

Operations Systems Engineering for the Lunar Flashlight Mission

Michael Hauge
Georgia Institute of Technology
620 Cherry Street NW, Atlanta, GA 30332; (908) 612-6192
mhauge3@gatech.edu

ABSTRACT

Lunar Flashlight, a 6U CubeSat developed by NASA’s Jet Propulsion Laboratory (JPL) and operated by students at the Georgia Institute of Technology (GT), was launched in December 2022 with a mission to demonstrate novel small satellite technologies, including a first-of-its-kind green monopropellant system, and to map surface water ice in permanently shadowed regions of the lunar south pole using near-infrared laser reflectometry. While responsible for tactical operation of the spacecraft, including commanding and telemetry monitoring, GT students have also taken on strategic roles as operations systems engineers. The team has maintained, developed, and refined models of spacecraft subsystems as well as coordinated the project’s approach to anomaly response and fault protection. This paper reports how flight data and post-launch experiences have influenced the development and refinement of these models and approaches, and how in turn this systems engineering work has allowed the team to make more efficient use of the spacecraft’s capabilities, especially in dynamic anomalous situations, by taking advantage of margins, synthesizing data, and adapting flight rules and constraints. In-flight anomalies have required substantial rework of the mission’s concept of operations, and anomaly management and resolution has leaned heavily on modeling and predictions from the operations systems engineers. Working closely with JPL subject matter experts, the GT operations team has made full use of available data, including telemetry and observed system behavior, to swiftly recognize and address anomalies, support strenuous recovery efforts, and make possible a realignment of the concept of operations to achieve mission success despite significant challenges.

INTRODUCTION

The Lunar Flashlight (LF) mission was originally conceived in 2013 as a way to achieve both novel technology demonstration and lunar science in a CubeSat form factor.¹ While the initial idea included a solar sail, similar to the LightSail or later NEA Scout spacecraft, which eventually was replaced by a green monopropellant propulsion system, the overall concept has remained constant for the duration of the project: a 6U CubeSat with a 2U near-IR reflectometer instrument which uses advanced COTS components along with new small-sat technologies to both demonstrate such technologies and map the distribution of surface water ice in permanently shadowed regions (PSRs) of the lunar south pole.^{1,2} The Lunar Flashlight spacecraft was developed by the NASA Jet Propulsion Laboratory (JPL) under the Space Technology Mission Directorate (STMD), and its official classification is a “technology demonstration.”³

Integration and test (I&T) for the spacecraft was performed at the Georgia Institute of Technology (GT) and the Georgia Tech Research Institute. Detailed information on the integration and test cam-

paign can be found in Reference 3. In addition, the GT Space Systems Design Laboratory (SSDL) was contracted to serve as the primary Mission Operations Center/Ground Data System (MOC/GDS) for the mission, and a team of graduate and undergraduate students have supported the mission throughout Phase D and E.

Spacecraft Overview

An overview of the Lunar Flashlight spacecraft and its components is shown in Figure 1. The instrument payload is comprised of four near-IR lasers with their own battery, along with a detector and receiver to record reflections off the lunar surface; the payload also includes a radiator to achieve temperatures well below 0°C on the detector electronics for enhanced IR sensitivity. Four solar arrays provide power to the spacecraft; two deploy automatically when the spacecraft exits its dispenser, and two are deployed via burnwire actuation. The avionics stack is made up of a Sphinx radiation-hardened flight computer and an Iris Deep Space Transponder, both developed by JPL, as well as a custom electric power system (EPS) card and interface board. Two sets of X-band low gain antennas (LGAs), on the

$\pm Z$ faces of the spacecraft, allow for communication with Earth via the Deep Space Network (DSN) in any orientation.

The spacecraft is three-axis stabilized, and the attitude control system (ACS) is an XACT-50 made by Blue Canyon Technologies (BCT). The ACS includes three reaction wheels, and no magnetorquers, since the Earth’s magnetic field in cislunar space is negligible; it also includes four sun sensors, an inertial measurement unit, and a stellar reference unit (SRU). The propulsion system, which occupies around 2.5U, was designed and developed in partnership between the NASA Marshall Space Flight Center (MSFC) and the GT SSDL.⁴ The propulsion system uses ASCENT “green” monopropellant, which is significantly less toxic than the more commonly used hydrazine. Four thrusters, canted inwards from the +Z axis by 12°, provide Δv and reaction wheel desaturation capabilities. The propulsion system is commanded directly by the XACT.

Mission Objectives

As a technology demonstration mission from STMD, Lunar Flashlight’s objectives are twofold: perform measurements with its instrument payload, and demonstrate novel small satellite technologies. In particular, the mission seeks to demonstrate:

- The Lunar Flashlight Propulsion System (LFPS), an advanced miniaturized green monopropellant propulsion system, including additively manufactured titanium components and a COTS-based controller, designed to provide over 2500 N · s of total impulse.^{4,5}
- The Iris Deep Space Transponder radio in cislunar space, communicating with the DSN at various uplink and downlink data rates, and novel Pseudorandom-Noise (PN) Delta Differential One-Way Ranging (DDOR) capabilities.
- The instrument, including four high-power diode lasers and a multi-band reflectometer, designed to detect water ice absorption peaks.⁶
- The Sphinx radiation-hardened processor, which serves as the flight computer.
- The spacecraft’s flight software (FSW) architecture created in JPL’s F Prime framework, which utilizes a significant amount of autocoding and was created concurrently with FSW for the NEA Scout mission.²

The mission’s two Level 1 requirements are:

L1-01: Lunar Flashlight shall have the capability to address a key strategic knowledge gap at the moon.

Full Success: Detect and map surface water ice on the moon with a spatial resolution of 1 km over 10% of the permanently shadowed and occasionally sunlit regions poleward of 80°S latitude.

Minimum Success: Demonstrate the ability to detect surface water ice content with a spatial resolution of 10 km or better with multiple measurements in permanently shadowed and occasionally sunlit regions poleward of 80°S latitude.

L1-02: Lunar Flashlight shall be in a 6U CubeSat form factor compatible with a NASA provided CubeSat deployer for launch on a NASA provided launch vehicle.

Concept of Operations

Lunar Flashlight was originally envisioned as a secondary payload for the first test launch of NASA’s Space Launch System.¹ While this mission, later designated Artemis I, would launch with ten CubeSats, delays in integration caused Lunar Flashlight to be demanifested. It eventually launched in December 2022, one month after Artemis I, on a SpaceX Falcon 9 rocket as a rideshare payload alongside the Japanese private company ispace’s Hakuto-R lunar lander.

Figure 2 shows Lunar Flashlight’s nominal concept of operations. Shortly after launch, the spacecraft would deploy and detumble into a sun-pointed attitude with its reaction wheels; after commissioning the propulsion system and performing a desaturation burn, a series of Trajectory Correction Maneuvers (TCMs) would be performed over several months. A Lunar Orbital Insertion (LOI) maneuver into a Near Rectilinear Halo Orbit (NRHO) would be performed over the lunar south pole, beginning a series of highly elliptical orbits with perilune altitudes of only 10-20 km. Each orbit would require three Orbit Trim Maneuvers (OTMs) to maintain the period and target perilune altitude. Ten orbits, each about six days long, would result in ten observation opportunities of PSRs at the lunar south pole, followed by a deorbit maneuver that would dispose the spacecraft on the lunar surface.

Unfortunately, only a few days into the Launch and Early Operations Phase (LEOP), significant issues with the propulsion system’s performance were

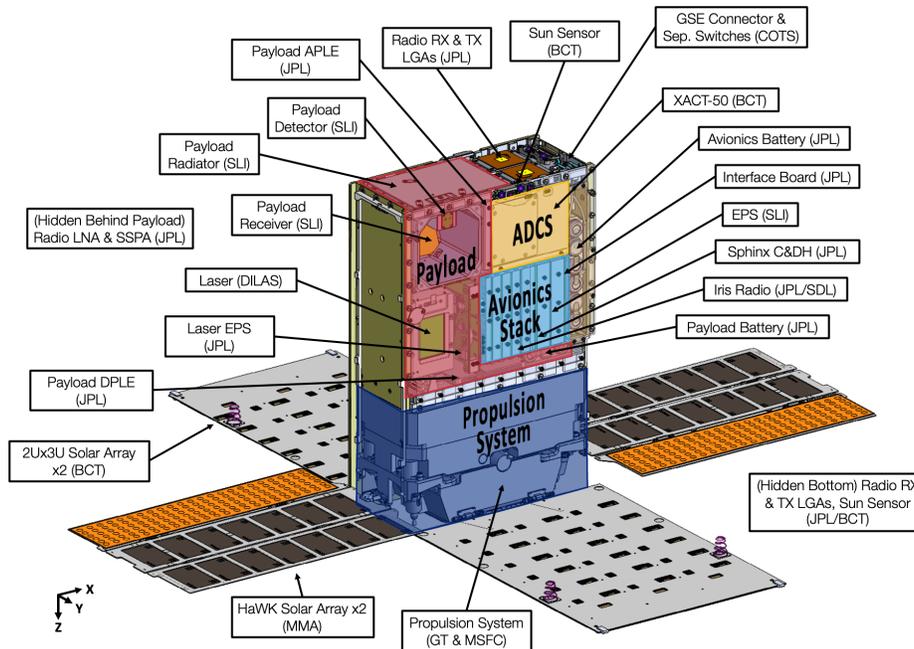


Figure 1: Lunar Flashlight spacecraft overview and body frame axes.³

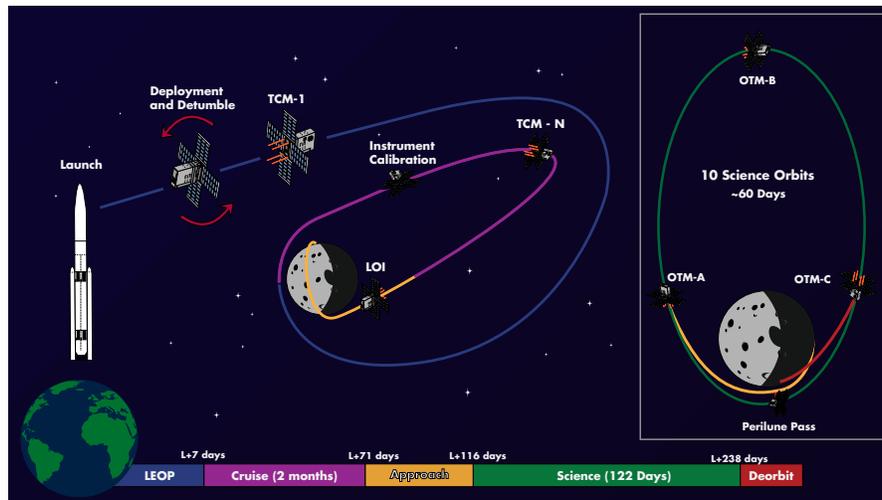


Figure 2: Concept of operations for the Lunar Flashlight mission.

discovered, which began a series of characterization and recovery efforts which resulted in multiple trajectory redesigns, seismic changes to maneuver schemes, and an eventual change of mission concept to target a series of lunar flybys rather than insertion into an NRHO. It still remains a possibility that propulsion system performance will degrade and resist recovery efforts in such a way that a lunar encounter is impossible, leading to a realignment of mission objectives entirely towards technology demonstration and eventual escape of the Earth-moon system.

Mission Operations and Operations Systems Engineering

In the summer of 2021, JPL awarded the GT SSDL the contract to operate Lunar Flashlight. A team of student operators, starting with graduate students and eventually including undergraduate students as well, were trained by JPL engineers and subject matter experts (SMEs), established the GT MOC and GDS, participated in spacecraft I&T, and prepared for launch with rigor. A brief timeline of GT mission operations involvement is given below:

July 2021: Contract for operations awarded

August - December 2021: MOC facility buildup

October 2021: First benchmark mission operations documents delivered

November 2021 - March 2022: Mission operations team supports spacecraft I&T

January - March 2022: GT mission operations team leads Day-in-the-Life (DITL) and fault protection testing

March - November 2022: Operational Readiness Tests (ORTs); procedure and tool development

August - December 2022: Training of new operations personnel

October 2022: Final mission operations documentation released; passed Operational Readiness Review (ORR)

December 2022 - May 2023: Active operation of the Lunar Flashlight spacecraft

While the GT mission operations team started with only four students, it has grown to seven graduate and seven undergraduate students, most of whom joined the project at the beginning of the academic year in August 2022. A GT Mission Operations Lead coordinates the activities of the operations team with project management at JPL and other teams, including the propulsion team (which includes representatives from JPL, MSFC, and various vendors) and the science team. JPL provides oversight, expertise, and mission assurance in the form of project management, a Mission System Manager (MSM), SMEs for all spacecraft subsystems, a Mission Operations Assurance Manager (MOAM), and the project’s Mission Design and Navigation (MDNAV).

Since the launch in December 2022, GT mission operators have staffed consoles during DSN tracks, usually a few hours long, one to three times a day, commanding and monitoring the spacecraft in real-time “tactical” operations. Additionally, new activities, including command sequences and procedures, have been developed in “strategic” time between contacts to exercise the spacecraft’s capabilities and respond to the dynamic anomalous environment. Every new activity is subject to verification and validation on the LF testbed, which is housed in the GT MOC, and reviewed by SMEs and project management.

Finding a balance between strategic and tactical operations for a small project in which the mission operators have distributed responsibilities has been an evolving challenge. With JPL SMEs often not closely involved in day-to-day operations or the development of new activities, operators took on roles as “subsystem leads for operations” as part of training and preparation; however, these same operators were often relied on for their tactical experience as well. Reference 7 describes the processes and approach that have allowed the mission operations team to overcome many of these challenges.

Subsystem Models

As operations systems engineers, the GT operators are the custodians of subsystem models, maintaining, refining, and sometimes developing them from scratch to make predictions about spacecraft behavior and assess performance. Before the launch, these models were used to inform the development of activities such as the deployment/boot-up sequence, initial acquisition, propulsive maneuvers, and instrument operations.

For power modeling, the GT team has made use of JPL’s Multi-Mission Power Analysis Tool (MMPAT), configured based on test data of the LF EPS. With a power equipment list (PEL) of the spacecraft, provided by JPL, which laid out the expected power draw of each subsystem in each of its modes, power profiles for critical events like the deployment and initial boot-up could be constructed, and were fed into MMPAT for simulation with wrapper scripts developed at GT. For example, Figure 3 shows the results of simulating the deployment and initial boot-up of the spacecraft, as it detumbles into a sun-pointed rotisserie and runs through its initial Safe Mode Sequence (SMS).

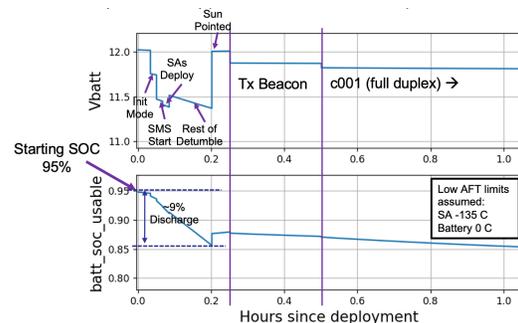


Figure 3: Power profile simulation of deployment and initial boot-up.

For the ACS, two simulation environments were used. First, the engineering model XACT on the LF testbed (shown in Figure 4) was connected to

a BCT Realtime Dynamics Processor (RDP), which could simulate sun sensor and SRU inputs to the XACT based on propagated orbital and attitude dynamics. This setup allowed the operations team to test slews, maneuvers, and command sequences related to the ACS in the testbed with realistic performance in terms of timing, stability, and reaction wheel speeds. In addition to the RDP, which simulates in real time, a BCT software simulation was used for faster iteration, which was essential for rapid activity development. While the software simulation did not replicate the commanding interface between FSW and the XACT, setting up precise initial conditions was feasible in a way that it was not for the RDP and XACT, making the two simulation environments complementary.

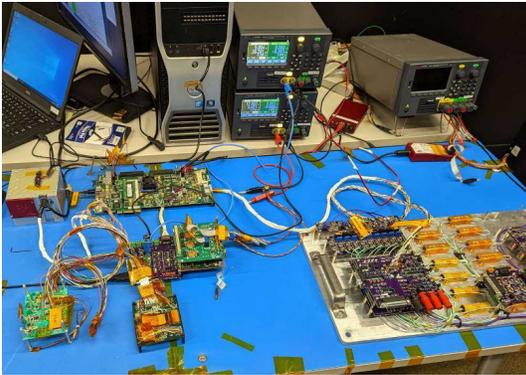


Figure 4: LF testbed in the GT MOC.

Because neither simulation environment included calculations for solar radiation pressure (SRP), angular momentum buildup predictions were made using a tool developed by the operations team in MATLAB. Figure 5 shows the predicted momentum buildup with worst-case SRP and gravity gradient torques with deployment tipoff rates of $10^\circ/\text{s}/\text{axis}$, indicating that under these assumptions it would take more than 21 days to exceed the overall system momentum limit of 66 mNms.

The spacecraft thermal model was created and maintained at JPL, with the operations team able to request simulation results based on power dissipation profiles and use this information to inform constraints. Prior to launch, most thermal concerns were centered on the Iris, and particularly its solid-state power amplifier (SSPA), which dissipates significant heat while transmitting. Figure 6 shows the thermal model results for the initial SMS configuration of the radio, toggling between receive-only, transmit-only, and full duplex modes. This simulation result was used to demonstrate that the chosen configuration would not result in radio components exceeding allowable flight temperatures.

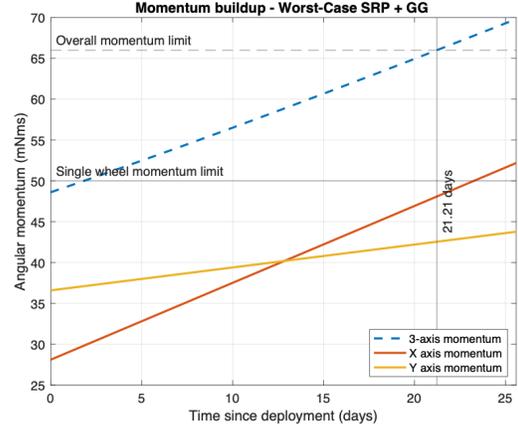


Figure 5: Pre-launch worst-case momentum buildup prediction.

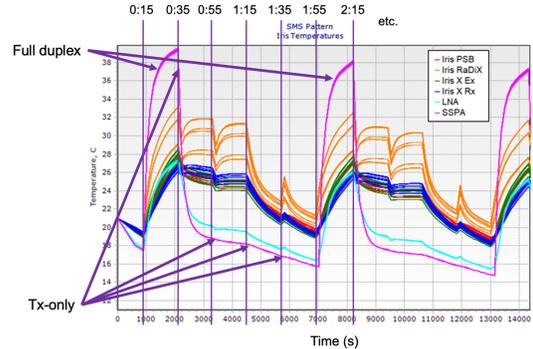


Figure 6: Thermal model simulation results of initial SMS radio configuration.

Modeling the performance of the propulsion system, in terms of estimating Δv from spacecraft telemetry, was done pre-launch using the rocket equation. Using a predicted mass flow rate provided by the propulsion team, telemetry indicating how long each thruster valve had been commanded open over the course of a burn could be used to estimate a mass delta, which in turn was used to estimate Δv based on the specific impulse of the propellant.

All of these subsystem models have been refined since launch based on flight data and spacecraft behavior, from small improvements to thermal and ACS simulations to wholesale restructuring of performance modeling of the propulsion system. In turn, the improved models have been used to inform the development of new activities that, prior to the launch, would have been thought impossible or outside the scope of the originally planned operations.

Flight Rules and Constraints

The operations team has also been responsible for collecting and tracking operational constraints on the flight system. Subsystem constraints based on modeling and predicted performance, such as minimum gaps between DSN tracks, minimum times between propulsion system firings, and minimum battery voltages required before power-intensive activities, were compiled and added to design- and heritage-based constraints in the project's flight rules. Over 100 draft flight rules were provided by JPL, and a large part of pre-launch systems engineering work for the operations team was the clarification and enforcement of these rules. Following edits, additions, and proposed enforcement schemes by the operations team, each rule was reviewed by SMEs and project management, and finally either accepted or rejected. A waiver process was also developed, establishing the criteria and approvals required to allow for the violation of a flight rule in specific or general circumstances, categorized by subsystem driver and criticality. DITL testing during I&T and ORTs during operations team preparation were vital for the identification and refinement of these constraints, as well as vetting enforcement methods.

Operationally, the team has also identified and mitigated constraints on commanding and telemetry monitoring. LF FSW does not support branching or conditional logic in command sequences, so careful ordering and timing is necessary to ensure desired behavior. In addition, observability of certain important data points is limited by either the lack or necessary throttling of telemetry channels. While these kinds of constraints were not given requisite focus prior to launch, they have been of vital importance in flight. In particular, the lack of representative readings on many telemetry channels on the LF testbed, along with other testbed limitations such as a low-data-rate serial connection, lack of engineering model hardware for all subsystems (including battery and solar power emulation), and slightly different Sphinx firmware (resulting in different pinouts than the flight unit) required the operations team to document these "testbed-isms" and develop mitigations for how they could affect preparation, training, and verification and validation once in flight.

Fault Protection Configuration

LF has a detailed catalog of faults originating from subsystems or from the FSW telemetry monitoring system, called GenMon. For LF, GenMon is configured to monitor currents and voltages from the EPS and Sphinx, temperatures from around

the spacecraft, and many of the payload instrument telemetry channels. GenMon compares the value of a telemetry channel against upper and lower limits, and triggers the fault if the value is outside the defined limits for a defined persistence period. Each fault has an associated ID, and many faults have associated responses (essentially subroutines of commands) intended to isolate, and in some cases correct, the fault. While many fault monitor thresholds are configurable, the responses are not; the modification or addition of responses requires a FSW update, unlike other F Prime fault protection implementations.⁸ Each individual fault ID can be enabled (i.e., if that fault condition is present, FSW will execute the associated fault response) or disabled.

FSW-defined faults fall into the following categories:

- ACS faults (from the XACT)
- Iris faults (from the Iris radio)
- Propulsion system faults (from the LFPS)
- Commanding faults:
 - Command failed in sequence fault
 - Command loss timer expired fault
- GenMon faults (from FSW monitoring of telemetry channels)

The EPS also has under-voltage lockout protection, but this is implemented in hardware and is not a FSW-defined fault that can be toggled.

During I&T, the operations team tested every defined fault response. For those for which it was possible, the fault condition was created on the testbed or flight unit (e.g., setting GenMon temperature limits outside ambient temperatures), and for others, the fault IDs were triggered via commands. Prior to launch, the fault IDs to be enabled by phase and activity were defined, and a script was developed to manage the commanding of enabling and disabling fault IDs.⁹

The response to many FSW-defined faults is to transition to Safe mode. LF does not have a state machine, so there are no fundamental differences between its Normal mode and Safe mode; however, upon a transition from Normal mode to Safe mode, any currently executing command sequence is stopped and the Safe Mode Sequence (SMS) is loaded from a protected file system location and executed. The SMS, which puts the spacecraft into a low-power, sun-pointing state and configures the radio into a predictable pattern of modes, swapping

between the two LGA pairs to allow for commanding and recovery, was one of the first sequences to be developed by the operations team. For the launch, three SMSs were prepared: the initial SMS, which would run on initial boot-up; a standard SMS; and a propulsion-specific SMS, which would be used when the propulsion system was active and which included extra steps to configure the LFPS to a safe state.

Anomaly Response

The logistics for documenting and responding to spacecraft and MOC anomalies were laid out prior to the launch in a project-level Anomaly Response Plan (ARP) written by the GT operations team. Upon the occurrence of an anomaly, its criticality would be assessed, along with the required speed of response: immediate (within the track it occurred), imminent (within 24 hours of occurrence), or delayed (more than 24 hours after occurrence). Operators would be responsible for documenting the anomaly, including information on the failure itself, the timeline of events leading up to and immediately after it, and the spacecraft state following the anomaly.⁹

The ARP was refined through the operations team’s experience in the designated “off-nominal” ORT, which was held in late October 2022. The period from launch to the execution of the first TCM was rehearsed over the course of several days, with simulated spacecraft and MOC anomalies, from drastically increased initial tipoff rates to operator illnesses to GDS failures, injected to test the team’s ability to respond and recover.

SUBSYSTEM MODELING

Since the launch in December 2022, the operations team’s understanding of LF’s behavior, capabilities, and limitations has evolved considerably. In addition to improvements in operations processes, strategic systems engineering work in using flight data to inform subsystem modeling and adapt flight system constraints accordingly has enabled the team to respond to significant anomalies and introduce new modes of operation that had not been considered prior to the launch.

Power Subsystem

Significant refinement of the spacecraft power model has been possible with measurements of loads derived from telemetry. The pre-launch spacecraft PEL, which collected predicted power draw information from each subsystem, had a number of conservative assumptions: many subsystems had applied

subsystem-level margins to estimated power draw, and an overall 10% system-level margin had been used. In addition, input profiles were given to MM-PAT as constant-power loads, with current draw assumed to increase as battery voltage decreases.

Voltage and current telemetry to and from the solar arrays and battery allowed the operations team to compute the actual load on the system as the sum of the loads on the regulated 5V bus and the unregulated 12V bus, and compare this number to the prediction in the PEL. The first opportunity for this analysis was in the first few days of the mission, regarding the predicted power draw of the heaters for the four thruster catalyst beds of the propulsion system.

ASCENT green monopropellant, unlike hydrazine, cannot cold-fire through catalyst beds at ambient temperatures; the catalyst must be above around 400°C in order for the propellant and catalyst to react fully. Therefore, each catalyst bed on the LFPS has its own heater. As part of the design of the LFPS, the total power draw of the propulsion system was required to be less than 47 W; it was estimated that having all four heaters on at the same time would violate this limit, so a derating scheme was implemented which would cycle the heaters at 5 Hz to be able to heat all four catalyst beds at the same time but at effective duty cycles of 25%, 50%, or 75%.¹⁰ Since the heaters draw more power at lower temperatures, this parameter (the “thruster heater derate max,” or THDM) was planned to first be set to 2, and then to 3 once the catalyst beds had all reached 300°C, as shown in Figure 7.

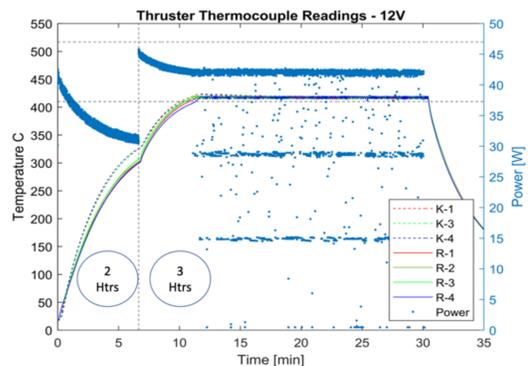


Figure 7: Catalyst bed heating curves from LFPS I&T.¹⁰

In flight, after the propulsion system was primed by venting the pad pressure nitrogen gas from the fuel lines, the first planned firing was a commissioning burn: a series of short pulses on all four thrusters simultaneously. When this was attempted, it was found that, with a THDM of 3, the catalyst beds

could only be heated to around 415°C, rather than the planned 440°C, due to the voltage sag on the unregulated bus due to the power draw from the heaters. Analysis of the power profile over the course of the catalyst bed heating and firing from commissioning, however, showed that the load on the EPS was significantly lower in flight than predicted by the PEL, as illustrated in Table 1. Overall, loads in flight were between 20-30% lower than predicted in the PEL, which reflects the 15% subsystem-level margin applied to LFPS loads measured in I&T and the 10% overall system-level margin in the PEL.

Table 1: Commissioning Power Profile

State	PEL Load	Flight Load	% Diff.
Idle	46.1 W	34.1 W	26.0
THDM=2	90.6 W	67.3 W	25.7
THDM=3	90.6 W	71.4 W	21.2
Firing	98.9 W	70.9 W	28.3

Based on data from commissioning (see Figure 8), the project determined that it was feasible to use a THDM of 4 to achieve temperatures of 440°C on all four thrusters. For the initial desaturation burn, which was attempted in the very next contact, all four thruster heaters were enabled when the catalyst beds had reached 390°C, and thermostatic control was established between 440-450°C. The maximum load during firing was computed to be 79.5 W, 30.1% lower than predicted by the PEL (see Figure 9). This analysis, which allowed the operations team to take advantage of margin in the EPS design and respond to the slight underperformance of the heaters, was performed strategically between two DSN tracks that were only 12 hours apart.

Moving forward, the measured loads in different spacecraft states were used in MMPAT simulations of proposed activities, particularly those involving the propulsion system. When it became clear that the pre-launch concept of operations for Δv maneuvers would not be feasible given thruster underperformance, exploration of a maneuver scheme design space began, including one-, two-, and four-thruster maneuvers with the corresponding catalyst bed heaters staying on throughout the burn (as opposed to nominally turning off when the thrusters would begin to self-heat). Table 2 shows the results of preliminary power analyses identifying the maximum allowable fully off sun-pointing maneuver duration with different numbers of heaters and different Iris modes. The maximum allowable duration was limited by two constraints: keeping the battery state of charge (SOC) above 30% and battery voltage above 9.5 V.

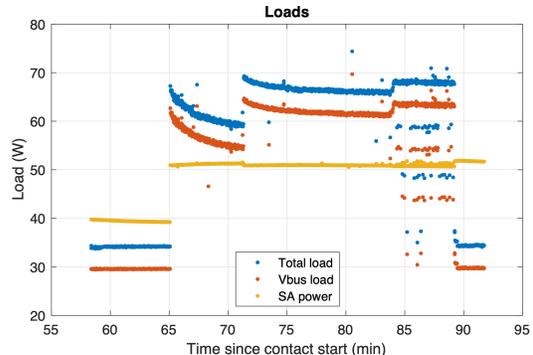


Figure 8: Commissioning power profile loads.

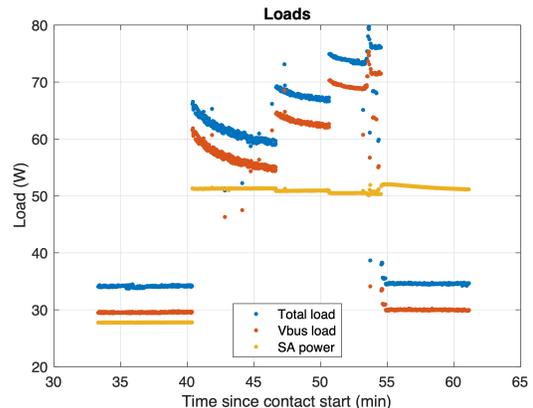


Figure 9: Initial desat attempt power profile.

Table 2: Off-Sun Maximum Maneuver Duration Predictions.

Number of Catbed Heaters	1	2	4
Iris TX/RX	590 s	470 s	140 s
Iris RX	1400 s	1100 s	620 s
Iris Off	2000 s	1520 s	920 s

The pre-launch estimated maximum maneuver duration of 1200s (20 minutes), which was constrained by the thermal limitations of the propulsion system, included turning the Iris radio off completely, a significant risk which was deemed necessary by power and thermal considerations at the time. The refinement of the MMPAT model based on flight data allowed the operations team to identify alternate, less-risky maneuver schemes, keeping the radio on in full duplex, with around the same maximum maneuver duration even with the significant additional load from the catalyst bed heaters.

The first sets of maneuvers, which were developed as single-thruster burns performed while rotating about the thrust force direction to achieve zero

net angular momentum buildup (described in detail later in this paper), were performed at partially sun-pointing attitudes, so power considerations were less limiting. When a new burn arc two months after launch required fully off-sun attitudes, an iterative approach involving MMPAT simulation and incremental changes to maneuver duration was adopted. From the analysis underlying the results shown in Table 2, the preliminary estimate for maximum off-sun maneuver duration was around 10 minutes; however, the first test maneuvers, with durations of one and five minutes, resulted in much less severe voltage drops than anticipated.

Further testing of 10-minute maneuvers showed a large margin to the 30% SOC and 9.5 V constraints. The maximum maneuver duration acceptable under propulsion system thermal constraints was 20 minutes, which was attempted and executed without issue. Extrapolation of EPS data from this maneuver campaign indicated that the 9.5 V limit would be reached only after an approximately 35-minute long maneuver, as shown in Figure 10.

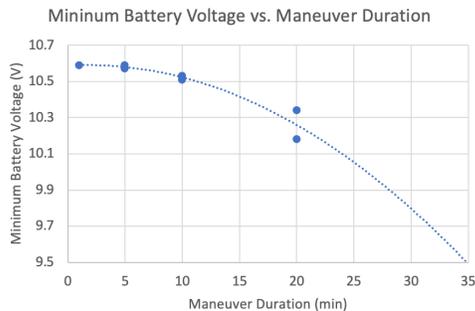


Figure 10: Minimum battery voltage reached during off-sun maneuvers.

The transition from constraints based on simulation, even with updated flight loads, to confidence-building based on incremental testing was motivated by mismatches between the predictions made about off-sun behavior under large loads by MMPAT with what was observed on the spacecraft. Further investigation has identified an overly conservative equivalent series resistance in the MMPAT model of the spacecraft battery, which is in the process of being reconfigured.

Attitude Control Subsystem

Over the course of the mission, the operations team’s usage and interpretation of ACS data has matured dramatically. Prior to launch, the XACT’s control algorithms were used to perform Δv and reaction wheel desaturation burns in tests and re-

hearsals, with very simple commanding and verification. Flight rules dictated that when the spacecraft’s angular momentum magnitude was predicted to exceed 25 mNms, a desat burn should be commanded, with a deadband of 12 mNms. However, in the first attempted desat, due to thruster issues, the spacecraft momentum magnitude increased rather than decreased, leading to a wholesale rethinking of ACS telemetry interpretation and strategic planning.

The XACT reports reaction wheel speeds, body rates, and a computed system angular momentum vector in telemetry, but the quantization on the momentum vector components is significant, with each vector component only reported to the nearest 0.2 mNms. As a result, the first step was to correlate reaction wheel speeds and the computed momentum, and then make predictions about what wheel speeds corresponded to dangerously high momentum states, and vice versa. Figure 11 shows the result of that correlation. The nominal capacity of each reaction wheel, 50 mNms, was found to correspond to around 6700 rpm, and the software-enforced overall momentum magnitude limit, 66 mNms, would correspond to around 8900 rpm on a single wheel. With additional information from BCT that the wheels were safe to run at even upwards of 9000 rpm for extended periods of time, this allowed the operations team to revise the threshold for requiring a desat burn higher, thereby decreasing the frequency of burns. This was beneficial to the safety of the mission as a whole, as each firing of the propulsion system in its anomalous state represented a potential danger.

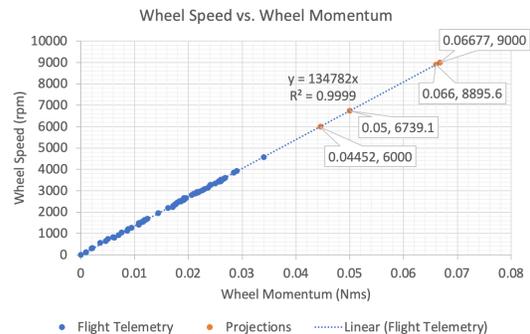


Figure 11: Correlation of reaction wheel speed and computed momentum telemetry.

Using wheel speed and momentum telemetry, following the first contact with the spacecraft, the operations team was also able to estimate the deployment tipoff rates that resulted in its initial momentum state, at approximately $6.52^\circ/\text{s}/\text{axis}$.

Approximately one week into the mission, it became clear from angular momentum buildup that one of the solar arrays (on the spacecraft $-X$ side)

had not deployed following the launch (as discussed further below). As a result, the SRP torque on the spacecraft was between one and two orders of magnitude larger than expected prior to launch. This meant that even though the desat threshold had been raised the spacecraft would reach it much more quickly than anticipated (on the order of weeks, rather than months). In addition, asymmetrical testing of the propulsion system resulted in significant angular momentum buildup.

To address these issues, as well as the inability to rely on the XACT’s desat algorithm (which assumed consistent thruster performance), two desat options were developed. The first used pulses from a single thruster to reduce the angular momentum magnitude, computing an attitude that aligned the angular momentum vector in the body frame as close as possible to opposite the torque produced by firing a particular thruster while respecting constraints on off-sun and off-Earth angles. This approach has been used successfully numerous times to reduce spacecraft angular momentum. The other approach was to use SRP as the external desaturating torque. A model of the spacecraft with one solar array undeployed, assuming certain optical characteristics, was developed and correlated with changes in momentum and wheel speed telemetry to find the direction and magnitude of SRP torque in the body frame; this could then be aligned opposite the current momentum vector in much the same way as the single-thruster desat. This scheme has also been successfully demonstrated (see Figure 12), with spacecraft momentum magnitude being decreased from 39 mNms to below 4 mNms over the course of several days. Both desat approaches have been used to achieve momentum deadbands much smaller than the 12 mNms baselined for the BCT desat algorithm prior to launch.

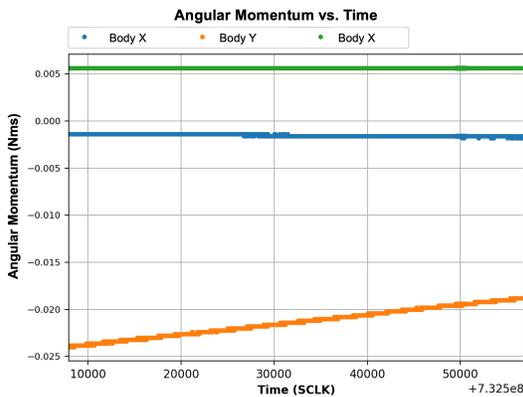


Figure 12: System angular momentum during SRP desat.

The most challenging ACS analysis and modeling tasks came as part of the effort to develop alternate maneuver concepts using fewer than four thrusters. The torques due to any three of the four thrusters are linearly independent, meaning that all four have to be functioning in predictable ways in order for a mapping of duty cycles to exist across them to achieve Δv with no net angular momentum buildup. In the absence of predictable performance, creative solutions were required.

First, when two of the thrusters still showed nominal performance, analysis was conducted to estimate how long they could both fire before saturating the reaction wheels. The results, shown in Figure 13, were not promising: even starting from an ideal momentum state (magnitude 30 mNms, in the exact opposite direction of the net torque), the two thrusters could only fire at nominal performance for around 40 seconds before reaching the Z-axis wheel’s limit of 50 mNms. As a result, a two-thruster burn would have to involve many slews between the desired Δv attitude and a desat attitude, a complex and inefficient mode of operation.

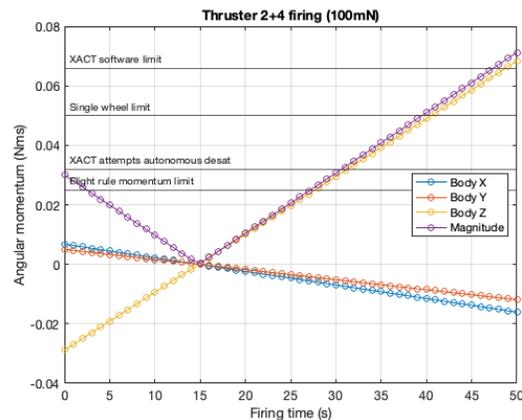


Figure 13: Angular momentum change from firing two thrusters simultaneously.

Eventually, the project settled on an approach, proposed by JPL SMEs, for rotating maneuvers performed using a single thruster. By rotating about the thrust force vector, which is orthogonal to the torque produced by that thrust force (since $\tau = \mathbf{r} \times \mathbf{F}$), the net angular momentum buildup over the course of a firing (for an integer number of rotations) could be zero. If the thrust magnitude and rotation rate were matched appropriately, then the angular momentum vector coordinated in the body frame will remain constant, as the external torque due to the thruster firing rotates it in the inertial frame exactly opposite the rotation of the body frame itself.

It was the responsibility of the operations team to implement this scheme in command sequences and verify the predicted evolution of the spacecraft’s momentum state. Both ACS simulation environments were critical to this task, as iterative changes to timing and command parameters were checked in the BCT software simulation, while command sequences were run on the LF testbed with the RDP simulation active.

Over the course of the development of the rotating maneuver scheme, problems with off-axis momentum buildup and attitude errors were encountered. With the help of BCT, updates to controller parameters were proposed and tested in software simulations. By increasing the gyroscopic torque compensation limit and the integral error limit of the controller, ACS performance while rotating at up to 6°/s improved dramatically, at the cost of slightly longer settling times when slewing between attitudes. Further improvement was seen by the injection of feed-forward torque into the control loop.

While these updates were straightforward to test and verify in software simulations, their implementation on the LF testbed and spacecraft was more complex. The table update and feed-forward torque commands needed had not been implemented in LF FSW. For most XACT commands, LF FSW managed the mapping between a command stem with arguments and a set of bytes to send over a serial bus to the XACT itself, as shown on the left of Figure 14. However, for the commands in question, the operations team had to develop tools to turn BCT-provided command strings (originally intended to be sent directly to an XACT through COSMOS, an open source command and control software for embedded systems) into raw bytes. These binary files would be uplinked to the spacecraft file system, and then commands sent to interpret the contents of those files as bytes to be sent to the XACT, as shown on the right of Figure 14.

During a rotating propulsive maneuver, any misalignment of the rotation axis and the thrust force vector would result in off-axis momentum buildup, requiring periodic desaturation burns. As rotating maneuvers began to be performed successfully for longer durations, the operations team could track this off-axis buildup and from it derive the rotation axis offset. The difference between the true thrust force direction and the nominal direction (as derived from CAD measurements) for the thruster used for the majority of the rotating maneuvers was found to be 6 mrad ($\approx 0.3^\circ$), within assembly tolerances but still enough to affect momentum buildup. With

the correct thrust vector now computed, the commanded rotation axis was corrected and the rate of momentum buildup was significantly decreased.

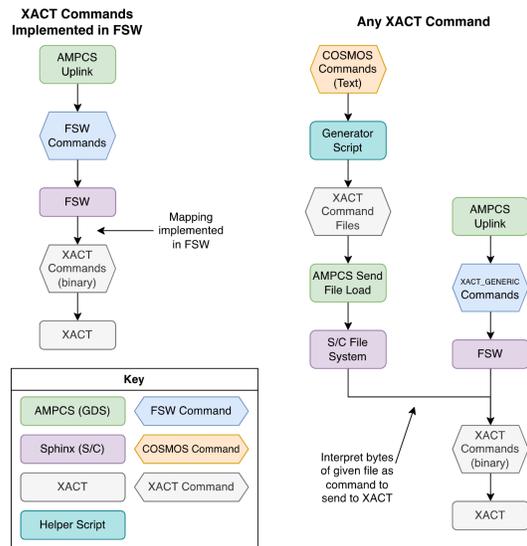


Figure 14: Commanding the XACT.

Thermal Subsystem

The operations team’s evaluation and trending of spacecraft thermal behavior over the course of the mission has led to significant changes to constraints. As mentioned above, temperatures for the Iris radio were the driver for constraints on track length (maximum two hours) and time between tracks (minimum six hours) prior to launch. In flight, however, the thermal environment was more benign than originally anticipated. Compiling Iris telemetry while the radio was in full duplex (TX/RX) across multiple contacts, illustrated in Figure 15, led the operations team to suspect that, since the temperature rise was much less than anticipated, longer and more frequent tracks should be possible.

The allowable flight temperature (AFT) of Iris components is nominally 50°C. Analysis of temperatures was at times hindered by poor resolution on many temperature telemetry points (for example, Iris SSPA temperature is reported to the nearest 3.05°C), but this was overcome by taking weighted moving averages and fitting to those data sets. Providing temperature data to JPL to refine the thermal model resulted in both longer and more frequent tracks, with the maximum track time extended to two and a half hours and the required time in between tracks reduced to only one hour. These changes were communicated to the DSN for scheduling and allowed for much greater flexibility.

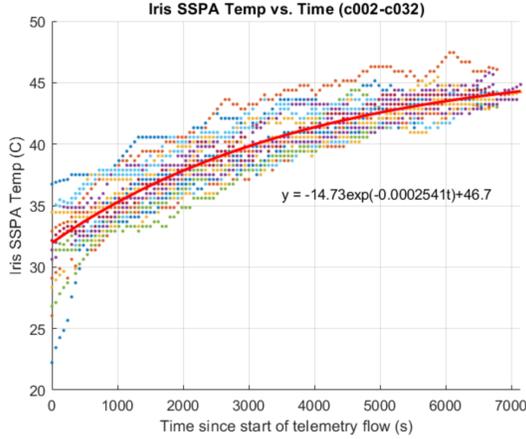


Figure 15: Iris SSPA temperatures across multiple contacts.

Just as with the EPS, the operations team was also able to find and take advantage of margins on temperatures. In several instances, discussion with SMEs lead to revisions of AFTs upward, as the original values were overly conservative and had been constraining desired operations. Even before the launch, the payload battery upper temperature limit had been raised to the hardware limit of 40°C as part of analysis done in preparation for thermal vacuum testing. Soon after launch, with the Iris SSPA temperature steady and close to 50°C by the end of most tracks, that component’s flight limit was revised to 60°C. Similarly, after initial high-power-draw propulsion activities led to higher than expected temperatures on the spacecraft EPS (due to increased power dissipation from diodes in the path between the solar arrays and the unregulated bus), detailed analysis of datasheets for all circuit board components on the EPS card allowed the limit to be raised from 50°C to 75°C. On the LFPS, similar datasheet analysis and conversations with vendors allowed the controller and pump temperature limits to be revised upwards as well.

Many of these changes were driven by the desire to push the envelope on new spacecraft activities. For example, the extra power draw that caused heating on the spacecraft EPS was part of the effort to raise the catalyst bed THDM to 4, as described above; LFPS temperature limits were the driver behind time limits on propulsion activities, so revising them higher allowed for more flexibility in testing and maneuvers. All the changes became important, however, when a command file error (CFE) resulted in the Iris staying in full duplex, its highest-power mode, for almost 65 hours straight over a weekend of unattended autonomous operation.

Figures 16 and 17 show spacecraft temperature telemetry over the course of the first 40 hours of the thermal anomaly. Around 20 hours in, the spacecraft has reached an approximately steady state. This ‘hot’ case, while entirely anomalous, demonstrated the thermal design of the spacecraft well: in this configuration, none of the steady-state temperatures reached by spacecraft components exceeded their flight limits, as illustrated in Table 3.

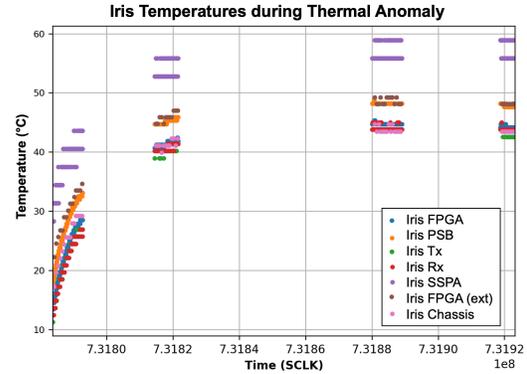


Figure 16: Iris temperatures during weekend ‘hot’ case.

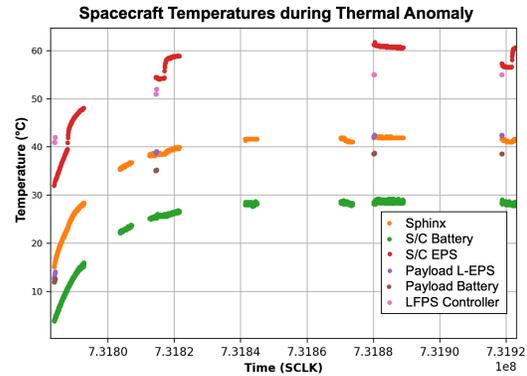


Figure 17: Additional spacecraft temperatures during weekend ‘hot’ case.

Recovery from the Iris thermal anomaly was simple: the Iris was set to receive-only mode for 24 hours, and spacecraft components returned to their nominal temperature ranges. Through the work done by the operations team earlier on in the mission, finding margin in various component temperature limits, the criticality of the anomaly was accurately assessed as not mission-threatening.

Table 3: Thermal Anomaly Steady-State (SS) Temperature Summary.

Component	AFT (°C)	Flight Limit (°C)	SS Hot Temp (°C)
Iris FPGA	50	50	44.2
Iris PSB	50	50	48.2
Iris Tx	50	50	42.5
Iris Rx	50	50	43.8
Iris SSPA	50	60	55.8
Iris FPGA (ext)	50	50	48.1
Iris Chassis	50	50	43.4
Sphinx	50	50	41.2
S/C EPS	50	75	56.5
S/C Battery	30	30	28.3
Payload L-EPS	50	50	42.1
Payload Battery	20	40	38.4
LFPS Controller	70	85	55.0
LFPS Pump	50	70	50.3
LFPS Tank	45	45	44.0

Propulsion System

The pre-launch model of propulsion system performance, in terms of predicting required firing times and estimating Δv , was based on the estimated nominal thrust force T and mass flow rate \dot{m} for each thruster. These estimates were provided by the propulsion team, and were planned to be refined over the course of the mission. Using these parameters, as well as the thrust force cant angle with respect to the spacecraft body Z-axis, θ , the ‘effective’ specific impulse in the direction of spacecraft motion when all four thrusters were firing together, $I_{sp,eff}$, could be computed:

$$I_{sp,eff} = \frac{T}{\dot{m}g_0} \cos \theta \quad (1)$$

From this value, the number of ‘thruster seconds’ (total seconds of on time across the four thrusters), $\Delta t'$, required for a given desired Δv could be computed from the rocket equation:

$$\Delta t' = \frac{\Delta m}{\dot{m}} = \frac{m_0}{\dot{m}} \left(1 - e^{-\Delta v / (I_{sp,eff} g_0)} \right) \quad (2)$$

The number of thruster seconds could then be translated into an expected maneuver duration by accounting for the duty cycle and ramp time used by the XACT Δv controller.

After executing a burn, the operations team’s tools would compute, from thruster second telemetry reported by the XACT, the estimated Δv achieved by a given burn:

$$\Delta v = -I_{sp,eff} g_0 \ln \left(1 - \frac{\dot{m} \Delta t'}{m_0} \right) \quad (3)$$

These equations represent an inherently simplified model, as mass flow rate and thrust force (and therefore specific impulse) vary with temperature and chamber pressure. However, they were used in all pre-launch ORTs, and an automation pipeline for maneuver parameter and command sequence file exchange between the operations team and JPL MD-NAV was developed and tested rigorously.

As mentioned above, however, within a few days of the launch this model had to be abandoned. The initial desat attempt, which was intended to dump the angular momentum imparted to the spacecraft from deployment, actually increased the magnitude of the total spacecraft angular momentum, as shown in Figure 18. Clearly, the assumptions underlying the operations team’s calculations on the ground, and implicit in the XACT’s control algorithms, were no longer valid.

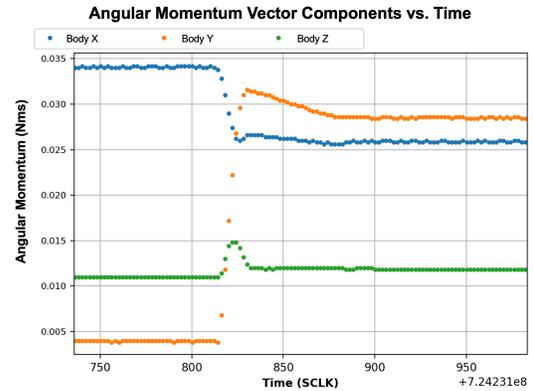


Figure 18: Momentum change during initial desat attempt.

No longer able to rely on pre-launch predictions, the operations team now had to develop a method to estimate the thrust force produced by each thruster. The core of this new model of propulsion system performance was the recognition that the change in the momentum state of the spacecraft, $\Delta \mathbf{h}$, due to a firing of the propulsion system was related to the torque produced by a thruster, $\boldsymbol{\tau}_i$. The momentum change is an inertial quantity, and the torque produced by each thruster is constant in the body frame, but if the relation between the body frame

and the inertial frame is assumed constant over the course of a firing (i.e., attitude transients are small), then the relationship between $\Delta \mathbf{h}$ and $\boldsymbol{\tau}_i$ is

$$\Delta \mathbf{h} = \sum_{i=1}^4 \Delta t'_i \boldsymbol{\tau}_i \quad (4)$$

where $\Delta t'_i$ is the number of thruster seconds fired on thruster i .

When only a single thruster is firing, Equation (4) allows for direct calculation of the thrust force, as the direction of the torque vector is constant in the body frame. This was used throughout the extensive thruster characterization and testing campaign, where sequences of pulses at different duty cycles on each thruster were attempted. When all four thrusters were fired (e.g., during the commissioning burn or some follow-on propulsion system tests), the following algorithm was devised to map the change in momentum to the ‘most likely’ group of three thrusters, and their associated thrust forces: For each group of three thrusters i , j , and k , all fired for $\Delta t'$ thruster seconds, a 3-by-3 matrix $\boldsymbol{\tau} = [\boldsymbol{\tau}_i \ \boldsymbol{\tau}_j \ \boldsymbol{\tau}_k]$ of the associated torque vectors was constructed, and then the thrust forces associated with the thrusters, $T_{i,j,k}$, were computed as:

$$\mathbf{T} = \begin{bmatrix} T_i \\ T_j \\ T_k \end{bmatrix} = \frac{1}{\Delta t'} \boldsymbol{\tau}^{-1} \Delta \mathbf{h} \quad (5)$$

The only valid solutions were those with $T_{i,j,k} \geq 0$.

Furthermore, supplementing spacecraft momentum telemetry with Doppler residuals compiled by JPL MDNAV, as well as the temperature response of the thruster catalyst beds (which is indicative of the effectiveness of the reaction), has allowed for synthesized assessment of propulsion system performance. In particular, while the spacecraft is rotating, the assumption of a constant relationship between the body and inertial frames is invalid, so the model of Equation (4) cannot be used. Doppler data, combined with knowledge of the Earth vector in the body frame, can be used to back out the Δv of a rotating maneuver, from which the average thrust can be derived as well. Figure 20 shows an example of correlated thermal response and Doppler residual data over the course of a 20-minute, fully off-sun rotating maneuver.

Using these models, it has been possible to track the performance of the thrusters, in terms of both thrust force magnitude and Δv per second of firing, over the course of the mission, and assess what tests and modes of operation have produced the best results under different circumstances.

Figure 19 shows the ‘efficiency,’ in terms of $\Delta v/s$, of various rotating maneuvers on the same thruster over the course of approximately one week. Higher efficiency was seen for shorter burn durations, indicating a dropoff in thrust and therefore lower average performance for longer burns. Overall, tracking the performance of the propulsion system, using models developed entirely after the launch, has been invaluable to strategic planning and the characterization of LF’s uniquely extended phase of anomalous operations.

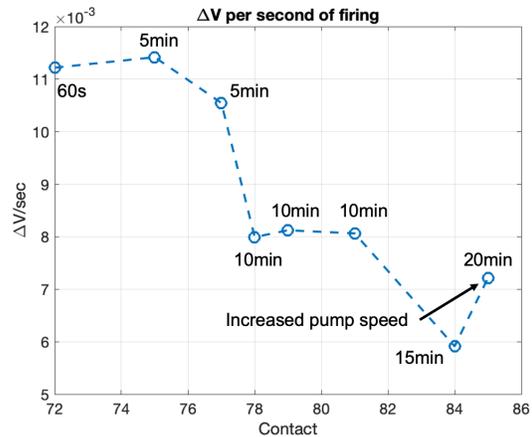


Figure 19: Performance of rotating maneuvers over time.

FAULT PROTECTION CONFIGURATION

The pre-launch strategy of enabling subsets of fault IDs depending on the phase of the mission and the activity in question has broadly carried over into operations, with minor adjustments. Most obviously, many of the anticipated fault protection configurations relating to the propulsion system have never been used, as the approach to commanding the LFPS has changed drastically. However, there have been other refinements to monitors and operational constraints.

In the first track of the mission, in which operators established initial contact with the spacecraft, telemetry and event verification records (EVRs) indicated that the XACT was being power cycled every 21 minutes. Solar array voltage was going to zero, body rates were spiking, and the Iris radio was reporting difficulties with forward error correction all at correlated intermittent times. Tactically, operators and SMEs discerned that the root cause of the power cycles was an improperly defined fault monitor, with FSW implementing a check on the overall XACT mode rather than the sun point algorithm state. The immediate response, accomplished

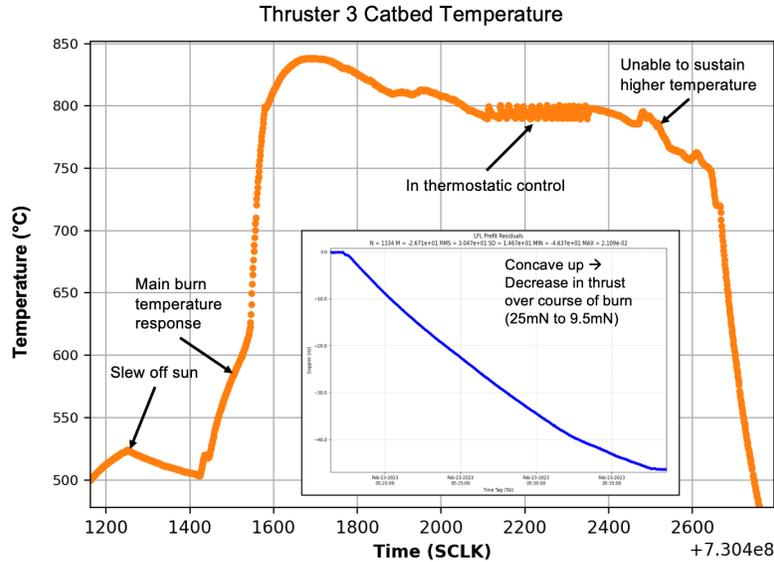


Figure 20: Synthesized temperature and Doppler data used to assess maneuver performance.

within the same track, was to disable that fault ID to avoid additional power cycles, each of which re-tumbled the spacecraft as the reaction wheels spun down.

Fixing the fault monitor would have required a FSW update, and was deemed a low priority in the aftermath of the anomaly. The Safe mode sequences were updated to avoid enabling the fault ID in question, and a flight rule was added to require the fault ID to be disabled at all times. The residual risk associated with leaving the fault ID disabled was deemed low; the fault condition the monitor had been originally planned to detect would require hardware failure of the reaction wheels or sun sensors.

The second adaptation came when it was found that two telemetry channels monitoring the Iris SSPA temperature reported floating, unphysical values when the radio was not transmitting. The decision was made to disable the associated GenMon fault IDs, as many other Iris temperature monitors were in place.

Additional changes were made to GenMon thresholds relating to current draws and temperature limits based on 'rediscovered' margin, as discussed above for the power and thermal subsystems. In particular, the many payload telemetry channels monitored by GenMon had to be carefully tuned to avoid unnecessarily canceling out of a firing of the lasers due to transients in voltages or temperatures. As the operations team's understanding of payload behavior grew, adaptations were also made to flight rules regarding which combinations of payload fault IDs should be enabled under different conditions.

All of these changes to fault protection configuration can be traced to limitations of the LF testbed. With no thermocouples, no temperature telemetry channels reported physical values, so floating readings like those from the Iris SSPA were not specifically noted. Without lasers or emulation of the payload battery, voltage transients during firing could not be observed. Finally, with so many idiosyncrasies chalked up to 'testbed-isms,' the improperly defined XACT fault monitor did not raise any alarm throughout the entirety of the pre-launch phase, and the fault ID had been disabled during all ORTs.

The most significant change in the operations team's approach to fault protection came with the addition of a new channel monitor. When it became clear that relying on the XACT to command desat or Δv burns would be impossible, the operations team had begun characterization and testing of the thrusters to investigate their underperformance and attempt recovery. These tests had to be performed very conservatively, as at any point it was feared that the thruster being fired might return to nominal performance and produce a large torque. Pulse train lengths that could be tested were severely limited by ACS considerations, as exceeding the momentum capacity of the wheels would tumble the spacecraft, likely resulting in loss of mission. However, the propulsion team required longer pulse trains to assess performance and develop a path forward.

To address these competing considerations, the project determined that a new monitor which would cut power to the propulsion system when any of the telemetered momentum vector components exceeded

bounds which could be configured by the operations team was necessary. This “momentum safety net” would stop any ongoing firing before the momentum magnitude reached dangerous levels, allowing for much longer pulse trains to be attempted.

Computing the desired thresholds involved expertise in ACS, FSW, and C&DH. The primary task was to estimate the delay between the true spacecraft momentum exceeding the threshold and the thruster valves closing as a result of power being cut, which in turn determined how long the thrusters would continue to fire, and therefore how much additional momentum could build up. Adding up worst-case delays in the XACT reporting the momentum state, FSW reacting and dispatching the command to cut power to the propulsion system, the execution of the command, and the time the valve would take to close resulted in an upper bound of two pulses, or up to two seconds, of additional firing time. The thresholds were then set based on the worst-case momentum buildup in two seconds of firing.

The operations team assisted in the formal definition of the fault monitor to the FSW team, and tested the new fault monitor on the LF testbed. The procedure to update the spacecraft FSW was adapted from a DITL test procedure, which had been developed by the operations team nearly a year before. Since the introduction of the new fault monitor, it has successfully activated numerous times, keeping the spacecraft in a safe momentum state while allowing progressively more ambitious operation of the propulsion system.

ANOMALY RESPONSE/RESOLUTION

Apart from the anomalous performance of the propulsion system, which has been a dynamic, evolving situation, other spacecraft and MOC anomalies have been identified, documented, analyzed, and resolved by following the project ARP. Throughout these processes, the operations team has identified both limitations and opportunities for improvements. Anomaly resolution is coupled with the closure and approval of the associated Anomaly Record (AR), which is often tied to a specific action: performing follow-on analysis, adding a new flight rule, updating a procedure, or adding a new safeguard.

The number of ARs, which are separated into MOC-related anomalies and spacecraft-related anomalies, was significantly higher early on in the mission, as the operations team worked through GDS issues and adjusted to the differences between testbed behavior observed in ORTs and spacecraft behavior in flight. As the mission has gone on, how-

ever, ARs have been filed less frequently (see Figure 21), as a more experienced operations team and a more robust GDS have together retired or resolved many early concerns.

In this section, four anomalies will be discussed for their significance to the spacecraft and mission system, and how the experience of responding to them allowed the operations team to exercise tools and analytical skills.

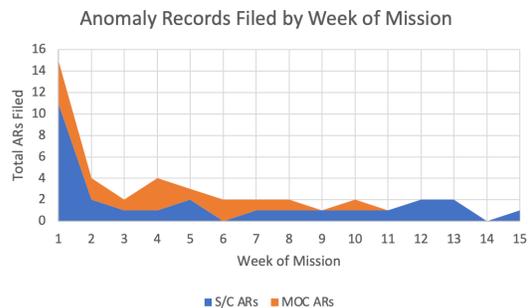


Figure 21: History of spacecraft and MOC ARs.

Solar Array Deployment Anomaly

In the first few days after launch, LF’s power margins were significantly better than expected, as described above. However, trending of spacecraft momentum over the first week of the mission found that, for the amount of time spent in an inertial sun-pointed attitude, the momentum buildup due to SRP was significantly greater than expected. When the spacecraft is in its sun-pointed rotisserie, in which it aligns the solar array normal with the sun vector and rotates at $0.2^\circ/\text{s}$, there is no overall change to angular momentum because the SRP torque vector rotates in the inertial frame. Two avenues of inquiry were opened: one into telemetry relating to the initial deployment of the solar arrays, and another into the model of SRP effects.

Investigation into recorded EPS telemetry from the time of solar array deployment (just a few minutes after launch, before contact had been established with the spacecraft) indicated that one of the two burnwire release mechanisms had not fired due to a driver tripping an overload condition. Essentially, the commanded duty cycle for the burnwire had been too high, and the driver had stopped current from flowing through. This had initially escaped notice because solar array current telemetry was in line with pre-launch expectations, which were (as discussed above) 20-30% too conservative. Further testing, in which the spacecraft was oriented in

such a way as to expose the supposedly undeployed solar array to the sun, confirmed its status through observed voltage and current measurements.

Based on spacecraft CAD, the difference in spacecraft center of mass due to the undeployed array would be around 0.3 cm, since the array itself is extremely light, while the difference in the center of pressure would be 7 cm, greatly increasing the magnitude of SRP torque. The rate of angular momentum buildup was still quite manageable, and this did not present a concern for the operations team or for the project.

With the root cause identified, it would have been straightforward to deploy the array, using a smaller duty cycle to avoid the driver overload condition. However, two considerations made this choice not so obvious. First, deploying the fourth solar array would increase the solar array current, thereby increasing the power dissipation from diodes on the EPS card and in turn resulting in higher temperatures. In addition, with a highly anomalous propulsion system, the project was interested in taking advantage of the increased SRP torque magnitude to keep an option open to perform reaction wheel desaturations in case of a total propulsion system failure. As power margins were still surpassing expectations, the project made the decision to leave the solar array undeployed and close out the AR, with the knowledge that future power considerations (including degradation) might necessitate its eventual deployment.

Corrupted File System Anomaly

In early January, the operations team performed a FSW update on the spacecraft to introduce the “momentum safety net” fault monitor described above. This process required a soft reset of FSW, in order to boot to the new version. After the file system partitions mounted following the soft reset, warning EVRs began to appear when file system operations were being attempted on the */eng* partition, where recorded telemetry is saved.

Operators worked with FSW SMEs and the MSM to tactically respond to the anomaly and ensure the spacecraft was configured safely. The rate of file system statistics reporting was increased to 1 Hz for greater visibility, and a file close command for another application process identifier (APID) was sent to test whether a new file could be opened successfully. As more and more warning EVRs were received indicating that FSW was unable to read from or write to the */eng* partition, it became clear to the operators that the partition had been corrupted. All

attempts at recording, including channelized telemetry and EVRs, were stopped by command before the end of the contact in which the anomaly was discovered.

Immediately following the contact, the team began investigating the root cause and path forward. Correlating warning EVRs and commands with spacecraft telemetry resulted in finding large numbers of double bit errors (DBEs) on the */eng* partition, as shown in Figure 22. The rate of new DBEs was increasing with each additional file system operation failure, until all recording was set to zero. All other partition DBE values were zero, reinforcing the conclusion that the partition was corrupted.

Based on the recommendation of FSW SMEs, the path forward was identified as a format, or complete wipe, of the */eng* partition. This would require an additional soft reset. As a precaution, all recording would remain disabled until after the soft reset and */eng* partition format. A procedure was written, and run on the testbed, to be ready for execution in the next contact. The format was successful, and nominal operations were resumed in the following contact. In all, the analysis, documentation, and resolution of the anomaly were performed overnight in under 8 hours of strategic time by two operators.

Further discussion with FSW and C&DH SMEs identified the root cause of the partition corruption: the soft reset occurred right as a file system operation was in progress, cutting it off and resulting in anomalous behavior. It was a documented flight rule to stop all recording prior to a hard reset (i.e., full power cycle) of the spacecraft to avoid partition corruption, but it had not been documented that this was possible upon a soft reset. Adding a flight rule to account for this consideration was part of the closure of the AR. Since this anomaly, the operations team has had to format the */eng* partition two further times due to different but related file system issues, and has used the same format and soft reset procedure developed as part of this anomaly response.

XACT Power Cycle Anomaly

In early February 2023, operators contacting the spacecraft were surprised by a 5.5° error between the solar array normal and the sun vector. From recorded EVRs, it was discovered that the XACT fault response for “Refs Invalid” (indicating that the XACT could not properly evaluate its onboard ephemeris, in the form of Chebyshev polynomial fits, to compute position and velocity) had executed, power cycling the XACT and sending commands

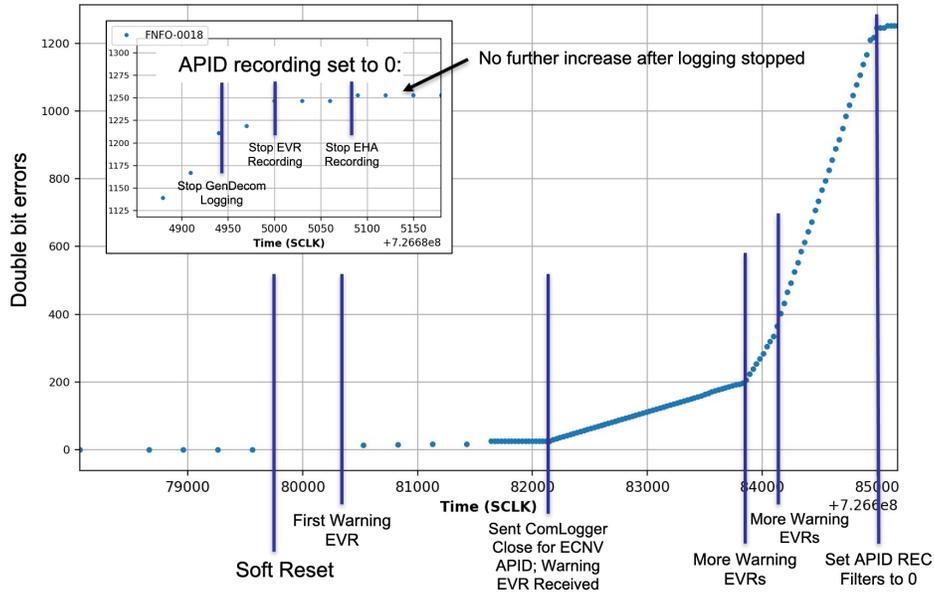


Figure 22: Double bit errors on /eng partition during anomaly.

to set the Chebyshev polynomial coefficients and XACT internal time to the values to which they had last been commanded.

It seemed logical to connect the power cycle and the sun point angle error, and this did turn out to be the case. The XACT internal time is not reported in channelized telemetry, but commands can be sent to dump the entire XACT internal telemetry buffer to a file and downlink that file, which can then be parsed on the ground. The XACT time was found to be around 5.6 days behind the current time, which accounted for the approximately 5.5° error, as the sun moves around 1° per day in the sky. The XACT has an internal model of sun position in the EME2000 frame versus time which it uses to determine how to sun point when in Fine Reference Point mode; it uses its coarse sun sensors only when in Sun Point mode. There had been 5.6 days between the last time the XACT internal time had been set via command and the fault protection power cycle and reinitialization, which explained this discrepancy.

Consultation with ACS SMEs, including BCT engineers, did not uncover any evidence in telemetry of identifiable errors; the root cause of the anomaly was presumed to be radiation effects. The XACT is not immune to single-event upsets, and a power cycle is known to clear them. With no way to mitigate radiation effects once in flight, the possibility of an XACT power cycle remains a low to medium residual risk for the project. Overall, FSW fault protection worked as designed and properly reset the XACT when it had encountered an error.

However, one important risk related to the power cycle was that, if the XACT internal time was not set via command for a significant period of time, a power cycle and reinitialization could result in the XACT time and real time being off by a large amount. If this was on the order of months, it is possible that the “sun pointed” inertial attitude commanded in preparation for a contact would actually point the solar arrays away from the sun, putting the spacecraft in a power negative state and potentially violating sensor keepout zones. To mitigate this risk, the sequences that run in between contacts have been updated to set the XACT internal time to the current time before every contact. Therefore, the difference between real time and the XACT internal time will never be greater than the time between contacts, which is always less than 24 hours.

Safe Mode Recovery Following CFE

During the campaign of rotating propulsive maneuvers in February 2023, new feed-forward torque command binary files were being generated and uplinked for each burn, calculated from the thruster being fired, expected thrust, and duty cycle. As discussed above, each command file was uplinked to the spacecraft file system, then FSW commands in the burn sequence would reference the file system location of the appropriate feed-forward torque command. At this point, the procedure for performing these burns was in flux, as different approaches and process optimizations were being trialed. For

one burn, the step to uplink the new feed-forward torque command file was omitted from the procedure. Since there was no file at the expected file system location, the command in the burn sequence failed, and the associated fault response executed, sending the spacecraft to Safe mode.

The Safe mode sequence that is used when the propulsion system is active first configures the LFPS to a safe state. In addition, it configures the Iris radio to receive-only mode at a low data rate, allowing the spacecraft to cool down, before sending a series of transmit-only beacons in between periods in receive-only, swapping between the two pairs of LGAs. The immediate tactical response to the safing event is detailed below, as operators followed the Safe mode contingency plan developed pre-launch and included in operator procedures (in T+H:MM:SS):

- T+0:00:00** Command failed fault. SMS executes.
- T+0:00:15** Cause of safing identified.
- T+0:00:30** Additional personnel called in to support recovery. Began executing Safe mode contingency procedure.
- T+0:05:00** Identified risk of instrument keepout zone violation during transition to Sun Point mode.
- T+0:06:00** Requested station to configure for 62.5 bps uplink rate.
- T+0:08:00** Team discusses possibility of commanding back to full duplex through the currently active ‘topcap’ antenna pair, which are pointed away from Earth.
- T+0:18:00** SMS configures topcap antenna pair to transmit-only; station reports RF observables, indicating topcap pair can transmit to Earth.
- T+0:20:00** SMS configures prop tank antenna pair (in view of Earth) for receive-only.
- T+0:24:00** Operators request station begin an uplink sweep.
- T+0:30:00** With team concurrence, commands are sent to cancel out of the SMS and configure the Iris for full duplex.
- T+0:34:00** Performed full spacecraft health check-out. Observed spacecraft momentum state was slightly different from prior to fault. Observed high LFPS manifold pressure. Determined several pulses of burn executed before fault response executed.

- T+0:42:00** Began re-enabling fault responses.
- T+0:44:00** Reset telemetry reporting rates.
- T+0:49:00** Began downlinking EVRs from ‘black-out’ period in receive-only.
- T+1:00:00** Transitioned back to Normal mode.
- T+1:04:00** Powered on propulsion system to relieve pressure in manifold.
- T+1:09:00** Powered on payload to assess detector current. Observed nominal value, indicating no instrument keepout zone violation.
- T+1:14:00** End of track.

The decision to recover the spacecraft in the same contact of the safing event was made with concurrence from MOAM, MSM, GT operators, and other team members on the tactical line. Given that the cause of the safing event was well understood, there was no risk in a timely recovery.

There were two important takeaways from this event for the operations team. First, the successful execution of the SMS, and of the Safe mode recovery procedure, gave confidence in the team’s ability to arrive at consensus and make prompt tactical decisions, and validated pre-launch preparations for anomaly response. Second, the CFE that resulted in the safing event was avoidable and caused by a procedural error. After this anomaly, the procedure for all propulsive maneuvers was generalized so that it could be reviewed and confirmed to cover all necessary steps, rather than requiring the operations team to develop a bespoke procedure each time.

CONCLUSION

Operations systems engineering for LF has taken place in a unique, dynamic environment. As the first group of students to operate a JPL spacecraft, the GT team’s preparation focused on tactical responsibilities, like commanding and monitoring, and developing a systems engineering understanding of the spacecraft itself. The team has made major post-launch adjustments to workflow and the balance between tactical and strategic priorities, as the operations systems engineering work detailed in this paper has grown in scope and importance to the mission.

LF’s operations systems engineers are part of a holistic mission operations system (MOS) that includes not only the spacecraft and its tactical operators but also the GDS, communications infrastructure, and a host of supporting engineers and administrators. The MOS in turn interacts with the

outside world, as the laws of physics determine trajectory options and pragmatic, personal constraints demand flexibility from the people involved. As operations systems engineers, the GT team has been able to mediate relationships within and external to the MOS: closing the loop between model-based predictions and observed spacecraft behavior to inform activity development; adapting constraints to extract performance from margins; and responding, tactically and strategically, to anomalies to keep the spacecraft safe with minimal impact to schedule.

LF's story is still unfolding. Its months-long period of anomalous operation is uncommon among interplanetary spacecraft, and has demanded attention and resilience from the entire project. Throughout, continued efforts to recover and adapt the mission have relied on the successful application of systems engineering concepts, from modeling and simulation to constraints, technical budgets, and failure analysis, in the context of operations. It is the hope of the GT mission operations team that the experience of the LF project can serve as a model for future student-led operation of interplanetary spacecraft, and that lessons derived from that experience can inform expectations and implementations for future projects.

ACKNOWLEDGEMENTS

The work presented in this paper would not have been possible without years of work by the LF team at JPL, MSFC, GT, and around the country. I am grateful for the heroic efforts of my colleagues on the GT operations team, and for the continued support of my advisor, Dr. E. Glenn Lightsey.

This research was conducted at the Georgia Institute of Technology under contract from the NASA Jet Propulsion Laboratory.

REFERENCES

- [1] P. O. Hayne, B. A. Cohen, R. G. Sellar, R. Staehle, N. Toomarian, and D. A. Paige, "Lunar Flashlight: Mapping Lunar Surface Volatiles Using a CubeSat," in *Proceedings of the Annual Meeting of the Lunar Exploration Analysis Group*, 2013. <https://sservi.nasa.gov/wp-content/uploads/2014/04/7045.pdf>.
- [2] A. Rizvi, K. F. Ortega, and Y. He, "Developing Lunar Flashlight and Near-Earth Asteroid Scout Flight Software Concurrently using Open-Source F Prime Flight Software Framework," in *Proceedings of the Small Satellite Conference*, Advanced Technologies, SSC22-VIII-03, 2022. <https://digitalcommons.usu.edu/smallsat/2022/all2022/104/>.
- [3] N. Cheek, C. Gonzalez, P. Adell, J. Baker, C. Ryan, S. Statham, E. G. Lightsey, C. Smith, C. Awald, and J. Ready, "Systems Integration and Test of the Lunar Flashlight Spacecraft," in *Proceedings of the Small Satellite Conference*, Beyond LEO, SSC22-II-06, 2022. <https://digitalcommons.usu.edu/smallsat/2022/all2022/149/>.
- [4] D. Andrews, G. Huggins, E. G. Lightsey, N. Cheek, N. Daniel, A. Talaksi, S. Peet, L. Littleton, S. Patel, L. Skidmore, M. Glaser, D. Cavender, H. Williams, D. McQueen, J. Baker, and M. Kowalkowski, "Design of a Green Monopropellant Propulsion System for the Lunar Flashlight CubeSat Mission," in *Proceedings of the Small Satellite Conference*, Propulsion, SSC20-IX-07, 2020. <https://digitalcommons.usu.edu/smallsat/2020/all2020/155/>.
- [5] N. Cheek, N. Daniel, E. G. Lightsey, S. Peet, C. Smith, and D. Cavender, "Development of a COTS-Based Propulsion System Controller for NASA's Lunar Flashlight CubeSat Mission," in *Proceedings of the Small Satellite Conference*, Advanced Technologies, SSC21-IX-06, 2021. <https://digitalcommons.usu.edu/smallsat/2021/all2021/196/>.
- [6] U. Wehmeier, Q. Vinckier, R. G. Sellar, C. G. Paine, P. O. Hayne, M. Bagheri, M. Rais-Zadeh, S. Forouhar, J. Loveland, and J. Shelton, "The Lunar Flashlight CubeSat instrument: a compact SWIR laser reflectometer to quantify and map water ice on the surface of the moon," in *Proc. SPIE 10769*, CubeSats and NanoSats for Remote Sensing II, 107690H, 2018. <https://doi.org/10.1117/12.2320643>.
- [7] M. Starr and E. G. Lightsey, "Development of Tactical and Strategic Operations Software for NASA's Lunar Flashlight Mission," 2023. Master's Report, Georgia Institute of Technology.
- [8] A. Paletta and E. G. Lightsey, "Development of an Autonomous Distributed Fault Management Architecture for the VISORS Mission," 2023. Master's Report, Georgia Institute of Technology.

- [9] D. McDonald and E. G. Lightsey, “Fault Management in Small Satellites,” 2022. Master’s Report, Georgia Institute of Technology.
- [10] C. R. Smith, L. M. Littleton, E. G. Lightsey, and D. Cavender, “Assembly Integration and Test of the Lunar Flashlight Propulsion System,” in *AIAA SCITECH 2022 Forum*, AIAA 2022-1731, 2022. <https://arc.aiaa.org/doi/abs/10.2514/6.2022-1731>.