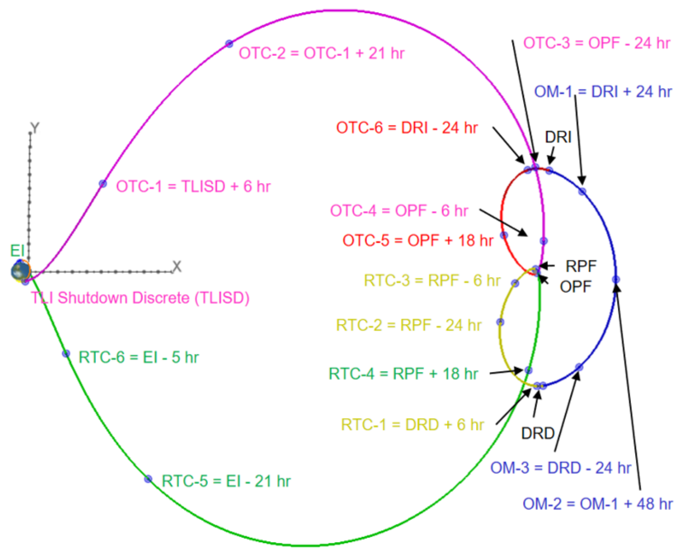

Figure 1: Example Artemis I translational burn timing for an arbitrary launch date, plotted to scale in an EMRP frame. The major burns, in black text, allowed their times to be optimized. For each phase, the trajectory and the events occuring along it are indicated by the following colors: magenta, outbound Earth departure phase; red, post-OPF phase; blue, DRO phase; yellow, post-DRD phase; green, Earth return phase.

gives Orion the necessary information to target and execute each remaining burn in the plan without otherwise requiring ground intervention. This level of automation, to provide a spacecraft the capability to target and complete a full mission in a total loss of communications, is novel in human spaceflight. While individual burns could be targeted onboard or future burns could be scheduled on previous spacecraft, Orion's ability to solve for and maintain a sophisticated cis-lunar trajectory entirely on its own is a tremendous step forward to ensuring safe crew return on increasingly complex Artemis missions. This unprecedented capability is accomplished through the usage of Two-Level Targeters (TLTs)[5] in trajectory operations.

Figure 2 illustrates the entire workflow of mission design and launch preparations for the launch on November 16, 2022. The design reference mission dates back more than a decade,[1] but this timeline and subsequent sections focus on the processes resulting in the provenance of the specific trajectory flown for this successful and historic launch.

The final Δv magnitude of each major burn executed by Orion was within 2 ft/s of the corresponding optimized burn in final 3-Degrees of Freedom (DOF) reference trajectories produced by the On-Orbit Performance Team (OPT), in some cases well within that difference as early as the preliminary mission design.

This paper will introduce the process of transitioning Artemis I from conceptual trajectory design into applied operational spaceflight. The collaboration between many teams across the Artemis program to accomplish the mission design is explained in *Trajectory Responsibilities*. The reference trajectory design cycle is discussed in the *Design Phase* section. Significant attention is given to the concepts of developing operational targeting plans to uplink to Orion in the *Operational Burn Plan and Ephemeris Generation Process* section. The *Pre-Launch Processing* section covers the development of operational products in the Mission Control Center (MCC) following the hand-off of trajectories from design teams to operations as well as covering activities in the weeks leading up to a launch attempt. The *Real-Time Operations* section addresses the real-time trajectory operations support from launch day through the end of the mission. Finally, the *Mission Performance* section provides insight into how trajectory operations progressed for the actual mission as well as further details evaluating the Orion Δv performance. A *Terminology* section is also provided at the end for reference.

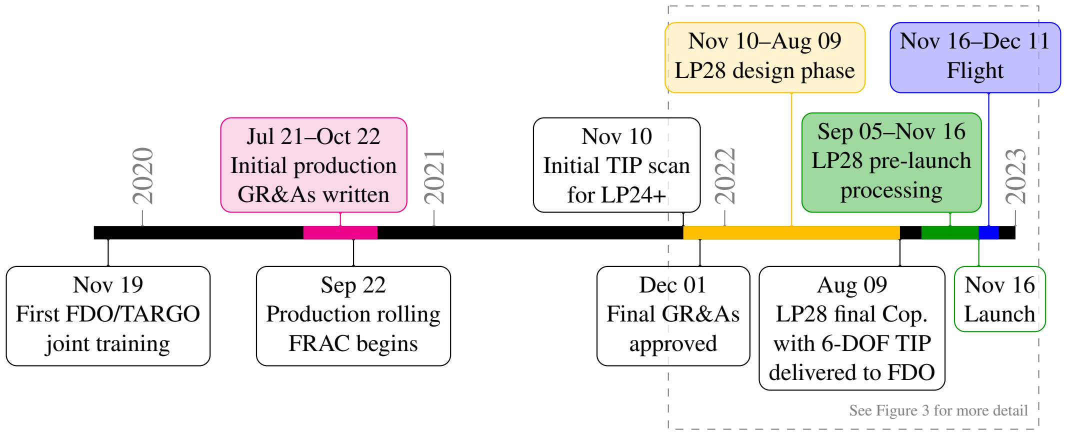

Figure 2: Long-term product development for LP28 mission design and trajectory data (dates in US Central Time).

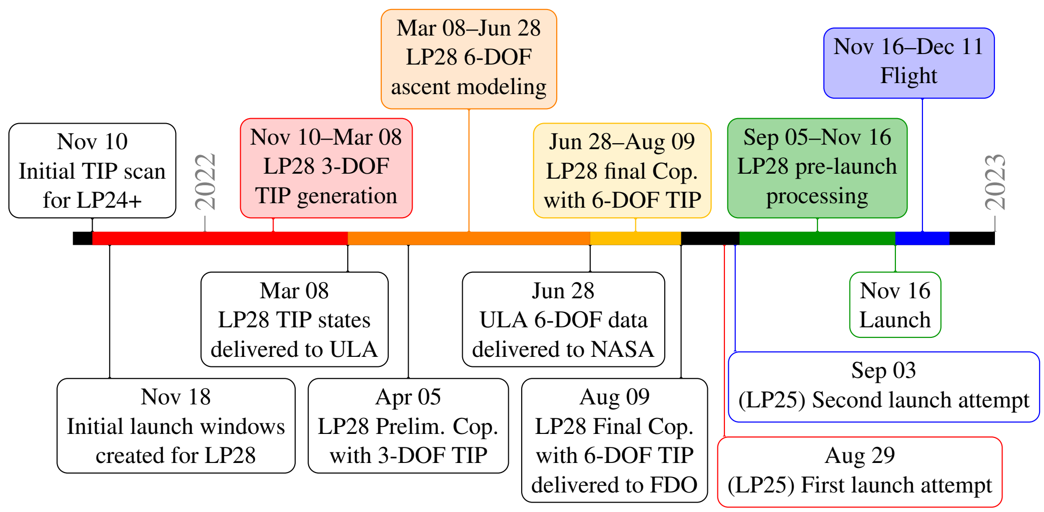

Figure 3: LP28 focused design and pre-launch data processing stages (dates in US Central Time).

TRAJECTORY RESPONSIBILITIES

The Artemis I mission design function was distributed among a number of different internal groups at NASA, with each phase of the trajectory design led by the corresponding vehicle trajectory design teams. Mission designers from the SLS team at the Marshall Space Flight Center (MSFC) led the design of ascent trajectories from liftoff to the Core Stage separation, after which the ICPS team from United Launch Alliance (ULA) was responsible for the on-orbit design through ICPS/Orion spring separation and ICPS disposal. A team of entry trajectory and guidance engineers in the JSC Engineering Directorate developed the Orion trajectory from atmospheric Entry Interface (EI) through splashdown. The Orion OPT, primarily consisting of astrodynamicists from the JSC Engineering Directorate, led the trajectory design for the remaining Artemis I mission, known as the *on-orbit phase*. Overall mission design and integration for Artemis I across the enterprise was managed by the Exploration Mission Planning Office (EMPO) at JSC, who reported to Exploration

Systems Development Mission Directorate (ESDMD) Mission Management at NASA Headquarters in Washington, DC.

Trajectories in each phase are designed in collaboration with the JSC Flight Operations Directorate (FOD) trajectory team members. The Flight Dynamics Officers (FDOs) are the flight controllers responsible for trajectory design, planning, and execution in the MCC in all phases of flight, including the execution of all Orion translational burns. The FDO console is shown in Figure 4. For Artemis missions, the FDO is supported in flight by a team of "backroom" personnel, including the Trajectory Officer (TRAJ) console and additional offline data processing support known as TRAJ Support.

Figure 4: FDO console during Artemis I. (Photo: NASA/Bill Stafford)

The process of a complete mission Design Analysis Cycle (DAC) requires technical coordination across all these trajectory design teams. The End-to-end Mission Performance Team (EMPT) is the cross-program team consisting of subsystem engineers from the EMPO, EGS, SLS, and Orion programs collectively responsible for Artemis trajectory design. The EMPT integrates the trajectory design for all phases of flight, coordinating between engineers developing the mission profiles for ascent, in-space, and entry. Each team performs their design and analysis with different software suites and varying levels of fidelity. The scope of EMPT includes the series of increasing-fidelity design and analysis cycles as well as the production and delivery cycles prior to pre-launch processing.

Since Artemis I trajectory operations were largely dominated by the Orion on-orbit trajectory, it was critical to guarantee a seamless process from the Orion reference trajectory design to operational in-space burn execution. In that sense, trajectory operations effectively began long before launch when the mission transitioned from conceptual development to production for particular launch dates. This necessitated particularly close collaboration between the full FDO team and the OPT at all stages of Orion trajectory design. In addition to designing the nominal Artemis I mission, this integrated trajectory design and analysis team also developed and maintained operational databases of off-nominal trajectory options to initiate aborts, early returns, and to respond to missed or failed translational burns.[6]

In order to ensure that reference trajectories were representative of feasible Orion burn plans supported by the vehicle, a set of formal Ground Rules and Assumptions (GR&As) were created. These GR&As were maintained by the EMPT so that all teams generated continuous, realistic trajectories. The 3-DOF reference trajectories had to be replicable by two different implementations of a TLT,[5] which will be discussed in the *Operational Burn Plan and Ephemeris Generation Process* section.

Copernicus, a generalized spacecraft trajectory design and optimization tool, was used to generate a propellant-optimal in-space trajectory that satisfies the numerous constraints of the Artemis I mission (e.g., minimum lunar flyby altitudes, required entry ranges and speeds, etc.).[1] Copernicus was the common software used throughout the Artemis I trajectory design community as the centralized representation of the 3-DOF trajectory design throughout its lifecycle.

Copernicus was also installed in the MCC such that trajectories could be re-optimized by the FDO in real time, if necessary. The complexity of the trajectory optimization workflow required members of the OPT to be able to support the real-time mission by providing the FDOs in the Flight Control Room (FCR) and other analysts in the Orion Mission Evaluation Room (MER), inside the MCC, with engineering expertise in trajectory design and optimization. This subset of OPT members, who had additional training to support operations from their console in the MER, were referred to as the Trajectory Analysis, Retargeting, and Optimization Officers (TARGOs). The TARGO console is shown in Figure 5.

Figure 5: TARGO console during Artemis I. (Photo: NASA/Robert Markowitz)

During the mission, the FDO team evaluated upcoming translational burns using actual vehicle conditions prior to each burn. In parallel, the TARGO team would ingest the vehicle's state and mass information into an updated Copernicus mission file and attempt to re-optimize the trajectory to inform FDO of any benefits of re-optimization. The FDO and TARGO teams supported the vehicle in this capacity continuously for the entirety of the mission. TARGO real-time mission design support became necessary when a long-duration translational burn was inserted into the mission timeline as a previously unplanned Developmental Flight Test Objective (DFTO). In addition, TARGOs would evaluate the abort solution space, based on available propellant loads and current operational plans.

DESIGN PHASE

3-DOF TIP Generation and SLS Performance Windows

The mission design process began with generating simplified Copernicus trajectories from which Target Interface Points (TIPs) could be sampled. The TIP is designated as a state vector on the Orion trajectory shortly after separating from the ICPS and is the point SLS and ICPS target during the ascent phase. TIPs were collected into reference databases representing these targets for a sweep of launch opportunities used in later mission design processes.

TIPs, derived from trajectories that started in a designated parking orbit, were solved for an optimal Artemis I mission lasting about 26 days for a launch every day of the year. These Copernicus files initially used a simplified model of the ICPS performing the TLI. The initial design trajectories also excluded the OTC-1 burn. OTC-1 would be inserted later to avoid recontact with the ICPS along the outbound leg and lunar flyby.

The baseline model of this Copernicus trajectory was designed by the OPT at JSC and then sent to the Mission Analysis and Design Team at MSFC. This team used Program to Optimize Simulated Trajectories II (POST2) to optimize ascent trajectories for a given sequence of events.[7] This was used to generate a series of tables representing the insertion state vector after the SLS Core Stage separation and other related ascent parameters known as the *hypergrid*.

A Copernicus plugin[8] parsed these hypergrid files into multi-dimensional splines of the data to interpolate the parking orbit insertion state vector. These interpolated hypergrid state vectors served as the initial state for the TIP Copernicus trajectories. The MSFC team also added higher-fidelity

3-DOF models of ICPS trajectory segments, including adding the PRM that targets the pre-TLI parking orbit. Finally, the MSFC team generated optimal TIP trajectories in one minute increments for every SLS performance-feasible minute throughout a year. This was done using Damocles, a collection of Python scripts for running multiple Copernicus instances in parallel on a computer cluster to generate Artemis trajectory "scans" written by the OPT at JSC.[2] The spans of SLS-feasible minutes resulting from this TIP scan were known as the *SLS performance windows*.

3-DOF-TIP Orion Performance Windows

Once the MSFC team generated the TIP Damocles scan, the SLS hypergrids and Copernicus files from the scan were delivered back to the OPT to refine the Orion trajectories into the "true" missions from which the launch windows could be selected. At JSC, the TIP states were extracted into a data file and splined with another Copernicus plugin. The Copernicus optimization was then updated to use the splined TIP states as the initial state.

Now that the TIP states were known, and thus the trajectory from which the ICPS would dispose, OTC-1 was added to the Copernicus optimization problem to avoid recontact between Orion and ICPS. OTC-2 was also added to provide additional control variables for flexibility in the optimization of the outbound trajectory, but was heavily weighted in the objective function to discourage its use except when additional controls were needed for smooth convergence.

Finally, Misson Selection Logic (MSL)[2] was activated in Damocles. For Artemis I, MSL was an algorithm that provided the capability to adapt the mission duration from a minimum of approximately 26 days to a maximum of approximately 42 days. MSL then pared down the viable mission duration and selected a mission duration based on a set of criteria, which included an assessment of Orion vehicle constraint violations. It used a simple machine learning model to autonomously determine the likelihood of success at mitigating these violations. If successful mitigation was likely, MSL autonomously inserted trajectory mitigations into Copernicus to attempt to conform to Orion constraints and recover those missions.

A new Damocles scan with these updates, referred to as the "MSL scan," resulted in a set of Copernicus files representing Orion-feasible Artemis I trajectories of varying mission duration, each optimized for minimum propellant usage. The span of resulting Orion-feasible minutes from this scan were known as the *Orion performance windows*.

The performance windows could range from instantaneous to a natural maximum of about four hours in duration. However, due to operational limitations, for any given launch attempt, the longer windows were trimmed to a maximum duration of two hours. A set of elimination and selection criteria were developed based on priorities from stakeholders across the Artemis enterprise in the EMPT. Based on this criteria, an automated launch window selection process was developed called Window Shopper. This trimmed down the Damocles MSL scan dataset to only include the *selected windows*. This subset with selected windows represented the preliminary mission designs using the MSFC-modeled 3-DOF ICPS TIP.

The first and last valid minute of every launch window are referred to as the window *open* and *close*, respectively.

Rolling FRAC and 6-DOF TIP Generation

The Flight Readiness Analysis Cycle (FRAC) was initiated sequentially for each *Launch Period (LP)* to generate its operational reference trajectories. Launch Periods (LPs) are defined as a consecutive set of SLS performance windows in the same sidereal lunar month. Each LP is numbered; for Artemis I, the LP number was incremented when there were more than seven days between SLS performance-viable launch windows. This results in LPs lasting about two weeks each sidereal lunar month. LP1 began November 6, 2020; Artemis I launched in LP28 on November 16, 2022. The process of evaluating multiple LPs in parallel, with each new LP cycle beginning approximately monthly, was known as *rolling FRAC*. Since the launch window scans described in the above sections only spanned one year at a time, these scans were repeated as necessary to ensure continuous production during the rolling FRAC period.

For a given LP, FRAC began approximately 8 months prior to the first window of the next LP. The OPT extracted the TIP states from the preliminary selected windows dataset for that cycle's LP and delivered them to MSFC. The SLS mission design team re-packaged the selected TIP states with their corresponding 6-DOF SLS ascent trajectory data and delivered the full data package to ULA. ULA then used this package to generate the 6-DOF ICPS trajectories that target the selected TIP. These were used to develop the targets for the ICPS FSW guidance. After generating the targets to use for launch, ULA sent back the 6-DOF trajectories to MSFC and JSC from which final TIP states and times could be extracted.

6-DOF-TIP Orion Performance Windows

After final 6-DOF TIP states were delivered about three to four months before a given LP, the OPT would update the TIP data files and ingest into new Copernicus nominal mission files re-optimized for each valid launch time at a one-minute density. These final design nominal reference trajectories were delivered to the FDO team for processing into the MCC.

The final nominal reference trajectories were also sent to mission analysts at Lockheed Martin (LM), the Orion prime contractor, for power generation analysis. This "power screening" was used to flag launch dates that may result in conditions with unfavorable power generation, especially when considered in various battery failure scenarios evaluated to ensure power system robustness. Further detailed analysis of mission opportunities with degraded power could ultimately result in elimination of additional launch dates from the trajectory-feasible selected windows.

The selected launch windows were communicated to the EMPO in a document produced by OPT analysts known as the Launch and Landing Table. New revisions were published throughout rolling FRAC as the launch windows were refined. Dates recommended for elimination due to power screening were also noted in this table.

Additionally, these final nominal trajectories were used to build initial databases of off-nominal Copernicus solutions and subsequently delivered to FDO. These databases were generated based on the final nominal reference trajectories every 30 minutes through a launch window and for each window open and close.[6]

Correction Burns

Figure 1 illustrates the typical translational burn schedule for an approximately 26-day mission. Longer missions would contain additional correction burns and/or DRO maintenance burns so that Orion would never go longer than about 4.5 days without a trajectory correction. The actual Artemis I launch on November 16, 2022, resulted in an approximately 26-day mission, so the additional corrections were not added.

Correction burn scheduling was originally inspired by Apollo precedence and was informed by linear covariance analysis of Orion navigation performance of an early-stage conceptual design reference mission. The final timings were ultimately adjusted for operational considerations to make pre-flight timeline development more consistent between launch dates.

With the exception of the Orbital Maneuvering System Engine (OMSe) Checkout burn at Outbound Trajectory Correction (OTC)-1 and the optimization smoothing control OTC-2, correction burns were not modeled in Copernicus prior to flight. These were added as neeeded throughout the mission.

OPERATIONAL BURN PLAN AND EPHEMERIS GENERATION PROCESS

The Ground Two-Level Targeter (GTLT), used in the MCC to generate targeting inputs in the burn plan and verify onboard targeting results, was developed by the FDO team. Independently, the Orion FSW implemented the Onboard Two-Level Targeter (OTLT).[9] While based on the same underlying concepts, the independence of these TLTs afforded some protection against coding or logic errors by ensuring consistent results. Each is referred to in this paper distinguished by the "ground" or "onboard" qualifier or appropriate initialism.

While the burn plan contains many general vehicle configuration and automation parameters that are trajectory independent, the core of the operational burn plan consists of the inputs to the OTLT. As a differential correction scheme, a TLT requires an initial guess of the desired trajectory. These initial inputs are crucial to ensuring targeter convergence and an operationally reasonable and acceptable solution. The targeting inputs for an Orion burn plan are derived from the FDO burn targeting process used to develop the MCC trajectory *ephemeris*, or time history of state vectors. An ephemeris, in this context, always refers to an integrated trajectory simulated in nonlinear force models. Such a dynamical model is therefore sometimes referred to as an *ephemeris model* to distinguish it from analyticial models such as the Circular Restricted Three-Body Problem (CR3BP) or Keplerian two-body dynamics.

For Artemis I—though the process will likely be similar for future Orion missions—the process of building an Orion burn plan for uplink from a 3-DOF reference trajectory started with Copernicus. Nominal and off-nominal trajectories were generated by the OPT and delivered to the FDO team for review and processing. The FDO team would select a subset of reference trajectories for conversion into Orion burn plans and ground ephemerides, which would then be compared to the Copernicus-optimal reference trajectories.

A flowchart depiction of the entire process described in this section is shown in Figure 6. This process was employed both throughout real-time operations in flight as well as during pre-launch processing. During pre-launch processing, this was used to pre-generate the targeting plans for a selected subset of cases, so only the steps up to building the burn plan were followed. A pre-launch targeting plan could then be deployed during the flight by picking up from the burn plan building steps.

Optimization

The Copernicus mission files served as the input to the FDO burn targeting process. FDO made any necessary updates to the trajectory in Copernicus relevant to real-time operations such as current Earth orientation parameters or vehicle properties. Most importantly, FDO updated the Orion initial state vector with the best "current" navigation source, typically derived from ground tracking via

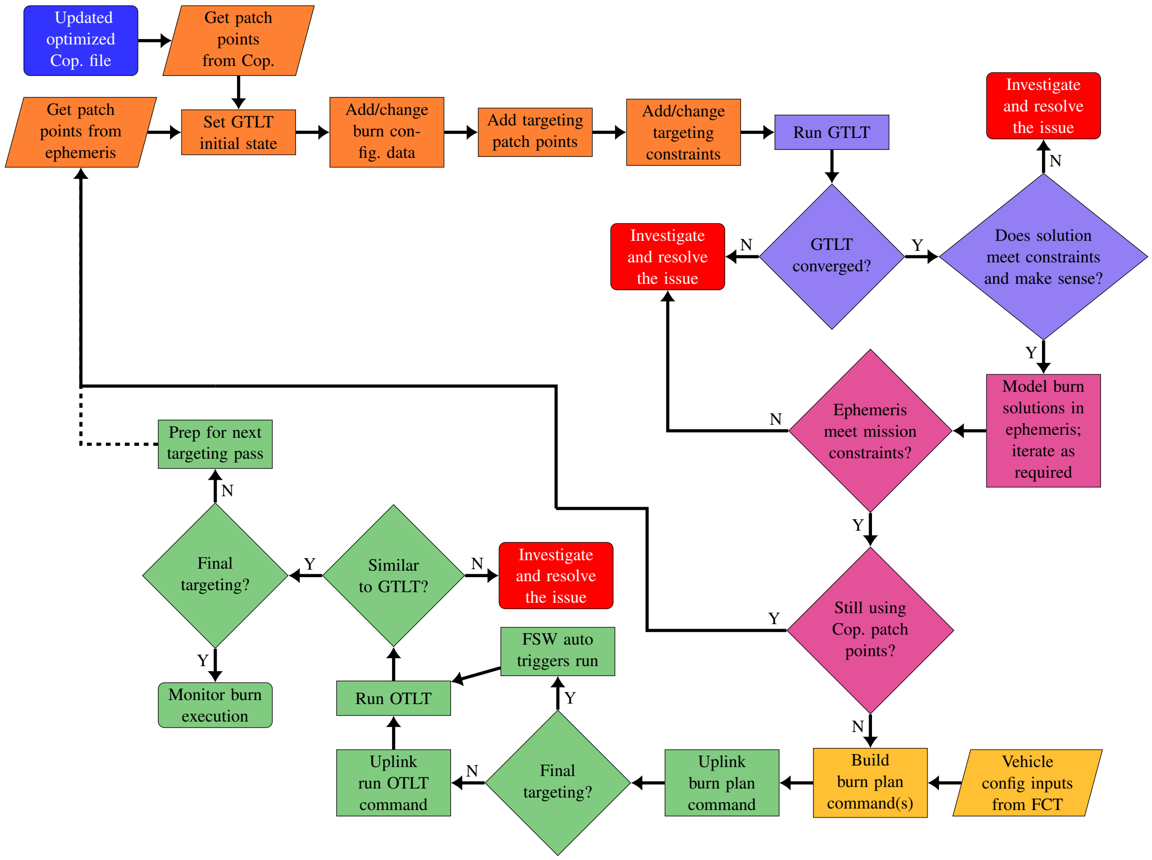

Figure 6: Flowchart of operational burn plan and ephemeris generation process. Colored blocks are discussed in the following sections: Blue, Optimization; Orange, Patch Point Selection; Purple, Linearized Trajectory Smoothing; Pink, Ephemeris Iteration; Yellow, Burn Plan Generation. Green blocks are part of the burn execution timeline discussed in the Post-TLI Operations section.

the Deep Space Network (DSN). If the file contained trajectory segments chronologically earlier than the time of the current navigation state, these segments were “frozen” so that their controls and constraints were removed from the optimization problem. Subsequently, Copernicus re-converged using the file with these updates to provide a more optimal 3-DOF idealized reference trajectory subject to the updated real-time conditions.

Patch Point Selection

The vehicle state and time was sampled from the 3-DOF reference trajectory at a handful of discrete points, referred to as patch points. Patch points represented the nodes in the TLT where constraints could be applied and evaluated. Burns could also be performed at designated patch points as control variables to connect the coast arcs between patch points. As such, these included planned upcoming translational burns or waypoints to guide the solution.

Since correction burn placement for Artemis I was timeline-driven rather than by trajectory geometry or shape, some launch dates could result in the arc between two patch points subtending a 180° angle inertially about the Moon while it was the dominant body. Targeting transfers through such angles about a dominant body are known to present convergence difficulties for targeters, a common example being a Lambert targeter, due to the planar ambiguity of solutions between the two points when no plane is specified. Gradient-based targeters, such as the TLT, are no exception to this. Thus, when near-180° conditions about the dominant body were identified for a given reference trajectory, an additional non-burn waypoint, called a "targeting patch point," was inserted along the offending arc.

The primary engine to perform the burn was also specified in the patch points. Engine selection depended on the estimated Δv. Copernicus did not provide estimates for correction burns, so these typically were set to a mid-level thrust option by default and re-evaluated during the FDO targeting process in the full ephemeris model before building the final burn plan.

Finally, constraints must be applied to the TLT patch points to provide the targeter conditions to achieve. In the OTLT implementation, position and velocity continuity constraints are assumed depending on the type of patch point.[9] Any necessary additional explicit non-continuity constraints, then, must be specified for the patch points in an Orion burn plan.

The constraints used in Copernicus may not be the same constraints applied to the TLT patch points. The TLT is not an optimizer and is fundamentally solving a different problem than Copernicus: the TLT seeks a feasible, rather than optimal, trajectory near the reference. Copernicus mission files from TARGO were designed with a set of standards and GR&As such that a mission design should still converge to a similar solution when modeled in TLT by accounting for these differences and limitations between their respective available trajectory constraints.

Linearized Trajectory Smoothing

Once an initial set of patch points were developed, they were run through the GTLT. This ensured that the constraints modeled in the GTLT resulted in a convergent feasible trajectory. FDO inspected the solution to verify it met vehicle constraints, certification limits, and mission objectives.

Due to modeling differences between Copernicus and the GTLT, continuity constraints may not always be immediately satisfied when seeding the GTLT. Therefore, to continue the targeting burn plan development process, a first pass through GTLT initialized by Copernicus-derived patch points smoothed out the trajectory solution in its linear space by satisfying the continuity constraints. This provides an initial guess for subsequent TLT passes so that they would require fewer iterations to converge.

Ephemeris Iteration

With the initial smoothed solution from GTLT based on the Copernicus trajectory, its computed burns were added to a high-fidelity ephemeris model. The trajectory was then evaluated to ensure that mission objectives were met and that vehicle constraints and requirements were satisfied. The buildup of constraint error tolerances in the TLT, numerical noise, and linearization introduced enough perturbation that the solution required additional differential correction. Even when constraints were satisfied by targeting patch points derived from an ephemeris, the entire process was always repeated at least once using the new ephemeris as the reference trajectory. This provided the vehicle the highest-fidelity initial guess. This series of repeated passes through the GTLT-to-ephemeris cycle in order to achieve satisfactory, convergent solutions was referred to as "ephemeris iteration."

Burn Plan Generation

After ephemeris iteration, the final solution was transferred to the active MCC trajectory and all trajectory products were re-delivered to various operational customers, such as the DSN. This final trajectory and the GTLT targeting data were then used to generate command parameters for the Orion burn plan. Additional vehicle configuration parameters were input by other systems flight controllers and collected into the final commands to uplink to the vehicle by FDO.

As time allowed, prior to uplinking the commands to the vehicle, a copy of the commands were sent to guidance engineers in the MER, who would run the commands through an emulation of the OTLT to assess how the actual FSW would respond to the command. This provided valuable insight in catching possible misconfiguration or poor TLT inputs that may cause the actual OTLT to fail to converge or other downstream guidance system issues.

PRE-LAUNCH PROCESSING

The OPT deliveries of the final Copernicus nominal mission files to the FDO team marked the transition from the design phase to the pre-launch processing phase. Data in this phase was also processed approximately monthly in batches per LP following the rolling FRAC deliveries.

Off-Nominal Trajectory Eclipse Screening

Once off-nominal trajectory databases had been developed for every 30 minutes of the launch window, TARGO analysts focused on off-nominal trajectories during the pre-TLI phase of the mission. During the pre-launch data processing for the first launch attempts in LP25 (August 23, 2022–September 6, 2022), TARGO analysts identified a vulnerability that a partially-completed TLI in this phase could result in particularly long periods of solar eclipse that could exceed Orion's solar array constraints and result in loss of the vehicle. Prior to LP25, this constraint had not been considered in determining launch availability. The Artemis I Mission Management Team (MMT) decided to accept the late-stage recommendation to eliminate launch dates for which this condition was present from further consideration, notably eliminating a previously-announced potential launch opportunity on September 4, 2022, and acknowledged at the post-scrub press conference for the first launch attempt. Subsequent launch periods were screened for this condition prior to entering the pre-launch data processing phase. Future missions will attempt to evaluate this condition earlier in the design phase rather than screening during final processing.

Flight Control Team (FCT) Product Generation and Resource Planning

The FCT generated products for the prime launch opportunity and for a subset of backup launch opportunities. The process used by FDO to generate the trajectories for these launch opportunities was very similar to the typical burn plan generation process discussed above in the *Operational Burn Plan and Ephemeris Generation Process* section. These trajectories were used by the FCT to develop mission products, such as the mission timeline, attitude timelines, power profiles, propellant timelines, communications and tracking schedules, etc. These products were also used externally by the communications and tracking network for resource scheduling and planning.

Due to the amount of time it took to run the pre-mission production process, only specific trajectories were generated by the FDO team, thus not every launch minute within a targeted launch period was produced. Additionally, pre-generated burn plans for a handful of early-mission abort scenarios were produced for each window open and close to have on hand and uplink quickly should they be needed.

In all, generating the suite of burn plans and trajectory products took the FDO team about two weeks to complete, with the final products distributed to the rest of the FCT and networks about two months prior to the targeted launch epoch.

Inertial and Earth-Moon Rotating-Pulsating (EMRP) Burn Plans

Ordinarily during flight, the state vectors associated with each patch point in the burn plans were sampled as Earth-centered inertial Cartesian states. In order for a single burn plan to be reasonably accurate across a full launch window, however, inertial states could not be used. To address this, both the GTLT and OTLT support burn plans with state vectors in the EMRP frame [10, p. 588] in addition to inertial frames. To improve numerical performance of calculations in this frame, the instantaneous Earth-Moon distance is scaled[11] by the mean lunar semi-major axis of 384,400 kilometers.

Burn plans in the EMRP frame allowed the trajectory to keep the same Earth-Moon relative geometry throughout a launch window. Since failure cases could result in a total loss of communications with Orion during the ascent/TLI phase of the mission, Orion needed to have a burn plan onboard that would work anywhere in the launch window. Although not optimal, burn plans in the EMRP frame allowed Orion to continue onward to the Moon regardless of launch time. If communications were regained with Orion, an updated, optimized inertial plan could then be uplinked to the vehicle to continue the mission in a more optimal fashion than remaining on a EMRP plan.

In addition to the inertial burn plans generated in support of FCT product generation, FDO generated burn plans in the EMRP frame prior to every launch attempt. The EMRP burn plan was uplinked to Orion during the pre-launch countdown sequence. If the MCC had communications with Orion prior to the first Orion burn, the MCC re-optimized the trajectory at the specific launch epoch and uplinked an updated inertial burn plan prior to Orion's first burn.

One noteworthy feature of the burn plan FSW is that the time of each patch point is stored onboard as seconds since liftoff, where the assumed liftoff time was not included in the uplinkable commands. Therefore, any software interfacing with burn plans must be configured with a liftoff time consistent with that used onboard. This had the benefit of having one EMRP burn plan at liftoff that could be used for any events specified relative to launch for that particular launch window.

However in practice, the Copernicus optimal reference trajectories had timing variations for each of the major burns, which could have impacted optimality through the launch window. In particular, the mission design GR&As specified that the OPF would be fixed in absolute time throughout the window. If a mission had launched at the close of a two-hour window with no update to the onboard burn plan from the launch pad, the OPF burn timing could have been two hours late relative to the optimal solution. However, because of the way the TLT constraints were configured for the OPF patch point, the OTLT would have likely still converged in most cases to a solution arriving at the target DRO, albeit more expensive, given that the TLT is only a minimum-norm solution method for finding feasible trajectories rather than optimal. FDO ran the EMRP burn plan using the expected post-TLI trajectory at the open and close of the daily launch window to ensure the GTLT convergence and propellant performance using that plan.

In the week prior to a specific selected launch day, both the launch window open inertial and EMRP burn plans were sent to Orion MER analysts to review and concur with FDO on the proper configuration.

REAL-TIME OPERATIONS

Day of Launch

Around-the-clock trajectory operations began 36 hours prior to the launch window open. TARGO used the timing of each correction burn to generate an initial TARGO Operations Report (TOR) for every 30 minutes through the launch window. The TOR was an interactive spreadsheet that contained a summary of every pre-generated off-nominal trajectory in the delivered databases and allowed FDO and TARGO to filter the options based on criteria such as remaining available propellant.[6]

FDO and TARGO teams reported to their stations in the MCC 24 hours prior to launch window open. The pre-generated EMRP burn plan was uplinked to the vehicle, and then both teams reviewed the available TORs throughout the launch window for the pre-TLI phase to enable making decisions quickly. Ascent abort trajectory options were also reviewed, though those trajectories were designed and analyzed in a separate process that did not use the workflows described in this paper.

TARGO analysts on console also evaluated a scenario with a passive Orion (i.e., no Orion translational capability) at various points in the mission, propagated for 100 years to get a sense for the character of these trajectories should the vehicle fail to be restored to nominal operation. These were single-run, low-fidelity studies with no dispersions or solar radiation pressure; the intent was merely to see the character of the long-term trajectory behavior, knowing that the true trajectory could vary significantly given the multi-body dynamics that govern the cislunar and heliocentric space. Results for different launch times, even in the same window, could vary simply because of the slight differences in Earth, Moon, and Sun geometry and in the TLI targets throughout the launch window. However, it was assumed the observed behavior at window open would provide a reasonable approximation for a propagated trajectory for any time in the window given the similar geometry of the day. For the November 16, 2022, launch window open, not performing the OPF burn would have entered heliocentric space with returns to the Earth-Moon system in 2064, 2087, and 2118 for the idealized trajectory, with that final encounter resulting in an Earth impact.

A similar analysis repeated with a nominal mission up to, but not performing, the DRI burn and propagated for a century showed chaotic behavior in the Earth-Moon system inside of the zero-velocity surface (i.e., "forbidden regions" defined by the Jacobi Constants in the CR3BP) before exiting into heliocentric space through the L2 portal after eight years with no observed returns to the Earth-Moon system. Finally, remaining in the DRO without performing DRD and no further correction showed a natural departure of the DRO after only three years before going heliocentric with no observed returns to the Earth-Moon system.

The TARGO team evaluated the minimum percentage of the estimated TLI burn duration that must be performed in order to meet the minimum re-entry speed to adequately assess the performance of the Orion heat shield representative of lunar returns. The exact percentage would vary through the launch window; for November 16, 2022, the percentage at window open was around 79%. It was later found after launch that the percentage decreased to between 76% and 77%.

Finally, the TARGO and FDO teams analyzed the minimum percentage of the estimated TLI burn duration that must be executed in order to complete the nominal mission should Orion need to perform a recovery from an "underspeed," the condition of insufficient Δv from a burn to otherwise achieve the mission on the current trajectory. For the November 16, 2022, window open, 99.2% and 99.4% were needed depending on whether the OMSe or eight +X Auxiliary engines (8+X) were used, respectively, with similar results for the eventual actual launch time.

Once all teams polled Go for launch and the launch sequencer started that triggers the choreographed steps for the launch, a precise target launch time was known, and all trajectory teams began a series of preparatory activities.

The FDO team retrieved the nominal Copernicus file for the nearest minute to the target launch time from the archive of delivered OPT trajectories. FDO then re-optimized the trajectory with the retrieved Copernicus file, performed ephemeris iteration, and built a new inertial burn plan for Orion. After liftoff, FDO updated the MCC ground software reference liftoff time to the Orion-sensed liftoff time. In parallel, FDO requested TARGO to regenerate a baseline nominal Copernicus trajectory with the updated reference liftoff time.

As the mission progressed through TLI, FDO updated the MCC trajectory by targeting Orion burns on the predicted post-TLI state to put Orion back onto the optimal path. After TLI was complete, FDO transferred the post-TLI plan into the MCC trajectory and uplinked the burn plan to Orion in support of the upcoming OTC-1 burn.

Post-TLI Operations

FDO consistently updated the MCC ephemeris for the nominal mission based on updated navigation and optimization results. Meanwhile, TARGO would independently continue re-optimizing the mission in Copernicus and comparing to the nominal burn plans. These re-optimized nominal trajectories would also be the reference trajectories for updated off-nominal databases of thousands of Copernicus mission files for a variety of permutations. These databases and TORs were regenerated and delivered to FDO throughout the mission based on actual mission performance. FDO would then select a handful of solutions from the TOR and pre-generate burn plans for those options to ensure there was always at least one abort burn plan that could be pulled and uplinked, if necessary. This required teams of additional OPT office staff and TRAJ Support teams working offline to design and generate the databases, trajectories, and burn plans for the abort options selected by FDO using the process outlined in the Operational Burn Plan and Ephemeris Generation Process section.

Five hours prior to a burn, FDO would begin executing the burn targeting and execution process in order to incorporate the latest navigation state into the OTLT initial guesses and update any vehicle configuration parameters. The burn execution timeline was designed to allow adequate time to perform three onboard targeting passes, known as “preliminary,” “intermediate,” and “final” targeting. Burn plan updates and re-uplink opportunities existed just prior to each of these targeting passes. The targeting pass executes the OTLT and, subsequently, the onboard Orbit Guidance flight software (OrbGuid) that computes burn steering parameters for the next burn in the plan. Preliminary targeting is always performed and must be manually triggered by ground command. Intermediate targeting is only performed as needed and, when performed, must also be manually triggered by ground command. Final targeting occurs automatically at a time specified in the burn plan prior to the burn, typically set to 31 minutes. FDO monitored burn targeting and guidance parameters in real time using Orion telemetry and notified the FCT of its progress when communication was acquired during the burn targeting process.

TARGO teams monitored FDO during the entire burn execution timeline, and were only engaged if there were issues with Copernicus during burn plan generation that required troubleshooting or mission design intervention. Once the burn completed, its performance was evaluated by FDO using measured acceleration as well as updated ground tracking as available. Immediately following the completion of a burn, FDO and TARGO teams estimated the next burn and re-optimized the downstream trajectory to prepare for the next burn plan generation cycle.

The FDO team assessed the need for re-optimizing the trajectory once a day. If the propellant savings warranted, the ground trajectory was updated with the new plan and sent to Orion. FDO also updated the upcoming abort opportunities for the next 48 hours to communicate to the FCT what options would be considered if a failure occurred requiring an abort to get Orion back to Earth quickly.

MISSION PERFORMANCE

Launch

Figure 7 shows a timeline of the key trajectory events of the mission as actually performed. The closest nominal mission file delivered to FDO by the OPT was for launch on November 16, 2022, at 06:48:00 Coordinated Universal Time (UTC). This would be the baseline of all subsequent trajectory products until a new file was generated post launch. After liftoff, FDO displays showed an Orion-sensed liftoff time of 06:47:43.643 UTC and was used by TARGO for subsequent re-optimization. It was later determined that a configuration issue in the MCC resulted in a slightly different launch time than was measured and used onboard, which was 06:47:44.119 UTC. This difference did not impact trajectory performance during the mission.

Immediately prior to liftoff, TARGO began assessing which off-nominal reference dataset to use based on the 30 minute increments. For the final launch attempt, the target launch time fell nearly halfway between the two thirty-minute sample points with a slight bias to the earlier one, thus the earlier solutions would be used to make decisions until new datasets could be generated with exact launch timing. Regenerating the off-nominal datasets had an estimated completion time of about 16 hours after launch. It was cautiously assumed that options could be assessed with both datasets and that solutions would be continuous between them should an abort be needed in the interim.

Star Tracker Abort Options

A number of issues with the star trackers early in the mission, particularly the poor performance later determined to be caused by "dazzling" from the Reaction Control System (RCS) exhaust plumes and larger than expected debris cloud,[12,13] initially led the FCT and MER teams to discuss whether the degraded navigation performance would require an early return should there be no resolution to the problems. It was determined that the onboard navigation filter would take about 5.5 days to degrade before safe return to Earth would not be possible.

TARGO was tasked to provide abort options that would result in splashdown before onboard navigation would render the vehicle unrecoverable. This would determine the latest decision point that would still result in successful Orion recovery. At that point, it was unclear to the rest of the team if abort options beyond the first abort option with a burn plan pre-generated before launch by FDO existed with returns prior to the navigation filter degradation. The TARGO team resorted to the assumption that the two thirty-minute off-nominal TORs could be bridged via interpolation to determine the latest decision gate to execute an abort. TARGO reported that, using the data in hand, there were dozens of viable options to perform an early return with decision points at least three days away, which would buy time for the trajectory teams to refine off-nominal trajectory

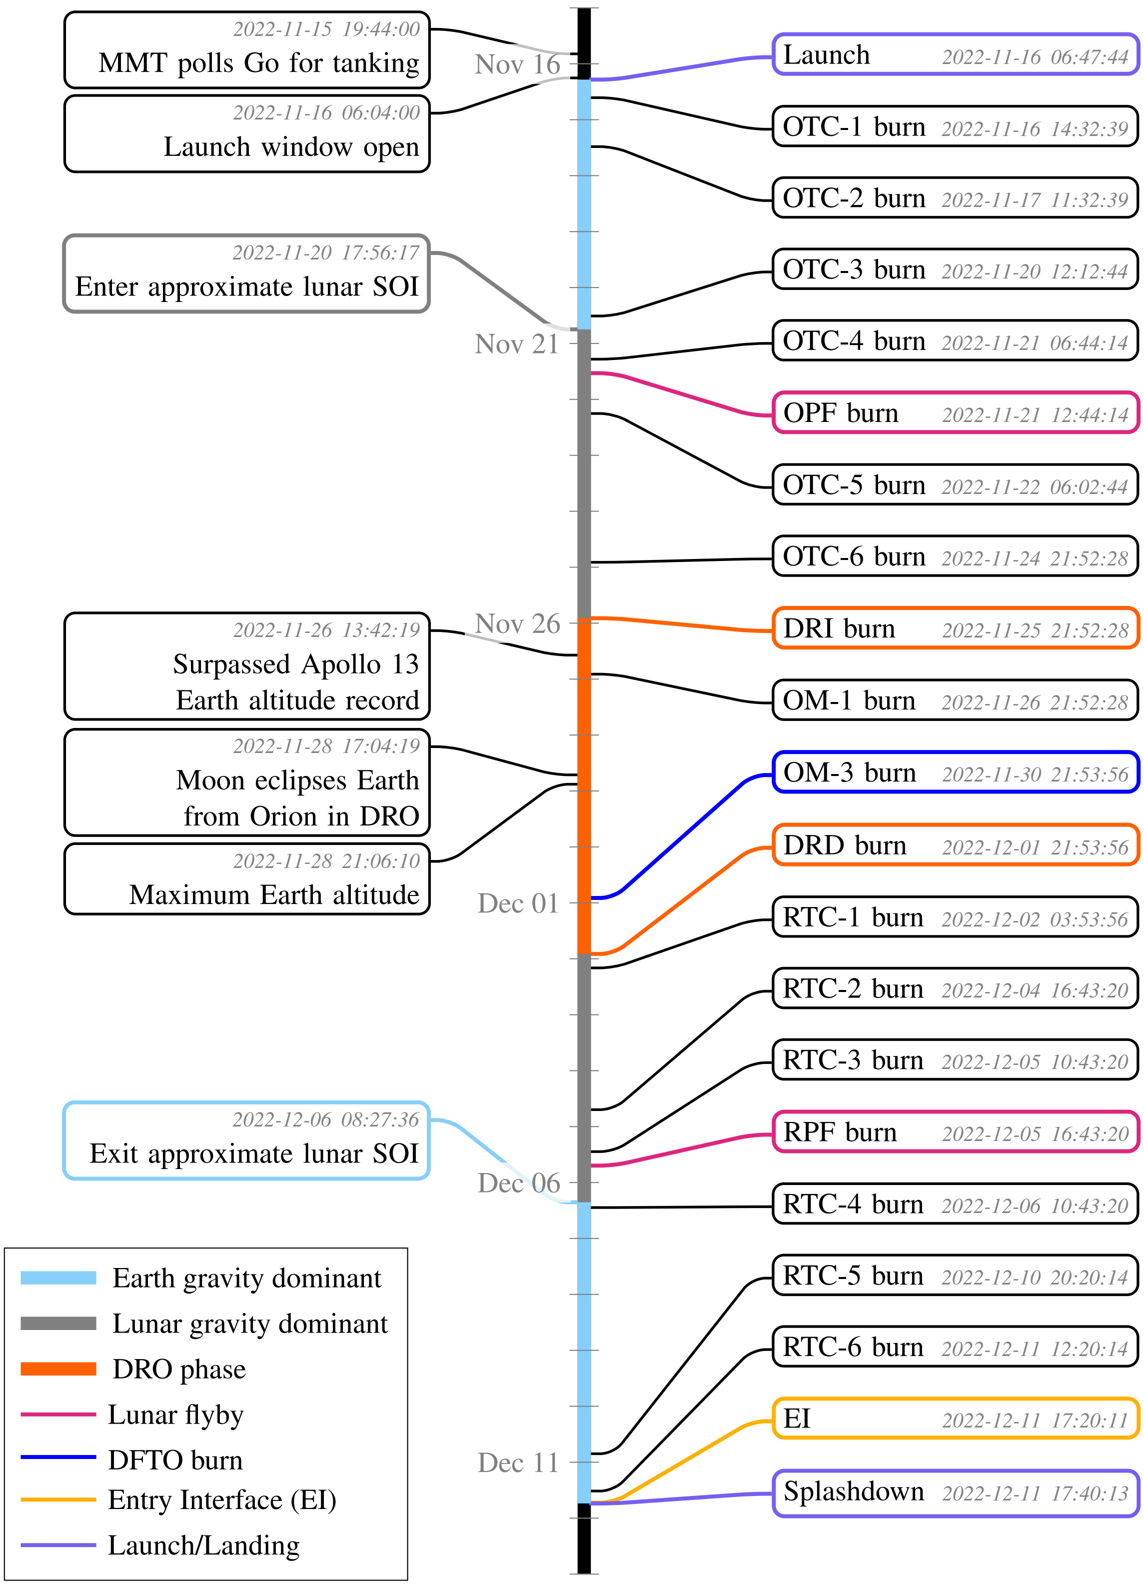

Figure 7: Artemis I mission timeline. Time of Ignition (TIG) is listed for burns. All dates/times are in UTC.

datasets and for the navigation teams to assess the star tracker and navigation filter performance. The MMT subsequently decided to forego the first pre-generated abort burn plan option and defer further decisions until the next business day. Ultimately, the issue was resolved and the mission progressed nominally.

Modifying the Mission Design in Flight

Most of the mission continued to use the same nominal baseline Copernicus mission design with the actual launch time generated right after liftoff with only minor tweaks based on as-flown performance. However, once Orion entered the DRO, the MMT invited teams to bring forward further proposals for improvised DFTOs that would buy down risk for future missions. One relevant MER proposal intended to test Orion automation paths that were not otherwise expected to execute during the mission, namely, demonstrating that Orion could successfully execute burns with only six of its +X Auxiliary engines rather than using all eight, as is the standard usage for this thruster configuration when used as a backup to the OMSe or alternative to the Service Module Reaction Control System (SMRCS) for correction burns. It was conceived that one of the upcoming correction burns be leveraged to test this, though it would require a lengthy burn duration to fully demonstrate the vehicle controllability during this mode of operation. To exercise the engines, it was proposed to fix the burn duration to 100 seconds.

TARGO analysis as part of the MER team assessed various proposed locations to perform the test. This analysis showed that performing such a large burn in the DRO would result in DRO escape in less than 180 days, while doing it after RPF would jeopardize meeting the re-entry corridor. Similar concerns existed about executing this after DRD, which could upset the powered flyby trajectory without adequate time to re-optimize. FDO recommended to the MER team that performing this DFTO while still in the DRO regime, particularly at the planned Orbit Maintenance-3 (OM-3) location scheduled 24 hours prior to DRD, would minimize the impact on the nominal mission and EI performance. FDO also requested that TARGO find a direction for the burn effective Δv that minimized the cost difference between the nominal mission and the mission with the DFTO.

MER planning teams decided the DFTO would be applied to the OM-3 burn per the FDO and TARGO recommendations. This, however, meant that Orion would effectively leave the planned DRO 24 hours earlier than expected. The fixed burn duration of 100 seconds also provided insufficient Δv alone to target a RPF, even in the most optimal direction. This required an impromptu redesign of the nominal mission to optimize OM-3 along with DRD and RPF to still satisfy the re-entry corridor and the heat shield demonstration. This demonstrated the importance of having active involvement of the mission design team in the MER during the flight and their tight integration with their FDO team counterparts.

Optimization, Targeting, and Execution Performance

Prior to launch, the OPT generated millions of reference trajectories, nominal and off-nominal, that made it possible to identify the viable launch opportunities. During the preparation for, and execution of, the final launch date alone, the OPT generated 2,502 new Copernicus mission files for off-nominal trajectories in addition to the four nominal baseline reference trajectories used in pre-launch, launch, post-TLI, and OM-3-onward operations. FDO generated hundreds of burn plans during the mission based on a subset of those reference trajectories, on top of the hundreds-to-thousands done for prior candidate launch dates.

Table 1: Summary of Orion Δv Performance for Artemis I (ft/s)

Burn	Design phase		Real-time optimization			Real-time targeting		Execution	
	Cop. with 3-DOF TIP	Cop. with 6-DOF TIP	Cop. with 6-DOF TIP	Cop. with OM-3 DFTO	Real-time Copernicus re-optimized	Final ground target	Final onboard target at TIG	Real-time sensed Δv	Post-flight review
UTC liftoff	*2022-11-16 06:48:00.000*		*2022-11-16 06:47:43.643*				*2022-11-16 06:47:44.119*		
USS-1[♣]	5.50	5.50	5.50	5.50	5.50	*Burn not targeted.*		5.09	n/a
OTC-1[♢]	113.24	113.24	113.24	113.24	113.24	114.72	114.50	114.65	114.67
OTC-2[♡]	0.45	0.45	0.44	0.44	0.70	0.73	0.88	0.73	0.73
OTC-3	0.	0.	0.	0.	0.01	2.96	2.95	2.86	2.87
OTC-4	0.	0.	0.	0.	1.32	0.98	0.85	0.71	0.71
OPF[*]	586.22	586.22	586.22	586.22	587.32	586.08	585.87	586.06	586.09
OTC-5	0.	0.	0.	0.	2.23	3.25	3.34	3.21	3.21
OTC-6	0.	0.	0.	0.	9.06	9.05	8.94	8.79	8.79
DRI	362.48	362.48	362.48	362.48	362.78	362.89	362.86	362.97	362.99
OM-1	0.	0.	0.	0.	0.51	0.42	0.27	0.06	0.05
OM-2	0.	0.	0.	0.	0.	*Burn was not performed.*			
OM-3[♠]	0.	0.	0.	42.82	43.34	43.36	n/a	43.38	43.37
DRD	476.38	476.69	476.69	433.16	454.51	454.46	454.22	454.33	454.34
RTC-1[*]	0.	0.	0.	0.	n/a	0.60	0.43	0.48	0.48
RTC-2	0.	0.	0.	0.	0.66	1.82	1.82	1.71	1.71
RTC-3[*]	0.	0.	0.	0.	n/a	2.29	2.17	2.04	2.05
RPF	961.41	961.31	961.31	961.33	961.07	961.39	960.86	960.93	960.97
RTC-4	0.	0.	0.	0.	0.88	0.63	0.66	0.54	0.50
RTC-5	0.	0.	0.	0.	7.95	5.16	5.18	5.04	5.00
RTC-6	0.	0.	0.	0.	0.97	1.11	1.09	0.99	0.98

♣ Burn was not targeted; only targeted burns included in post-flight review. Fixed explicit Δv in FSW automation.
♢ OMSe checkout burn ($\Delta t \geq 30$ sec).
♡ Non-zero burn used in Copernicus during launch window computation to improve convergence ($\Delta t < 10$ sec).
♠ During flight, this was repurposed as a 100 sec 6+X DFTO burn. No onboard targeting; forced ground targets.
* Did not have time to re-optimize with updated navigation state vector before burn.

The design, targeting, and execution performance is shown in Table 1. Ultimately, with the exception of DRD, each executed burn was within 2 ft/s of its estimated Δv by Copernicus. In some cases, this was true with respect to the original 3-DOF TIP mission design many months before flight.

Despite the introduction of the DFTO at OM-3, even many of the downstream burns still agreed with the original 3-DOF mission design. The largest outlier was DRD, whose deviation from its dedicated re-designed baseline is attributed to the six +X Auxiliary engines (6+X) thrust being higher than expected during the DFTO, improved ground tracking prior to DRD, and slight differences in modeled Orion mass. The real-time re-optimization of DRD with the new navigation state independently converged to the same targeting solutions as the other real-time processes.

Once the trajectory was updated after launch for the assumed liftoff time, the re-optimized trajectory variations compared to the integer minute solutions were minimal. The strong agreement in solutions increasing in model fidelity of the columns from left to right to the post-flight estimates provides validation of the accuracy of Copernicus 3-DOF modeling, the targeting and execution process, and Orion's capability to fly an optimized trajectory. It is also noteworthy that the trajectory design and operations processes are robust to variations in timing assumptions as shown with the slight differences in the assumed and actual liftoff times from the various stages.

CONCLUSION

The trajectory design and operations contributed significantly to the success of Artemis I, whose primary objective relied on the sophisticated trajectory. They also facilitated autonomous vehicle operations in the event of permanent communications loss scenarios, which will be important for future deep-space crewed missions. The scale of modern-day human spaceflight trajectory operations in cis-lunar space is vast and complex, but it is what enabled Orion to safely and successfully complete the Artemis I mission and create the operational infrastructure necessary for the new era of ambitious crewed space exploration.

ACKNOWLEDGMENT

The authors wish to acknowledge the following people for their contributions to Artemis I real-time on-orbit trajectory operations: Badejo "Toks" Adebonojo, TARGO; Ethan Ariyoshi, TRAJ Support; Charlie Barrett, FDO; Amelia Batcha, TARGO; Brennan Blumenthal, TRAJ; Colin Brown, TARGO; Phil Burley, TRAJ; Tim Dawn, TARGO; Randy Eckman, TARGO; Josh Fogel, TRAJ Support; Shane Fuller, TRAJ; Josh Geiser, TARGO; Jeff Gutkowski, OPT office support; Robert Harpold, TARGO; Matt Horstman, OPT office support; Brian Killeen, TARGO; Jacob Madalon, FDO; Mark McPherson, FDO; Liam Mortell, TRAJ Support; Natasha Peake, FDO; Brian Perry, FDO; Adrianna Ramirez, TRAJ; Julie Read, FDO; Scott Roe, TRAJ; Sarah Smallwood, TARGO; Tim Stuit, FDO; Max Widner, TARGO; and Jacob Williams, OPT office support.

TERMINOLOGY

6+X	six +X Auxiliary engines
8+X	eight +X Auxiliary engines
Cop.	Copernicus
CR3BP	Circular Restricted Three-Body Problem
DAC	Design Analysis Cycle
DFTO	Developmental Flight Test Objective
DRD	DRO Departure
DRI	DRO Insertion
DRO	Distant Retrograde Orbit
DOF	Degrees of Freedom
DSN	Deep Space Network
Δv	Delta velocity
EGS	Exploration Ground Systems
EI	Entry Interface
EMPO	Exploration Mission Planning Office
EMPT	End-to-end Mission Performance Team
EMRP	Earth-Moon Rotating-Pulsating
ESDMD	Exploration Systems Development Mission Directorate
FCR	Flight Control Room
FCT	Flight Control Team
FDO	Flight Dynamics Officer
FOD	Flight Operations Directorate
FRAC	Flight Readiness Analysis Cycle
FSW	Flight Software
GR&A	Ground Rule and Assumption
GTLT	Ground Two-Level Targeter
ICPS	Interim Cryogenic Propulsion Stage
JSC	Johnson Space Center
KSC	Kennedy Space Center
LM	Lockheed Martin
LP	Launch Period
MCC	Mission Control Center
MER	Mission Evaluation Room
MMT	Mission Management Team
MPCV	Multi-Purpose Crew Vehicle
MSFC	Marshall Space Flight Center
MSL	Misson Selection Logic
OM	Orbit Maintenance burn
OMSe	Orbital Maneuvering System Engine
OPF	Outbound Powered Flyby
OPT	On-Orbit Performance Team
OrbGuid	Orbit Guidance flight software

OTC	Outbound Trajectory Correction
OTLT	Onboard Two-Level Targeter
POST2	Program to Optimize Simulated Trajectories II
PRM	Perigee Raise Maneuver
RCS	Reaction Control System
RPF	Return Powered Flyby
RTC	Return Trajectory Correction
SLS	Space Launch System
SMRCS	Service Module Reaction Control System
SOI	Sphere of Influence
TARGO	Trajectory Analysis, Retargeting, and Optimization Officer
TIG	Time of Ignition
TIP	Target Interface Point
TLI	Trans-Lunar Injection
TLT	Two-Level Targeter
TOR	TARGO Operations Report
TRAJ	Trajectory Officer
ULA	United Launch Alliance
US	United States
USS	Upper Stage Separation burn
UTC	Coordinated Universal Time

REFERENCES

[1] J. Williams, T. F. Dawn, and A. L. Batcha, "A History of Orion Mission Design, Copernicus Software Development, and the Artemis I Trajectory," AAS/AIAA Astrodynamics Specialist Conference, Aug. 2023. AAS 23-241.

[2] A. L. Batcha, J. Williams, T. F. Dawn, J. P. Gutkowski, M. V. Widner, S. L. Smallwood, B. J. Killeen, E. C. Williams, and R. E. Harpold, "Artemis I Trajectory Design and Optimization," AAS/AIAA Astrodynamics Specialist Conference, Aug. 2020. AAS 20-649.

[3] R. Odegard, J. L. Goodman, C. P. Barrett, K. Pohlkamp, and S. Robinson, "Orion Burn Management, Nominal and Response to Failures," 39th Annual AAS Guidance, Navigation and Control Conference, Apr. 2016. AAS 16-113.

[4] R. Odegard, D. P. Dannemiller, C. P. Barrett, and K. Pohlkamp, "Orion Burn Automation Response to Failures," 41st Annual AAS Guidance, Navigation and Control Conference, Feb. 2018. AAS 18-087.

[5] B. G. Marchand, K. C. Howell, and R. S. Wilson, "Improved Corrections Process for Constrained Trajectory Design in the n-Body Problem," *Journal of Spacecraft and Rockets*, Vol. 44, No. 4, 2007, pp. 884–897.

[6] R. Harpold, C. Brown, B. J. Killeen, R. A. Eckman, T. F. Dawn, B. Adebonojo, and J. Williams, "Artemis I Off-Nominal Trajectory Design and Optimization," AAS/AIAA Astrodynamics Specialist Conference, Aug. 2023. AAS 23-129.

[7] R. A. Lugo, J. D. Shidner, R. W. Powell, S. M. Marsh, J. A. Hoffman, D. K. Litton, and T. L. Schmitt, "Launch Vehicle Ascent Trajectory Simulation Using the Program to Optimize Simulated Trajectories II (POST2)," AAS/AIAA Space Flight Mechanics Meeting, Feb. 2017. AAS 17-274.

[8] J. Williams, "A New Architecture for Extending the Capabilities of the Copernicus Trajectory Optimization Program," *Advances in the Astronautical Sciences: Astrodynamics 2015*, Vol. 156, 2016. AAS 15-606.

[9] S. S. Harris, C. M. Spreen, M. T. Stoft, G. J. Frank, and R. J. Stewart, "Performance of the Two-Level Targeting Algorithm on Artemis I," 45th Annual AAS Guidance, Navigation and Control Conference, Feb. 2023. AAS 23-056.

[10] V. G. Szebehely, *Theory of Orbits: the Restricted Problem of Three Bodies*. New York: Academic Press, 1967.

[11] C. Ocampo, "An Architecture for a Generalized Trajectory Design and Optimization System," *Proceedings of the Conference: Libration Point Orbits and Applications* (G. Gómez, M. W. Lo, and J. J. Masdemont, eds.), World Scientific Publishing Company, June 2003, pp. 529–572. Aiguablava, Spain.

[12] G. Holt, C. D. D'Souza, and M. Wasinger, "An Overview of the Artemis I Navigation Performance," 45th Annual AAS Guidance, Navigation and Control Conference, Feb. 2023. AAS 23-057.

[13] J. Sullivan and C. D'Souza, "Extended Kalman Filter Performance on the Artemis-1 Mission," 45th Annual AAS Guidance, Navigation and Control Conference, Feb. 2023. AAS 23-053.