

A DELICATE BALANCE OF TORQUE AND THRUST: HOW LUNAR FLASHLIGHT USED ROTATING MANEUVERS TO MAKE ONE THRUSTER DO THE WORK OF FOUR

Timothy P. McElrath^{*}, Steven M. Collins[†], Celeste R. Smith[‡], Kevin D. Lo[§],
Nathan Cheek^{**}, and Michael J. Hauge^{††}

After launch, Lunar Flashlight (LFL) was beset by thruster failures, ultimately leaving it with only one operable thruster. By developing the rotating Trajectory Correction Maneuver (TCM), the team made use of the reaction wheels to recover trajectory control capability. In this maneuver, the inertial momentum vector is rotated in a plane, while the spacecraft body momentum remains nearly stationary, by carefully balancing the thruster torque, spacecraft rotation rate, and initial momentum state.

INTRODUCTION

On December 11, 2022 a SpaceX Falcon 9 launched Hakuto-R Mission 1 and Lunar Flashlight (as a secondary payload) on a 3-month low-energy trajectory to the Moon. Lunar Flashlight is a JPL technology demonstration mission with a goal of investigating the distribution of ice deposits near the south pole of the Moon. The Georgia Institute of Technology's (GT's) Space Systems Design Laboratory (SSDL) operates the spacecraft, with JPL providing the navigation and mission design. GT designed and integrated the propulsion system with the guidance and support of Marshall Space Flight Center (MSFC), and the XACT-50 attitude control system (ACS) was obtained from Blue Canyon Technologies (BCT).

Starting with the first checkout burn (at L+2d), the propulsion system misbehaved, exhibiting severe thrust reductions on first one and eventually all of the thrusters. LFL has four canted thrusters that are designed to work together to provide ΔV and desaturate the reaction wheels. The thrusters use the ASCENT green monopropellant as an alternative to hydrazine. Functionally, they are generally similar to hydrazine thrusters, requiring the use of a catalyst bed heater to

^{*} Chief Engineer for the Mission Design & Navigation Section, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109.

[†] Guidance & Control Engineer, Guidance & Control Section, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109.

[‡] Propulsion Engineer, Flight Software And Avionics Systems Section, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109.

[§] Guidance & Control Engineer, Guidance & Control Section, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109.

^{**} Propulsion Engineer, Mobility & Robotic Systems Section, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109.

^{††} Systems Engineer, Planning & Execution Systems Section, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109.

achieve the desired performance. The propulsion system uses an electric pump to supply the required pressure. Each thruster has a nominal thrust of around 100 mN and can be commanded (by the XACT-50) to fire pulses as short as 50 msec at a 1 Hz rate.¹

The XACT-50 is a fully-integrated ACS system comprising: 1) 3 reaction wheels, each with a nominal storage capacity of 50 milli-Nms, and closely aligned with one of the spacecraft axes 2) a star tracker, 3) MEMS gyros, 4) 4 coarse Sun sensors, and 5) a CPU to manage the ACS hardware and command the thrusters. All of this is contained in a 1U form-factor package. Throughout the mission, the XACT system has performed well, and in particular, the reaction wheels have worked perfectly, a matter of critical importance as they are the only continuously-available attitude control actuator.

LFL spacecraft operations were of necessity built around DSN contact periods of about 2 hours or less. Primarily, the IRIS radio could not operate in full-duplex (two-way) mode for any longer without overheating, requiring about 8 hours between contacts to recover thermally.* The propulsion system likewise had a limited catalyst-bed operating duration of about 30 minutes due to thermal soak-back into the thruster valves, thus limiting the ΔV maneuver duration. The power system was sized to operate all the catalyst beds for only a few minutes when the solar arrays were not Sun-pointed. In between contacts, the spacecraft pointed the Z axis (and solar arrays) towards the Sun and rotated about the Z axis at a nominal rate of once per 30 minutes, thus limiting momentum buildup from solar radiation pressure torques. During the contacts the spacecraft stopped the rotation and held an inertial attitude (other than using the Sun as one reference direction).²




Figure 1: Lunar Flashlight thruster configuration (see reference 1, Smith & Cheek, 2023).

* Flight experience proved that the thermal models had been conservative, but the pattern of contacts lasting no more than 2.5 hours was maintained due to DSN pass scheduling.

Because it was a secondary payload, the nominal LFL launch trajectory required 77 m/s of ΔV within the first month to reach the necessary lunar arrival conditions to insert into the planned Near Rectilinear Halo Orbit (NRHO). This ΔV was divided into Trajectory Correction Maneuvers (TCMs) of no more than 29 m/s (for the reasons mentioned above), occurring every 4 days in the first half of January 2023. While the net ΔV of these maneuvers was not sensitive to timing changes of a few days, the cost increased significantly with a delay measured in weeks, and the planned trajectory became impossible with longer delays. Unfortunately, the propulsion system anomalies immediately threatened to impose a very long delay.

PROPULSION SYSTEM OVERVIEW AND EARLY OPERATIONS HISTORY

The thruster force vectors are all pointed within 12 degrees of the spacecraft -Z axis, so that the translational control (for ΔV maneuvers) can be accomplished efficiently. Figure 1 shows the thruster configuration. Each thruster is canted in a different direction, such that any pair of thrusters provides torque that is mostly aligned with the positive or negative side of each spacecraft axis, as shown in Table 1. The magnitude of these torques varies by a factor of 10 between the Y- and Z-axes. While the Z-axis moment of inertia is the smallest one, the moments of inertia only vary by 30%, and the primary reason for the low torque authority about Z is the need to generate ΔV efficiently.

Table 1: Thruster Force Directions and Torque Moment Arms

Thruster number	Thrust unit vector			Torque moment arm components, cm		
	X	Y	Z	X	Y	Z
One	-0.147	0.147	-0.978	-4.707	-11.113	-0.963
Two	0.147	0.147	-0.978	-4.707	10.775	0.912
Three	0.147	-0.147	-0.978	4.251	10.775	-0.981
Four	-0.147	-0.147	-0.978	4.251	-11.113	1.031

The nominal maneuver technique involved pulsing all four thrusters at 1 Hz, such that the net torque on the spacecraft was zero. Due to small thruster misalignments and center-of-mass offsets, the XACT controller would adjust the pulse sizes to first null the attitude rates, and then also to control the attitude. The reaction wheels were not used during a ΔV maneuver, since the thruster torques would overwhelm their torque capability. Apart from during a TCM, the reaction wheels were always used to maintain the spacecraft attitude.

The first two propulsive activities on the spacecraft consisted of a propulsion system checkout, which fired all 4 thrusters 30 times with pulses of 50 msec, and a reaction wheel desaturation (“desat”). The result of the desat was a higher spacecraft momentum state than the initial momentum state, so clearly something was badly wrong with the thruster performance. After a closer investigation, it was clear that Thruster 1 had produced barely any thrust (and consequently almost no torque), despite having been commanded up to the maximum pulse duration. Additionally, further review of the telemetry from propulsion system checkout revealed that this under-performance could be seen there too. Performing a TCM with a severely underperforming thruster was out of the question, and the pressing need became developing activities to understand (and possibly recover) the propulsion system performance, while not endangering the spacecraft.

The original configuration of the LFL system included an autonomous XACT-commanded desat as a response to any high momentum state while thruster usage was enabled. Tripping this response in the presence of unknown and possibly variable thrust output could result in saturation of the reaction wheels and the loss of the spacecraft, and so this behavior was disabled.

However, the reaction wheels still needed protection against excessive momentum states. During the propulsion system checkout activities described below, a Flight Software (FSW) update was developed to cut off power to the propulsion system (thus stopping further thrusting) if the system or wheel momentums exceeded parameterized limits. Since the propulsion system continued to behave erratically, the momentum safety net was uploaded to the spacecraft, and became an important part of all future propulsion activities.

TCM METHODS CONSIDERED

As thruster testing was being planned and executed, a parallel activity considered all the ways of generating useful maneuvers, to be ready for any outcome. The methods considered for performing TCMs can be split into two categories: 1) with all four thrusters generating useful thrust, using the XACT to control them in a closed-loop manner, or 2) with less than four thrusters firing in some pattern, using the reaction wheels to control the attitude while temporarily storing the momentum in between attitude changes, such that the average torque is canceled out, leaving the translational impulse. This second approach would necessarily be commanded in an open-loop manner (and not using the XACT to respond to dispersions). Due to the limits on the duration of propulsion system operation, the achievable ΔV per maneuver for each method became an important evaluation consideration, especially in view of the planned Lunar Orbit Insertion (LOI) magnitude of 9 m/s.

Table 2: TCM Methods Considered

#	Description	# of thrusters	ΔV in 20 min (m/s)	Implementation	
1	Original plan (fully functional thrusters, 90% duty cycle)	4	33	Nominal plan	Use XACT DeltaV command
2	Thrusters all > 20 mN	4	>6.6	Modified XACT usage	
3	Thrusters all > 6 mN	4	>2.0		
4	“Funny desat” (assuming 45s slews)	2	2.8	Use open-loop attitude and thruster commands	
5	Rotating TCM	1	2.5		

Table 2 shows all the methods considered. The first entry represents the nominal performance, with all thrusters operating within the unmodified parameters of the XACT control algorithm, at a relative thrust level within a few percent. If the thrusters were operating at different levels, the XACT usage probably could have been changed to allow for a larger initial thrust difference. In addition, there are two other significant thrust/duty-cycle thresholds, as listed in the next two entries in Table 2. Pre-launch testing had showed that the thrusters would self-heat at duty cycles above 20%, thus allowing the catalyst bed heaters to be turned off and permitting a much longer burn duration (due to power limitations). Finally, the XACT could not command pulse widths of less than 50 msec or more than 950 msec in TCM mode, so the thrust ratio from the highest-performing to the lowest-performing thruster could not exceed 19:1 and must be a bit lower to permit control margin. Even the methods requiring XACT usage changes would likely

have been easier to implement than the RW-enabled methods, but as further testing indicated that at least one thruster was producing no useful thrust, all 4-thruster methods were dropped.

The last two entries in Table 2 show the two reaction-wheel-enabled methods. The first method uses pairs of thrusters that together produce only small amounts of Z-axis torque and large amounts of ΔV (either Thrusters 1 & 3 or Thrusters 2 & 4). The Z-axis momentum is stored until it nears the limits of the reaction wheels, at which point the spacecraft re-orients by ~ 95 degrees to fire a single thruster to reset the spacecraft momentum, which now lies along (and opposite to) the thruster torque vector (near the X-Y plane).^{*} This pattern is then repeated until the ΔV is complete or the maximum propulsion system usage duration limit is reached. The overall average acceleration depends on the duration of the 2 slews in each cycle, when no useful thrusting is happening. More critically, the two thrusters used must have sufficiently similar and predictable thrust levels to avoid accumulating large enough momentum errors to trip the momentum safety net. During testing 12 days after launch, Thruster 2 suddenly dropped from full to negligible thrust, leaving Thruster 4 as the only fully-operational thruster, and ruling this method out.

The last method uses continuous pulses on one thruster while rotating under reaction wheel control, such that the body-fixed torque produced averages out over time inertially while the translational impulse accumulates. The properties of and experience with this “Rotating TCM” are the subject of the rest of this paper.

ROTATING TCMS: CONCEPTUAL DEVELOPMENT

The basic condition for rotating TCMS is that the spacecraft rotate about an axis that is normal to the active thruster’s torque vector. Since the torque vector is fixed in the spacecraft body frame, it will trace out a plane inertially as it rotates about a normal axis. Over a full revolution, the sum of the inertial torques from the thruster’s pulses will be zero. Rotating about the thruster force direction (as was used for LFL) is the most efficient choice[†], as there will be no cosine losses in the sum of the impulses, as shown in Figure 2.




Figure 2: Basic rotating TCM condition. Torques sum to zero inertially over a full revolution.

^{*} A common misconception neglects to remember that angular momentum is fixed in inertial space, until some external torque (such as a thruster firing) is applied. A reaction-wheel-equipped spacecraft can in principle reorient itself to align any direction in the spacecraft body frame along this inertial momentum direction. This technique was used by LFL to accomplish numerous desats using only one thruster.

[†] Co-author Steve Collins was the first person to note this in our development discussions.

The momentum vector in the spacecraft body also rotates, but this happens in between thruster pulses (since the inertial momentum is fixed until the next torque impulse is applied). For the right choice of an initial momentum state, the thruster firings produce a torque that is perpendicular to the momentum, and if the average torque applied by the thruster has the same magnitude as the cross product of the reaction wheel momentum and the spacecraft body rotation vector, then the reaction wheel momentum will remain essentially stationary (at time scales longer than the pulse frequency) in the spacecraft body frame. Note that the spacecraft body momentum from the spacecraft body rotation is perpendicular to the torque vector by construction, and so that momentum component, plus any reaction wheel momentum in that direction, do not participate with the thruster torque.




Figure 3: Simplified rotating TCM condition in the spacecraft body frame. The spacecraft is rotating about +Z (counterclockwise), so inertial vectors appear to rotate clockwise in the spacecraft frame. The reaction wheel momentum limit was enforced by the momentum safety net.

Figure 3 illustrates this condition, using a simplified system where the Thruster 4 torque lies entirely (instead of mostly) in the X-Y plane, and the rotation of the spacecraft body is about the +Z axis. The magnitude of the momentum impulse ΔH and the time between pulses is exaggerated for illustrative purposes. Additionally, Figure 3 shows the notional momentum limits of the X and Y reaction wheels, which must be respected during rotating TCM activities.

If this momentum condition is not perfectly achieved, the momentum state will follow circular paths in the spacecraft X-Y plane, with a radius proportional to the mismatch ratio. Consequently, the momentum will at times have a larger magnitude than the ideal condition, and hence possibly violate the X or Y reaction wheel limits. Since the spacecraft momentum in each axis is a spacecraft telemetry item, plots like this one were easily viewable in near real-time, and became an important analysis tool, as will be shown later in Figures 10 and 12.

Since the establishment of the correct initial momentum magnitude in the rotation plane is so important for a rotating TCM, how can this be accomplished? We considered two approaches before the one we ultimately adopted.

If the spacecraft has reached the rotating TCM attitude (with the thrust vector aligned with the desired ΔV direction), but has not yet started rotating, then a series of pulses could be fired to increase the size of the in-plane momentum to the appropriate value. (This is the equivalent of the single thruster desat mentioned earlier.) However, after the rotation begins, the timing of the start of the main burn pulse train would need to be phased correctly with the momentum direction in an open-loop manner. While this was our initial thinking, we realized that the momentum setup could instead be accomplished by modulating the initial thruster firings, without being concerned about the inertial rotation phase.

It turns out that in the inertial plane normal to the rotation axis (in which the torque vector moves), a series of rotating torque impulses will always cause the inertial in-plane momentum to follow a circular path, regardless of the starting point, as long as the reaction wheel limits are not exceeded. However, if this momentum circle is not centered about zero, then its magnitude must be limited to avoid hitting those limits, by as much as a factor of two if the circle passes through zero.

If the initial in-plane momentum state is nearly zero, then starting a pulse train at half the ultimately desired thrust level, for half a revolution, will produce half a momentum circle that ends at the centered main burn circle, thus setting up the desired inertial momentum state. At this point, the pulses can be increased to the full main burn thrust level, and the inertial momentum will progress around the efficiently-centered momentum circle. Predicting the thruster operation at two different pulse sizes is a significant drawback to this approach. However, we realized that there is a simpler scheme.




Figure 4: Establishing the correct inertial momentum state, with a 60-deg pulse train, followed by a 60-deg coast. The torque direction rotates with spacecraft body. The long pulse train is used for the main burn.

Figure 4 shows the inertial view of the ultimately chosen method used operationally to set up the initial momentum. If a pulse train (operating at 100% of the planned main burn thrust) begins at zero momentum and continues for 60 degrees (one sixth of a period), then it will have reached the intersection of the centered momentum circle and the starting momentum circle. After a pause in thrusting for 60 more degrees of rotation, the torque direction will now be tangent to the centered momentum circle. At this point, the main burn can be started and continued indefinitely without violating momentum constraints. When the end of the activity nears, this same pattern can be used in reverse to remove the momentum, with a 60-deg coast, followed by a final 60-deg take-down burn. Figures 10 and 12 below show an operational result of this same sequence in the spacecraft body momentum frame.

Using a setup and take-down burn in this manner has several advantages: 1) the momentum state for the main burn is quickly accomplished, 2) only one thrust level is needed, 3) the size of the momentum circle is automatically sized correctly for the achieved thrust level as long as the setup burn and main burn have the same thrust, and 4) the thrusting activities do not need to be phased with any external condition, but only timed correctly with respect to each other. All of these factors proved advantageous operationally.

ROTATING TCMS: IMPLEMENTATION

Based on this conceptual understanding, many implementation details needed to be chosen within the constraints of the LFL spacecraft, as follows:

1. The X-Y momentum magnitude must remain below the reaction wheel momentum constraints throughout the activity, with enough margin for the momentum safety net to operate. The safety net setting was usually 0.04 Nms, so a target momentum magnitude of no more than about 0.03 Nms seemed prudent.
2. The reaction wheels must have enough torque to be able to rotate the momentum at the desired rate, according to the simple formula $T = W * H$, where W is the spacecraft body rotation rate, H is the momentum magnitude in the plane perpendicular to W , and T is the gyroscopic torque magnitude. The maximum reaction wheel torque capacity is 7 milli-Nm, which for a momentum magnitude of 0.03 Nms would permit a rate of 0.233 rad/sec with no margin. The adopted rotation rate was less than half of this maximum for other reasons, so this was not a driving constraint. As will be discussed later, the XACT control system must also command the required gyroscopic torque level, which it was not configured to do initially, to our surprise.
3. The thruster pulse width choice combined with a 1 Hz firing rate produces an average torque level. The thrusters also came with a recommendation to avoid firing pulses of around 200 msec (within about 50 msec). Testing on Thruster 4 showed that shorter pulses produced more impulse than would be expected from the product of the nominal force and the on-time, with 150-msec pulses showing an impulse of 27 milli-Ns (for an effective average thrust of 180 mN. This impulse led to a workable rotation rate, while a 250-msec pulse would likely have been too large for the momentum limits and/or rotation rate. However, 250-msec pulses were used later, due to a recommendation to lower the pump speed.
4. The thruster moment arms are all about 12 cm. An impulse per pulse of 27 milli-Ns, occurring once per second results in an average torque of 3.24 milli-Nm. For a momentum magnitude of 0.03 Nms in the equation above, we would calculate a rotation rate of 0.108 rad/sec. This is very close to the more-conveniently-describable value of 6 deg/sec (0.1047 rad/sec), so 6 deg/sec was chosen as the nominal rate.

- The LFL propellant tank was designed to operate under a steady-state $-Z$ acceleration (since all the thrusters had a large $-Z$ thrust component), and so the outlet was in the middle of the $+Z$ face. Since it also had to operate in weightlessness for the first pulse, the tank contained an extensive Propellant Management Device (PMD), operating on surface tension, as shown in Figure 5. Rotating about any of the thruster axes would cause the propellant to migrate to the X-axis extrema of the tank, leaving the outlet uncovered (other than by the PMD sponge). At the desired rotation rate of 6 deg/sec, would the propellant tank be able to supply propellant to the thrusters? After a sleepless night and further analysis discussed below, we convinced ourselves that this would work.*




Figure 5: PMD configuration. Vanes are shown in blue, and sponge is shown in red. Tank bottom is towards $+Z$, and the tank's largest extent is in $\pm X$. Outlet is under the sponge.

For a rotation rate of 6 deg/sec, the radial acceleration at the largest distance within the tank from the spin axis (~ 12.3 cm) was 1.34 mm/sec^2 . However, with an average force of 27 mN, the acceleration along the thruster direction would be 2.03 mm/sec^2 . If the thruster force was constant, then the surface of the propellant would form a parabola normal to the local acceleration direction, with a maximum slope of ~ 33 deg. Although this could still be a problem for a nearly empty tank, it was not a concern initially.

In addition, the PMD system had been analyzed against a lateral acceleration of 3 mm/sec^2 (in the X-Y plane, although not rotating), significantly higher than the proposed acceleration. Since the spacecraft rotation would have to start well before the actual thrusting, the propellant would mostly not be in contact with the sponge assembly in the middle when thrusting started. The analysis had concluded that 3 mm/sec^2 would not drain the sponge. Since the centripetal acceleration near the outlet was only about 0.6 mm/s^2 at the intended rotation rate, this factor was retired as a concern.

Several other operational constraints that were not strongly tied to the rotating nature of the maneuver still had to be addressed, as follows:

- Firing a thruster did not automatically cause the XACT to command compensating feed-forward torque to the reaction wheels. Without the feed-forward torque, attitude and rate errors would have to build up until the controller gains generated the required reaction

* Since if it didn't, we were at the end of our options...

wheel torque. This attitude disturbance is undesirable, so a feed-forward torque was applied while the thruster was firing. The FSW interface with the XACT did not foresee needing to issue this command, but the feed-forward torque could still be implemented by using a binary file with a built-in duration, which was sent by the FSW to the XACT controller.

2. While rotating, the star tracker would not be able to maintain lock, and the attitude would be propagated on gyros. Once the rotation stopped and the star tracker reacquired, we could see a small jump in the attitude estimate due to the gyro drift. Even for the longest burns, we found that the gyro drift was not large enough to cause problems. The fault protection parameters for loss of attitude knowledge were updated to make sure that they would not activate unnecessarily.
3. Many of the later maneuvers (from February 2023 onwards) required turning the spacecraft completely away from the Sun, such that no power was generated on the solar arrays. The allowable duration in this state had to be carefully considered initially, and then adjusted as our experience with these attitudes grew.
4. The LFL instrument boresight (along spacecraft +Y) must be kept at least 10 degrees away from the Sun. Normally, it is 90 degrees away from the usual Sun-pointing direction (+Z axis), but by turning more than 90 degrees off-Sun, the potential existed for Sun exposure. The initial and final slews had to be checked to make sure that they did not violate this pointing constraint. For simplicity, the entire time spent rotating was constrained to be an even rotation period (1 minute for 6 deg/sec), so that these two slews would traverse about the same path. Furthermore, maneuvers close to 90 degrees off-Sun were prohibited, since rotating would necessarily cause the instrument boresight to sweep across the Sun. Any such maneuvers were done with a cosine loss and vectorization from the nearest acceptable attitude.
5. When turning away from the Sun, the spacecraft often turned far enough with respect to the Earth that an antenna swap was necessary. The time of the antenna swap was managed in the burn sequence based on the attitude time profile from simulated slews. While this required two commands which had to be spaced a second apart due to a FSW limitation (and thus were not at exactly at the desired point), the antenna swaps always worked well. This and other tests suggested that the low-gain antenna patterns had significant overlap.

While the sequenced commands handled a nominal maneuver, the fault protection did not respect the instrument pointing when slewing back to the Sun after a system safe mode event. The chances of violating the instrument pointing were small (20 degrees out of 360) while traversing the X-Y plane, but not zero. This risk was accepted, as it was difficult to retire, and it was also thought that the instrument could tolerate brief Sun exposure more readily than the letter of the constraint indicated.

ROTATING TCMs: OPERATIONAL EXPERIENCE

Now that we have covered the theory and practice of the rotating TCMs on LFL, we return to the operational history of the mission. As the flight software updates and thruster testing continued, mission design worked to adjust the trajectory to allow a later start of significant ΔV . Rather than 20-30 m/s per TCM, we now were planning for only 3-5 m/sec/day, with rotating TCMs executing up to thrice daily, in between contacts if necessary. The original pre-anomaly plan also included a 9 m/s Lunar Orbit Insertion, which was now impossible. The unstable nature of the NRHO provided the basis of a solution to both problems. The dynamics around the NRHO al-

lowed for trajectories that spent many revolutions (of about 6 days) approaching the target trajectory, such that very small ΔV s could be used to reach the target, at the expense of a longer flight time. This technique reduced LOI to achievable levels, as shown in Figure 6.

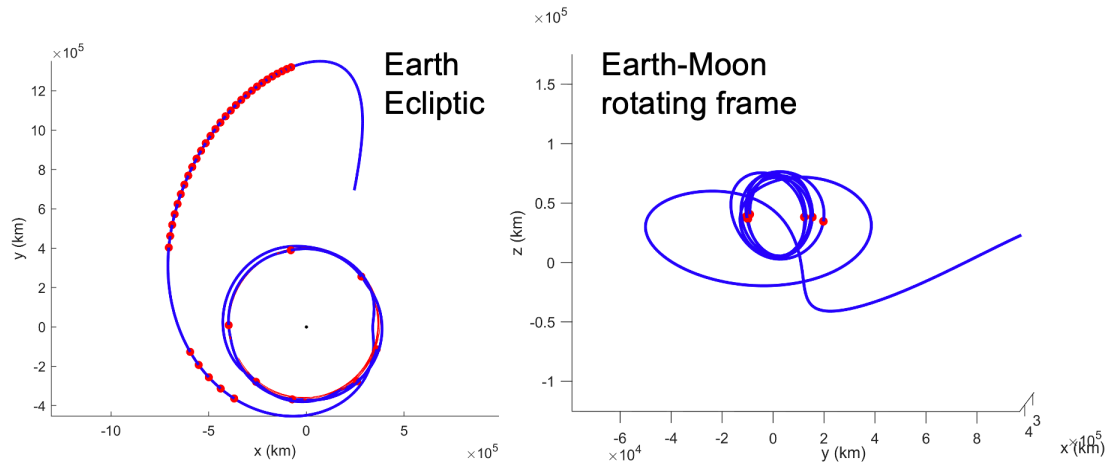


Figure 6: Later thrusting start trajectory example. Starting 1/27, thrusting is 3 m/s/day. LOI is 3 m/s. Red dots are ΔV s of up to 3 m/s. This trajectory costs +30 m/s compared to the pre-launch plan, but with later trajectories the cost was reduced to +10 m/s.

Before the approach to the NRHO, the trajectory spent over a month in the 3-body region surrounding the Moon, following a path that allowed the lunar arrival conditions to be connected to the NRHO approach. With thrusting starting later, the arrival condition changed, and we had to find a new trajectory to connect it to a slower NRHO approach. This was successfully accomplished with a LOI delay of only 2-3 weeks and a ΔV total of only 10 m/s more than the original plan.

Table 3: Rotating TCM Development History with Thruster 4

UTC date & contact start time (in 2023)	Contact #	Rotation rate (deg/s)	Duty cycle	Duration of thrusting (s)	Comments
1/18 03:00	c069	4	n/a	0	Attitude-only test (see Fig. 7)
1/19 01:15	c070	4	15%	20	Set-up and take-down burns only (see Fig. 8)
1/20 22:50	c072	6	15%	80	First main burn (60s) (see Figs. 9, 10)
1/22 12:40	c075	6	15%	320	First 5-minute burn (see Figs. 11, 12)
1/23 13:45	c077	6	15%	320	Increased automation
1/24 13:45	c078	6	15%	620	First 10-minute burn
1/24 21:45	c079	6	15%	620	Updated pump speed and FF torque
1/26 03:05	c081	6	15%	620	First with gyro torque compensation
1/27 13:00	c084	6	25%	920	First 25% duty cycle and rotation axis correction
1/28 12:00	c085	6	25%	1220	First 20-minute burn
1/29 08:45	c087	6	25%	1078	Planned for 20 minutes, but a severe thrust drop occurred near midpoint, so thrusting was halted

Notes: Desats on contacts c082 and c086. Thrusting duration includes setup and (if present) take-down burns.

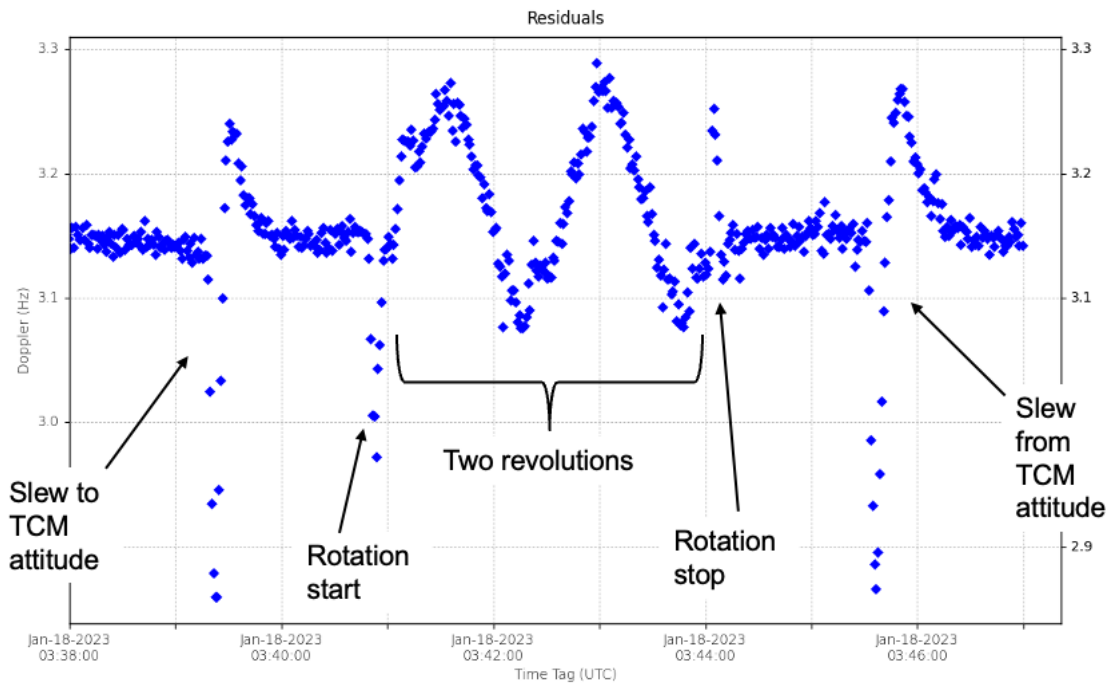


Figure 7: c069 Doppler residuals during the attitude-only test of a rotating TCM. In the Earth line-of-sight, 56.3 Hz is 1 m/s. The amplitude during the rotation period is ± 1.8 mm/s (± 0.1 Hz). The rotation period also has a 0.024 Hz bias due to the circularly polarized antennas used on LFL.

However, there was a limit to how late the thrusting could start and still achieve capture into the Earth-Moon 3-body region. For reasonable ΔV /day and total ΔV values, LFL needed to start thrusting reliably by late January 2023, and to continue to do so for 3 weeks. With this time constraint in mind, the LFL team developed a series of activities that incrementally but quickly built up to the full rotating TCM capability, as shown in Table 3.

The first rotating activity demonstrated the attitude maneuvers without firing the thrusters. LFL slewed to the target attitude, and then spun up to 4 deg/sec for 3 minutes. Afterwards, it slewed back to the nominal Sun-pointed attitude. During this time, the DSN generated two-way Doppler data, from which the Nav Team subtracted the motion due to a no-burn predicted trajectory. The resulting Doppler residuals (in Figure 7) clearly show these activities due to the offset of the antenna from the center-of-mass.

The following day, we added test setup and takedown burns, on Thruster 4, as shown in Figure 8. Since the rotation rate was still 4 deg/sec, these 10-second burns were not really the right length, but their timing was chosen to cancel out, and produce a total momentum change of zero.




Figure 8: c070 Doppler residuals during a test of 10-second momentum setup and take-down burns. Coast duration between burns is 35 seconds.

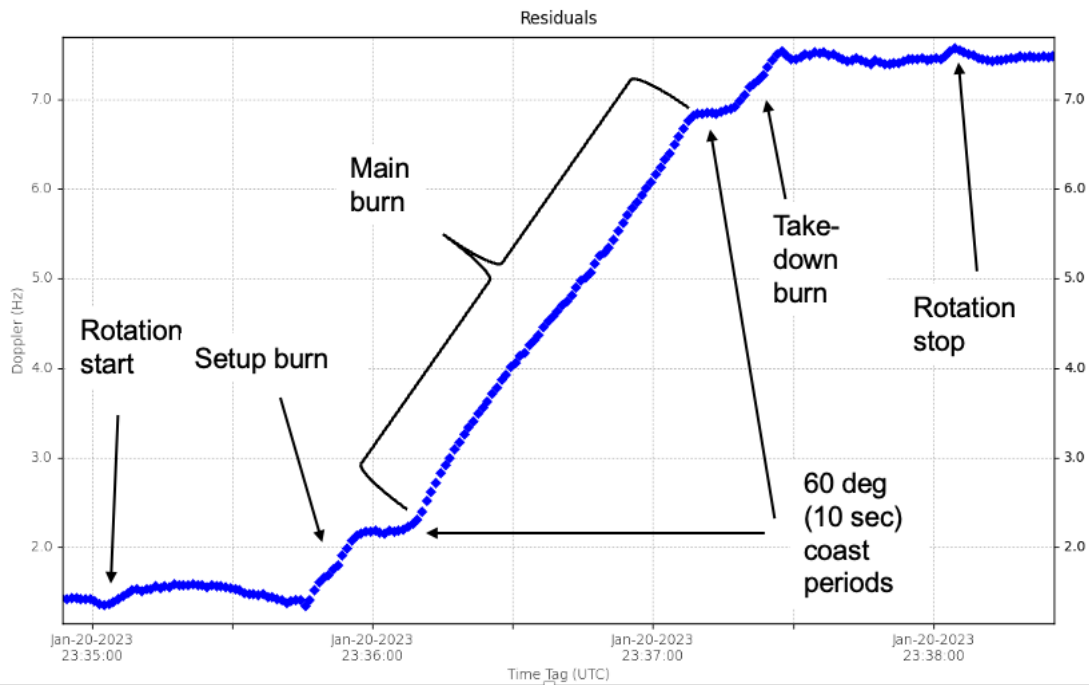


Figure 9: c072 Doppler residuals. Main burn is 60 seconds long. Setup and take-down burns are 10 seconds long.

One day later, the first main burn of 60 seconds was added, and the rotation rate was increased to 6 deg/sec. The Doppler residuals of this activity are shown in Figure 9. The corresponding LFL telemetry of the spacecraft body momentum in the X-Y plane is shown in Figure 10.

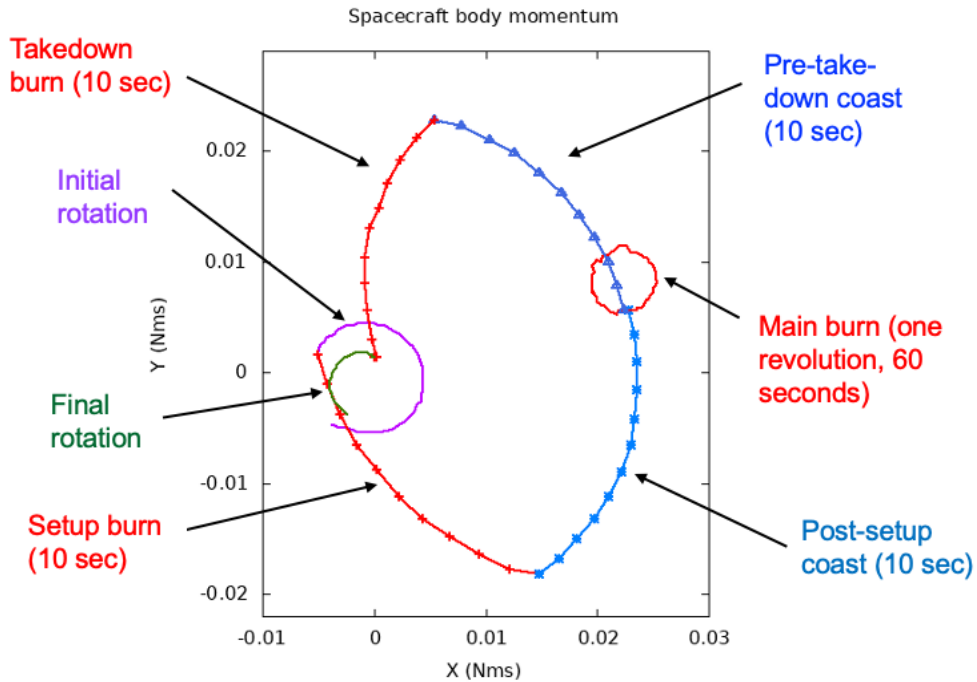


Figure 10: c072 X-Y spacecraft body momentum telemetry. Where shown, tic marks are every second. When starting from a non-zero X-Y momentum state, the size of the main burn circle is related to the initial momentum and the phase when the setup burn starts.

The next two rotating TCMs had main burns of 300 seconds, with increasing amounts of feed-forward torque applied. The command sequences used were refined to put more and more of the activities into one sequence, in preparation for autonomous TCMs outside of contacts. While this was successful, a troubling oscillation was seen in the spacecraft body X and Y attitude errors. This had also been seen in testbed runs, with the result that burns approaching the planned duration of 20 minutes significantly diverged. Thruster 4 showed a ~20% thrust increase in the first TCM due to self-heating, but this did not occur in the second one. The Doppler residuals and momentum circle telemetry for the first 300 second rotating TCM are shown in Figures 11 and 12.

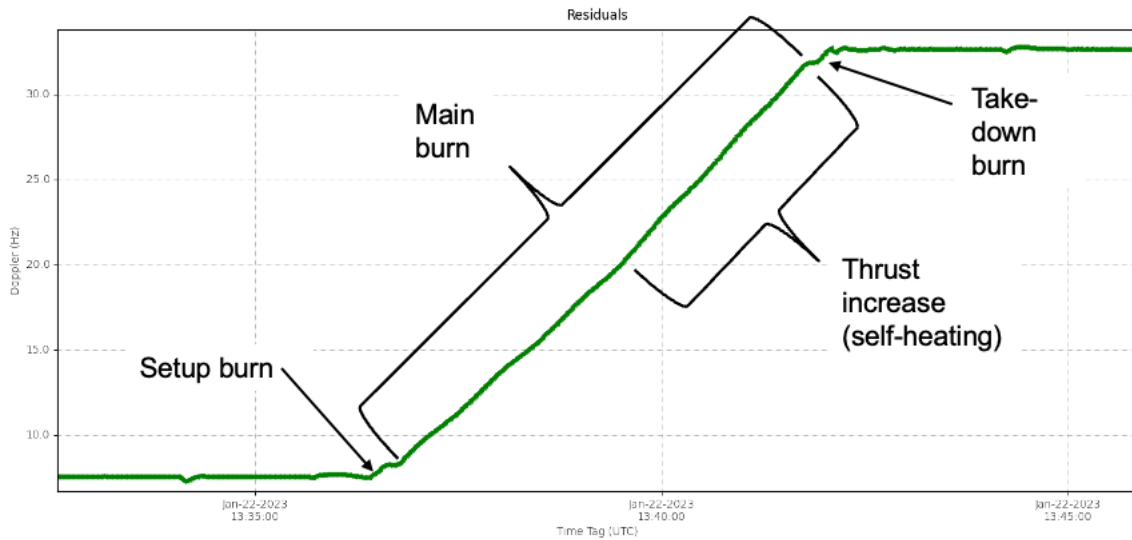


Figure 11: c075 Doppler residuals. Main burn duration was 300 seconds. Thrust increased at about the mid-point, due to self-heating on the thruster.

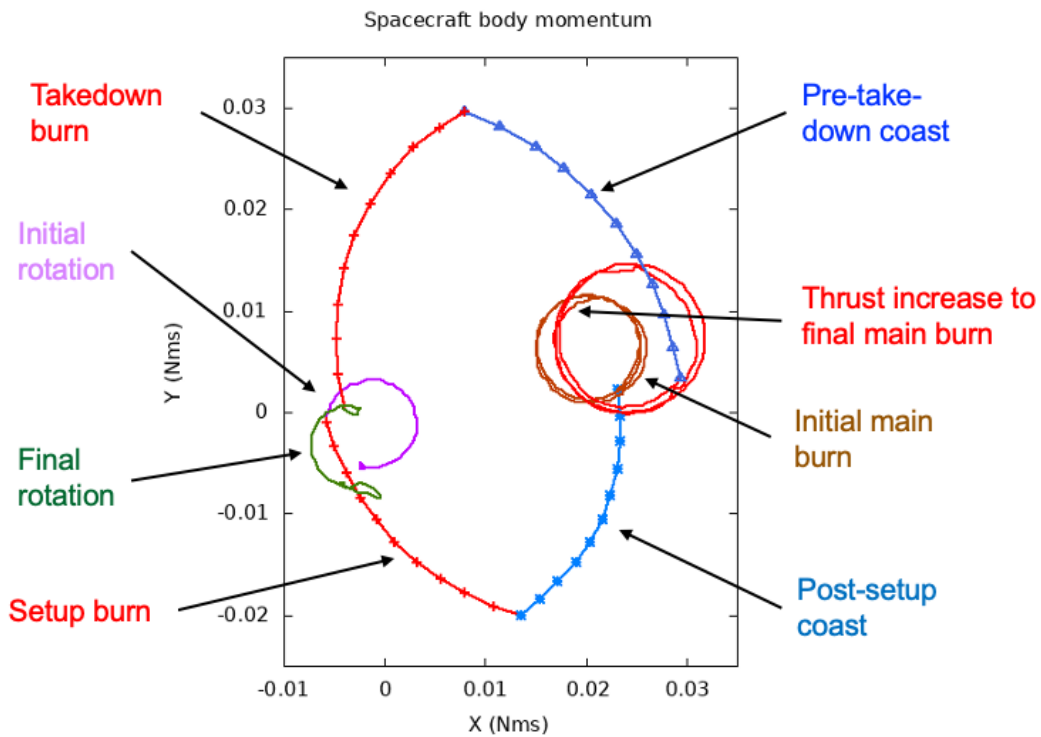


Figure 12: c075 X-Y spacecraft body momentum telemetry. Where shown, tic marks are every second. The thrust/torque increase halfway through the main burn changed the size of the main burn momentum circle.

Investigating and mitigating the attitude errors became the next priority. We contacted Blue Canyon Technologies (BCT), the provider of the XACT system, with our observations. Within only four days (including a weekend) they came back with an answer to what was happening, and solutions. It turns out that the controller had a limit on the amount of gyroscopic compensation torque it would allow in commanding the reaction wheels. The limit was much lower than what we needed during a rotating TCM, since they had not foreseen that we would need to continuously rotate the spacecraft when it had a large momentum state. (This is the torque due to $W * H$ as described above). Increasing this limit easily took care of the X and Y attitude errors. In addition, the integral term of their PID controller was also rather small. Increasing it to a medium or large value also allowed the controller to adjust to the required torque command level. These larger values caused spacecraft slews to take a little longer to converge to an acceptable attitude error but had the advantage that they covered a wide range of torque prediction errors.

With the gyro compensation and integral limits updated, we proceeded to longer main burns on Thruster 4. We had performed two 10-minute burns before this point, since the simulated attitude behavior was acceptable up to this duration. A 10-minute burn with the new XACT parameters behaved as expected. With negligible attitude control errors, we could now get a good estimate of the spacecraft dynamics during a burn. We observed a negative trend in the Z-axis momentum, which indicated that the true Thruster 4 thrust vector was not pointed exactly in the nominal direction in the spacecraft body axes.

If the spacecraft is not precisely rotating about a vector that is normal to the torque direction, then momentum will accumulate along the rotation axis, which is mostly in spacecraft Z. Thrust direction errors in the thrust-torque plane will cause this effect, while thrust direction errors normal to the thrust-torque plane will have very little effect (other than a tiny cosine loss), as shown in Figure 13. By estimating the rate of Z-axis momentum change, we calculated that the effective thrust direction was offset by about 6 milli-rad (mrad). Note that this is really an error between the thrust direction and the XACT-commanded rotation axis, which could be caused by a combination of thruster and XACT misalignments. In either case, the error magnitude was within expectations.

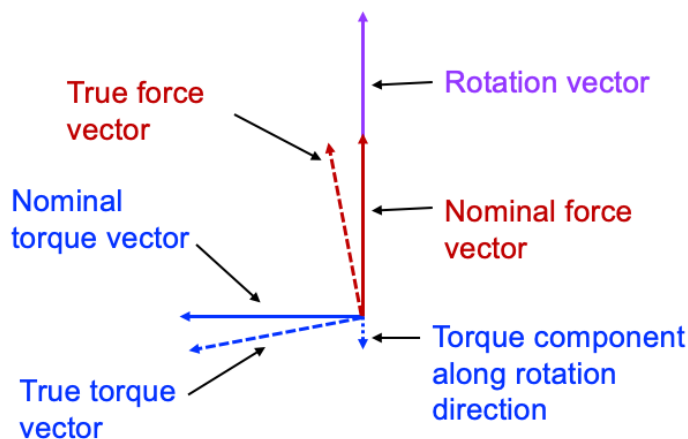


Figure 13: Spin axis correction geometry. Torque along rotation vector causes momentum build-up. Correcting the rotation vector to the true force vector corrects the momentum build-up.

entending the spacecraft such that the Thruster 4 torque vector opposed the inertial momentum. Then a series of Thruster 4 pulses were fired by ground command. After each set of pulses, the new momentum state was assessed to determine the next pulse train length to command. This allowed

the effective thrust direction was offset by about 6 milli-rad (mrad). Note that this is really an error between the thrust direction and the XACT-commanded rotation axis, which could be caused by a combination of thruster and XACT misalignments. In either case, the error magnitude was within expectations.

Going forward, the commanded spin axis was updated by the 6 mrad offset to neutralize the Z-axis momentum build-up. However, previous rotating TCMs had built up enough momentum that a single-thruster desat was needed. This was accomplished by ori-

us to reduce the momentum to below the magnitude of a single pulse, without needing a precise estimate of the impulse of each thruster pulse.

In addition to the rotation axis adjustment, we increased the duty cycle to 25% (firing 250 msec pulses) to compensate for an updated recommendation to use lower pump speeds. The maneuver duration was increased to 15 and then 20 minutes. The adjusted rotation axis successfully maintained a nearly-constant Z-axis momentum state, and at this point we had all the pieces in place to generate significant ΔV each day, including in off-contact maneuvers.

Unfortunately, the propulsion system had other ideas. During the 20-minute burn before the first off-contact maneuver, Thruster 4 thrust dropped precipitously (at about the halfway point). Evidently the problems with the other thrusters had finally affected Thruster 4. Further testing (with an attempted desat) showed that Thruster 4 was no longer generating any useful thrust.

With this setback, we made several changes in near-term activities and the trajectories we were trying to follow, which are summarized in the final section. However, Thruster 4 lasted long enough to allow us to fully demonstrate single-thruster rotating TCMs and single-thruster desats, thus fulfilling the main aim of this paper (as promised by the title). For more details of the activities of the propulsion system, and the efforts of the operations team, please see Smith & Cheek *et al* (reference 1) and Hauge *et al* (reference 2).

SUMMARY OF REMAINING ACTIVITIES

Since we were almost out of time margin to reach lunar orbit when Thruster 4 failed, it was clear that we could not develop the capability to use another thruster and still continue towards LOI. Fortunately, as the spacecraft was inbound towards the final low perigee, the ballistic trajectory crossed near the Moon's orbit while the Moon was only a few days away. Mission design found a trajectory that adjusted that timing to achieve a close flyby of the Moon on May 23 for only about 35 m/s, spread out over the next 2 months. After a controlled flyby, many trajectory options become available. We found a trajectory with close polar flybys at an average of every 1.5 months, which would still allow some opportunities for science, as shown in Figure 14.

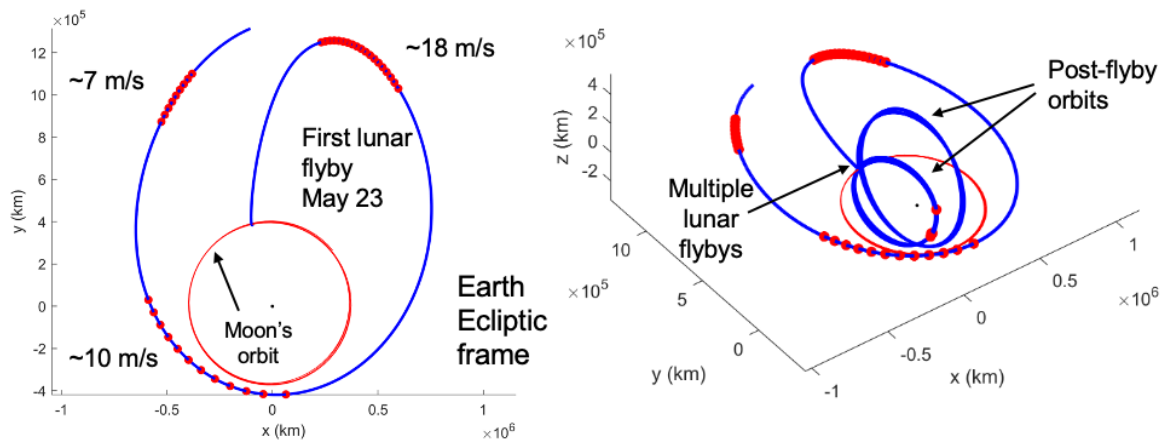


Figure 14: Flyby trajectory, which became the baseline after Thruster 4 failed. Starting point is beginning of February 2023. Lunar flyby is May 23, 2023. Red dots are daily 1 m/s ΔV s.

Unfortunately, the required ΔV was in roughly the opposite direction from the previous ΔV s, and completely off-Sun. Incremental tests were performed over a few days to understand how to

implement maneuvers in this direction, and the first short rotating TCM was performed on Thruster 3 only eleven days after the Thruster 4 failure.

During the next two weeks, Thruster 3 generated some ΔV , but eventually became unreliable and then stopped working entirely. The team started performing high-risk propulsion activities to try to recover some useful thrust, mainly involving running the pump backwards and then forwards to hopefully dislodge debris. This was somewhat successful, but the resulting thruster behavior was very unpredictable, which made operating a rotating TCM at the desired momentum state very difficult with sequences.

Consequently, we decided to develop a Ground-In-The-Loop (GITL) capability to command the propulsion system (thruster + pump) to manage the thrust unpredictability. This required measuring the delay between commands being sent and telemetry being received, to manage the phasing of the start of a pulse train with respect to the spacecraft body momentum state. While operator input was required, the timing of each command was managed by ground automation, closing the loop on the current telemetry state. This was much easier than requiring on-the-fly operator calculations, which were probably unworkable. The GITL tools were developed and worked as intended, but unfortunately the thrusters continued to trend towards zero thrust.³

During this time, an even lower ΔV trajectory was developed, that did not try for a lunar flyby in May 2023. With just a few m/s around perigee, the trajectory would linger in the Sun-Earth L1 region, instead of departing through that region to a heliocentric orbit. The leverage provided by the unstable L1 region allowed us to find a trajectory that returned to a lunar flyby on January 7, 2024, which would be followed by science flybys, as shown in Figure 15. However, this trajectory still required several m/s within a couple of days of perigee.

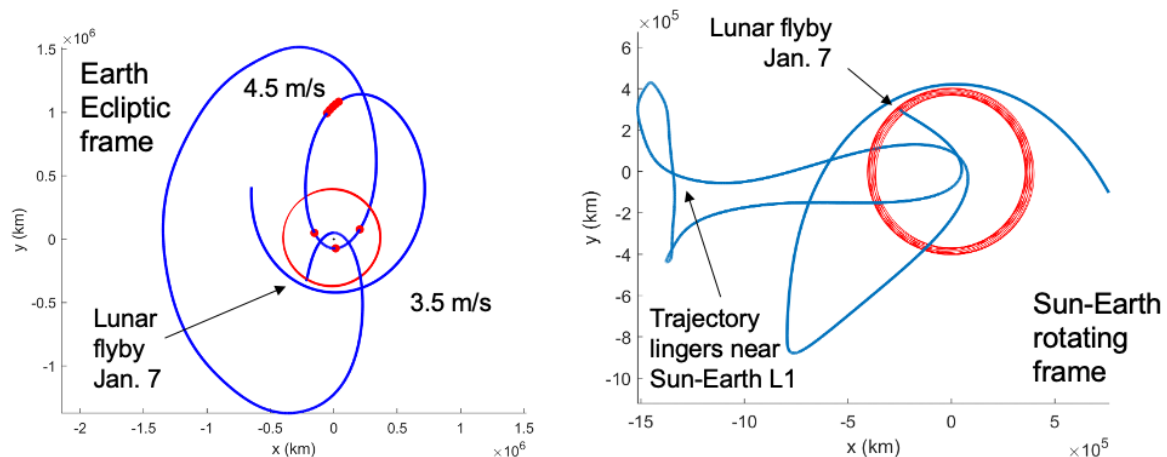


Figure 15: L1 trajectory, in inertial Earth-centered ecliptic and Sun-Earth rotating frames. Red dots are daily ΔV s of ~ 1 m/s. Achieving the required ΔV near perigee was critical.

The final propulsion activity was running the pump at very high speeds, to generate very high pressures, in the hopes that this would dislodge some debris before causing any other damage. The high pressures did show some positive thrust but did not reliably recover any of the thrusters. However, in the final rotating TCM attempt, the high pressure evidently ruptured a feed tube. After one further diagnostic propulsion system test, the project decided that we had done all we

could, and that reaching the Moon was impossible. Consequently, the LFL spacecraft departed the Earth-Moon system through the L1 region, after the last perigee, early (UTC) on May 17.

CONCLUSION

Despite severe propulsion anomalies, the Lunar Flashlight team successfully developed the means of maneuvering the spacecraft with only one thruster and the ever-faithful reaction wheels, within the constraints of the LFL system. Unexpected difficulties can catalyze remarkable creativity and perseverance – this was demonstrated by the LFL team in general, and particularly by the team of students operating LFL at the Georgia Institute of Technology.

ACKNOWLEDGMENTS

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004).

© 2023. California Institute of Technology. Government sponsorship acknowledged.

Every member of the LFL team contributed in some way towards the results reported herein. While not all are mentioned below by name, everyone’s efforts are still appreciated by the authors.

The LFL spacecraft was operated by GT graduate and undergraduate students, whose professionalism and excellence belied their educational status. Operations lead Mason Starr wrote much of the ground software, and diligently applied his previous operational experience to team procedures. Co-author Michael Hauge (who was still at GT during the period described in this paper) performed a general operations systems engineering role, while also specializing in attitude control and playing a critical role in the development of the rotating TCM. Shan Selvamurugan and Graham Jordan also made important contributions to attitude control operations, development, and testing. All of the graduate students functioned as both activity leads and flight directors, and included John Cancio, Marilyn Braojos Gutierrez, Conner Awald, and Robert Lammens, as well as those mentioned earlier. Undergraduate students (who performed the roles of command transmission (“ACE”) and data management) included Katherine Anderson, Shalomi Arulpragasam, Kiernan Barket, Emma Hanson, Mollie Johnson, Evan Leleux, Bryn Merrell, Micah Pledger, Katy Ryan, Cate Schlabach, Andrew Weatherly, and Alan Yeung. GT faculty members Dr. Glenn Lightsey and Dr. Jud Ready provided leadership and support to the operations team.

The LFL Mission Design and Navigation (MDNAV) team at JPL was led by Ted Sweetser in development and Stuart Demcak in operations. Andrew Cox performed the maneuver analysis, and also set up much of the LFL-specific navigation support software. Long-suffering orbit determination analysts fit through the plethora of large and small ΔV s and supported many midnight maneuver contacts. These included Stuart Demcak, Andrew Cox, Eric Graat, Mark Ryne, Jules Lee, and Julie Kangas, as well as the lead author. The mission design team, which produced a myriad of trajectories and trajectory products, included Gregory Lantoine, Jennie Johannesen, Julie Kangas, Stefano Campagnola, and Dan Grebow. Gregory Lantoine provided the plots used in Figures 6, 14, and 15 (along with many others), to assist in visualizing the new trajectories that he was able to develop, in response to LFL’s ever-changing prospects of achieving useful ΔV . The JPL Navigation Advisory Group (NAG) was at first led by the lead author, who was later kindly relieved in that role by Sumita Nandi.

Chris Burnside provided invaluable propulsion experience to LFL operations, along with several other MSFC colleagues. Chris was the primary interface to the propulsion system subcontractors and learned about JPL peanut-consumption rituals while supporting many contacts with propulsion activities. Additional propulsion advice came from Frank Picha and Matt Kowalkowski at JPL.

At JPL, the LFL project manager was John Baker, and his deputy was Philippe Adell, who also served as the Flight Systems Engineer. The Mission System Managers during operations were Christina Kneis and Jack Trinh, both of whom supported contacts at all hours with the rest of the team. DSN scheduling (and frequent re-scheduling) was provided by Ricky Cors.

The lead author is grateful for the comments and feedback on this paper that were received from the co-authors, as well as from colleagues Sumita Nandi and Shan Selvamurugan, and daughter Kerensa McElrath. Any remaining inadequacies of this paper are the responsibility of the lead author alone.

REFERENCES

¹ Smith, C., Cheek, N., Burnside, C., Baker, J., Adell, P., Picha, F., Kowalkowski, M., Lightsey, E. G. “The Journey of the Lunar Flashlight Propulsion System from Launch through End of Mission”, SSC23-VI-03, in *Proceedings of the Small Satellite Conference*, 2023.

² Hauge, M., Lightsey, E. G., Starr, M., Selvamurugan, S., Jordan, G., Cancio, J., Awald, C. “Operations Systems Engineering for the Lunar Flashlight Mission”, SSC23-WVII-02, in *Proceedings of the Small Satellite Conference*, 2023.

³ Starr, M. “Development of Tactical and Strategic Operations Software for NASA’s Lunar Flashlight Mission”, SSC23-XII-05, in *Proceedings of the Small Satellite Conference*, 2023.