LCROSS Flight Team Development and Operations Experience

Paul D. Tompkins*
Stinger Ghaffarian Technologies, Moffett Field, CA

Rusty Hunt†, John L. Bresina‡, Ken Galal§, Mark Shirley¶, Matt D'Ortenzio**
NASA Ames Research Center, Moffett Field, CA

James Munger††
Northrop Grumman Aerospace Systems, One Space Park, Redondo Beach, CA

Scott Sawyer‡‡
Orbital Network Engineering, Cupertino, CA

I. Introduction

LCROSS developed its flight team under particularly challenging conditions: the project had a small budget and a short development schedule, a unique and difficult mission and, owing to a 7-year gap since the last time NASA Ames Research Center (ARC) had led a mission, little available operations experience or infrastructure on which to base its flight. LCROSS assembled a diverse staff with varying degrees of operational experience, and transformed it into a capable operations team that performed well in nominal and emergency conditions. The remainder of Section I summarizes the conditions under which LCROSS team development began.

A. LCROSS Project

LCROSS was conceived as an economical approach to determining the nature of hydrogen observed in permanently shadowed craters at the lunar poles. LCROSS was selected in April 2006, under the Lunar Precursor Robotic Program (LPRP) of NASA's Exploration Systems Mission Directorate (ESMD), as a secondary payload to be launched with the Lunar Reconnaissance Orbiter (LRO) [1]. Under the earliest possible launch date of October 28, 2008, LCROSS was obligated to meet a 29 month development schedule from Authority-to-Proceed through launch readiness§§. The LCROSS project was constrained to a $79 million cost cap. The project's tight budget and fast-paced schedule provided a strong impetus for utilizing a small staff in creative ways to accomplish the huge body of work required to prepare for flight while managing risk.

To ease the budget and schedule challenges, NASA designated LCROSS as "Class D" under its risk tolerance scale. This classification allowed the project to accept greater risk relative to higher-profile, big-budget missions, and shaped the development strategy and operational approach. However, the definition of Class D was not well-documented, nor did it have a strong precedence in prior NASA missions. Hence, the LCROSS project team had to work out how the classification translated into actual programmatic, engineering and operational practice, in coordination with NASA ARC, the program office and NASA headquarters. Importantly, it had to balance the need to reduce cost and save time against the natural tendency to pursue the reliability standards of bigger-budget missions.

---

* LCROSS Flight Team Lead, Code TI, Intelligent Systems Division
† LCROSS Ground Data Systems Lead, Code TI, Intelligent Systems Division
‡ LCROSS Activity Planning & Sequencing Lead, Code TI, Intelligent Systems Division
§ LCROSS Mission & Maneuver Planning Lead, Code PX, Project Management Division
¶ LCROSS Payload Engineer, Code TI, Intelligent Systems Division
** LCROSS Flight Controller, Code TI, Intelligent Systems Division
†† LCROSS NG Mission Operations IPT Lead, System Definition Department
‡‡ LCROSS Mission Operations Test Conductor
§§ Though the LRO/LCROSS launch date ultimately slipped to June 18, 2009, LCROSS met the original deadline for flight readiness in October 2008.

B. The Mission

The LCROSS mission was designed to use a Centaur upper stage as a kinetic impactor to loft surface material from within a permanently shadowed lunar polar crater into sunlight. Meanwhile, a Shepherding Spacecraft (S-S/C) would observe the impact ejecta for signs of water and other volatile compounds using a combination of cameras, spectrometers, and a photometer. The S-S/C would remain attached to the Centaur for most of the mission, and perform all orbit maneuvers to target the Centaur for impact. Though, on a typical transfer orbit, the moon is only days from Earth, LCROSS could not achieve the required impact angle of incidence and other geometric constraints using a direct Earth-moon trajectory. Instead, the mission planned to perform a lunar gravity assist maneuver early in flight to throw the attached S-S/C and Centaur into an orbit about Earth, with a period of approximately 36 days and highly inclined with respect to the moon's orbit. The orientation of this phasing orbit would enable LCROSS to achieve a much higher impact velocity and a steeper impact angle into the polar target, thereby creating an ejecta plume with a greater volume and altitude. The impact was planned to occur after three revolutions about the Earth. On the final day of the mission, the S-S/C would separate from the Centaur. At Centaur impact, the relative position of the observing S-S/C would enable just four minutes of science data collection before the S-S/C would also impact the surface.

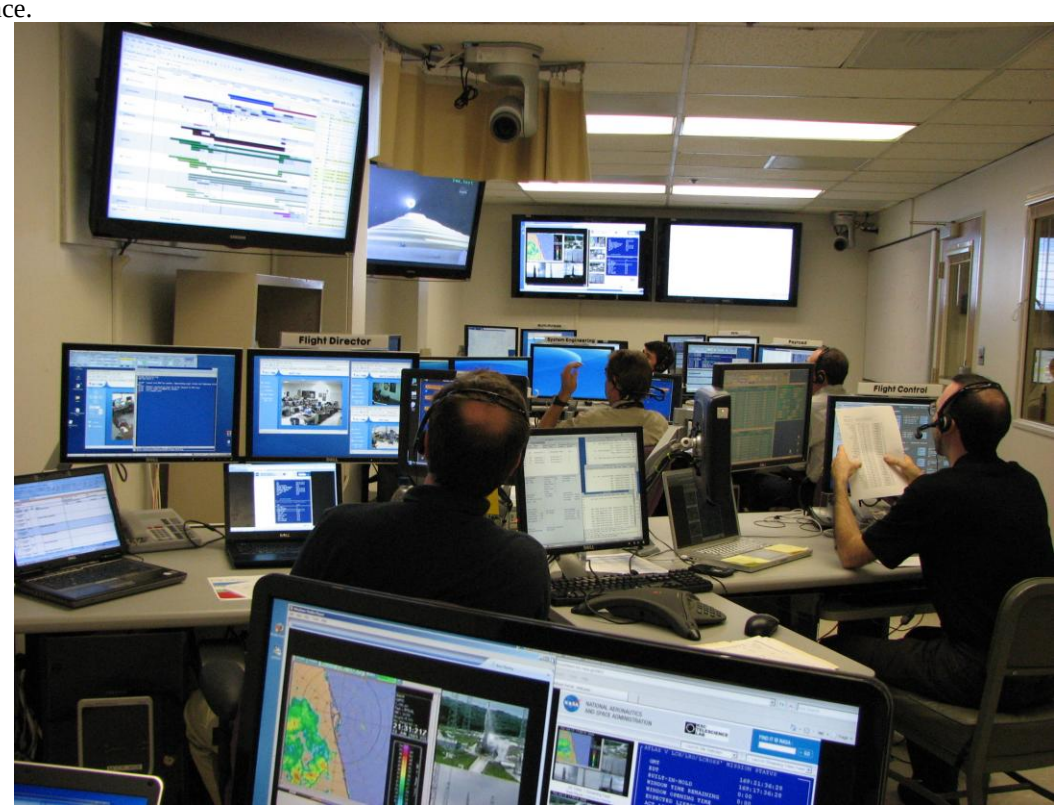

Fig. 1: Flight Team in the Mission Operations Control Room (MOCR). *The MOCR seated seven operators and was the focal point for execution and monitoring operations during the mission.*

The non-repetitive character of the LCROSS mission strongly influenced the makeup of the flight team and the design of operational practices. It was the first to operate with a Centaur upper stage far beyond launch. It comprised many unique event types (e.g. several types of science calibrations, lunar flyby, Centaur separation, lunar impact) and required extensive in-flight re-planning. It involved two periods of intense activity (launch week, lunar impact), separated by over three months of relative calm. LCROSS's success hinged on a single event – impact - at the very end of the mission; many other events were also time-critical. Class D programs take greater risks in spacecraft development (e.g. streamlined testing), but undetected problems often manifest themselves in flight, favoring an adept flight team or a spacecraft with a foolproof "safe mode" and plenty of time to analyze and recover from problems. Portions of the LCROSS timeline were not tolerant to such delays. Despite these challenges, the LCROSS mission was short, with no option for an extended mission, and therefore more forgiving in terms of flight team endurance.

C. Initial State of Mission Operations

The LCROSS mission operations campaign was the first NASA ARC had led since the Lunar Prospector (LP) mission, which ended in 1999 [2]. In the intervening years, ARC supported Space Shuttle payload operations and, concurrent with LCROSS, was preparing to support science operations for the Kepler mission [3]. However, neither was directly relevant to LCROSS. At its inception, LCROSS was provided with rooms and networks, but the majority of the ground system was undefined. Another consequence of the gap in flight missions was the small number of experienced operations staff available to LCROSS. Barring hiring operators from outside ARC, many team members had to be selected from other arenas of ARC business and trained to be flight operators. Also, there was no pervasive operations culture upon which to found LCROSS operational practices.

With this introduction, Section II describes how team development advanced in parallel with GDS and spacecraft development, and Sections III and IV cover how we composed our team and developed our operational practices. Training, both formal and opportunistic, is covered in Sections V and VI. Section VII provides a brief overview of LCROSS flight experience and team performance, from the human perspective, as a complement to descriptions of our experience with the spacecraft [4][5][6]. We point to our most valued lessons and consider the value of our approach for future missions in Section VIII.

II. Team Development Process

With little operational culture to draw from, the Mission Operations System (MOS) team had to define everything about how it would operate: its composition, its facilities, its general operational practices (e.g. voice loop protocols, Deep Space Network interaction, telemetry data archival, anomaly resolution processes) and LCROSS-specific practices (e.g. mission plan, team roles and responsibilities, team interactions and data transfers, command product generation and verification, procedures, flight rules, etc). It also needed to invent and execute a team training plan that would prepare its operators for flight.

Upon the selection of LCROSS, NASA ARC initiated a mission operations facilities and ground systems restoration effort. Hence, the MOS, GDS, and spacecraft developments proceeded in parallel. The MOS was responsible for defining requirements for generic and mission-specific tools and operations facilities. The implementation of mission-generic elements was handled under project-external ARC funding.

LCROSS adopted an iterative, “spiral” approach to both MOS and GDS development, recognizing that requirements would need to be refined through repeated, gradually more realistic testing. Training occurred on a gradually-developing ground system target, sometimes with serious limitations in capability, until late in the development schedule when all elements were in place. MOS procedures were refined gradually through collaboration and repeated simulations. LCROSS borrowed relevant, previously-successful operational practices wherever possible, as long as they fit within the lean LCROSS model. Given the team’s experience mix, it took specific inspiration from Lunar Prospector, the Mars Exploration Rover (MER) mission, Space Shuttle and International Space Station (ISS) operations, commercial communications satellite operations, prior Northrop Grumman (NG) missions, and even from ARC wind tunnel operational practices.

III. Composing the Team

Team composition was driven by two factors: the mission timeline (see section I.I.B) and the distributed expertise of the LCROSS team. Given the unique, non-routine character of the mission, the MOS had to cover a full range of operational disciplines. It needed to be lean to save costs during long-term low-intensity operations, yet be able to scale up for brief periods to support time- and mission-critical, high-intensity periods. The short mission meant that LCROSS could realistically maintain a single, lean team for the duration. To independently cover most operations, LCROSS composed a small core MOS team at NASA ARC, supplemented with maneuver design and navigation expertise from GSFC and JPL, and two systems engineers from NG. For the most intensive periods and for anomalies, an extended team of engineers from NG provided instant subsystem depth, reducing the risk of missing a critical event.

Importantly, under Class D, the LCROSS project recognized that it could not necessarily afford a comfortable staffing level for either development or flight. Staffing during development was enough to cover base work, but was often too small to comfortably address setbacks or unexpected expansions of project scope. The project recognized that during flight, the core MOS team would be barely large enough to address nominal operations (and during some phases, only through 13-hour shifts). It had very few backup operators should one of the prime operators become suddenly unavailable, and could not be expected to address anomalies without help from the extended team from NG, and at some risk to ongoing nominal mission events running in parallel. To enable team scale-ups during the

low-intensity Cruise Phase, the LCROSS project maintained a healthy operational reserve budget that it could use to support the full NG team for anomaly investigations.

A. The Core Team at ARC

A core team of ARC employees filled many of the primary MOS roles (see Fig. 2): planning and command generation (Scientists, Mission & Maneuver Design Engineers, a Link Analyst, a Command Sequencing Engineer, Engineering Analysts, and Simulator Engineers), event execution (Flight Directors, Flight Controllers, Systems Engineers, a Payload Engineer, and a Telemetry Data Engineer), and management (Mission Operations Manager).

Budget limitations dictated a small and efficient MOS staff. In most cases the head count was barely enough to cover the shift schedule (1-2 people per role), and did not allow for backup operators. Given a particular mix of skill and experience, some staff members were assigned to more than one role. Furthermore, LCROSS could not afford the luxury of employing independent MOS and GDS development teams - many operators on the team began work on LCROSS as GDS developers and later inherited additional responsibilities as mission operators.

NASA ARC flight team members came from a diverse set of backgrounds. Some had served in other mission operations roles, including in support of Lunar Prospector, Gravity Probe B, and other orbital science missions; the Mars Exploration Rover mission; Space Shuttle and International Space Station operations; and commercial satellite missions. Others had been employed as engineers in support of space and defense-related engineering projects. A number of staff had once led or supported wind tunnel operations at NASA ARC, and many were drawn from research and software development at ARC, in the fields of autonomous systems, artificial intelligence and robotics. Regardless of background, all team members were highly technically competent in their respective fields, and all were very excited and motivated by the prospect of supporting operations for a lunar impact mission.

B. External Team Members

LCROSS partnered with other organizations to bolster ARC experience in key areas, to create a reserve capacity for key mission phases, and to support spacecraft anomalies, should they occur. Due to the particular LCROSS challenge of precise lunar impact targeting, ARC augmented its own expertise in trajectory design and maneuver planning by partnering with the Jet Propulsion Laboratory (JPL) (small forces modeling and precision orbit estimation expertise) and NASA Goddard Space Flight Center (GSFC) (complementary trajectory design and maneuver planning capabilities). JPL also provided a small team for Deep Space Network (DSN) scheduling.

The MOS needed to possess deep spacecraft systems and subsystems expertise to assess spacecraft performance throughout the mission, and to diagnose and remedy anomalous spacecraft behavior, as necessary. LCROSS employed only two ARC spacecraft engineers - the Project Systems Engineer and his deputy - to oversee spacecraft development at NG. Both were ultimately assigned as Systems Engineers on the flight team. To create a deeper engineering team on short notice, a natural approach was to assign NG spacecraft engineers (leading design, integration and test) as part-time operators. It was unlikely that LCROSS could build a team with similar depth of knowledge in such a short time, and many at NG were excited at the prospect of participating in operations. Two NG sectors were involved in LCROSS development and later in the MOS (again, see Fig. 2): Aerospace Systems (NGAS) in Redondo Beach, CA, and Technical Services (NGTS) in Lanham, MD. Despite the clear benefits, this approach presented a number of additional challenges:

1. LCROSS would have to devote budgetary resources to train these engineers (through classes, simulations and rehearsals, etc), many of whom were not experienced in operations.
2. Because these engineers also served in leading roles in LCROSS spacecraft development, they were unlikely to be able to support all operational training activities (concurrent with spacecraft testing).
3. The spacecraft engineering team needed facilities and equipment to allow them to participate in real time (via telemetry monitoring). To co-locate them at ARC would entail a substantial increase in operational floor space, an enormous travel budget, and a significant travel burden on each of the participants (with high attrition likely). A distributed solution would require the build-up of dedicated remote operations rooms at each of the NG facilities, and would introduce the difficulties of distributed team coordination during rehearsals and flight.
4. The LCROSS project could not afford to employ the full set of spacecraft engineers full-time for the entire flight phase. Part-time involvement required that engineers work on other projects. During flight, without management support and careful staff planning, these engineers were at risk of being fully claimed by other NG projects.

C. Team Collaboration during Development and Test Phases

A key to MOS success was its tight integration and culture of open communications across organizations. Small team size enhanced the MOS's ability to communicate and to remain focused and coordinated. All organizations demonstrated full commitment to the LCROSS project, and to the streamlined approach required under Class D. NG

upper management allowed its LCROSS team far greater flexibility and autonomy to make decisions as compared to typical projects. Of critical importance, management at all participating organizations successfully established an atmosphere of collaboration, without barriers. Inter-organizational relationships were not strictly bounded by contract limitations. For example, NG openly voiced problems encountered during spacecraft development, and it involved ARC at every stage in working through these obstacles. As full-fledged MOS team members, NG engineers regularly contributed to internal ARC procedure reviews and provided valuable advice in support of spacecraft operational procedure development at ARC. By participating in procedure development and training exercises, NG team members could also better understand the operational effects of spacecraft design decisions. As described in Section VI, ARC MOS engineers were also included in key aspects of spacecraft development and test at NG, further solidifying this collaborative bond. Finally, tight cross-organizational team integration allowed external participants to communicate effectively. For example, the JPL Navigation team worked closely with NG attitude control engineers to develop models of the perturbation effects of thruster-based control modes on the trajectory. Through this body of interactions, and because of the tone set by management early in the project, the MOS (and the project at large) became a cohesive unit concerned principally with mission success, despite political and organizational boundaries.

In balance, a distributed MOS presented challenges. Though the team was well-integrated across organizational lines, communications breakdowns sometimes occurred due to physical separation. With the bulk of the team at ARC, informal discussions there occasionally unintentionally excluded external partners. Even full-team teleconferences suffered from poor audio quality and acoustics or occasionally poor network or graphics-sharing software performance. We recognized the importance of effective distributed team coordination early in the MOS development and reflected that both in MOS operational processes and in the GDS design (see section.IV.B).

IV. Development of Operational Practices

LCROSS operational practices were built on standard models of space operations, tailored to the LCROSS mission plan. The most fundamental of these practices are expressed in the choice of team organizational roles, described in the previous section. How these team members operated together is the subject of this section. The full set of practices is too large to describe in detail here, but the following sections highlight some of the salient features of LCROSS practice.

A. Workflow and Shift Scheduling

LCROSS operated under a cyclical workflow model with four phases: Planning, Command Generation and Verification, Execution, and Assessment. The first two phases were particularly important for LCROSS, in which plans and command sequences depended heavily on the outcome of previous events. Depending on the operation, a full cycle lasted anywhere from 7 to 48 hours:

1) Planning: During this phase, the Maneuver Design team designed trajectory correction maneuvers and spacecraft attitude changes to meet the requirements of impact targeting, science instrument calibrations, communications, and other activities, while minimizing propellant consumption and satisfying operational constraints. Systems Engineers designated housekeeping activities to be executed, and the Science Team and Payload Engineers designed payload observation sequences, as applicable. Link Schedulers coordinated DSN contact periods for LCROSS, and Link Analysts predicted link performance for future contacts based on past experience. Planning phase culminated in an Activity Selection Review, during which the team reviewed maneuver designs and selected and ordered the supporting activities to be undertaken during a future contact. The resulting fully-ordered list of activities and their associated parameters was called an Activity Plan. The Maneuver Design Engineers operated in multiple overlapping shifts during peak activity periods; all others supporting this phase operated in a single shift.
2) Command Generation and Verification: During this phase, the Command Sequencing Engineer converted the Activity Plan into a Command Plan, combining onboard command sequences and ground-based commanding procedures. This process was partially automated to minimize human error, but was sufficiently flexible to enable in-flight modifications. The Engineering Analyst and Simulation Engineer verified the correctness and safety of the command products using a combination of analysis (e.g. rudimentary automated flight rule checking) and simulation. This phase concluded with the Command Approval Meeting, during which the team reviewed all Command Plan elements and associated analysis and simulation results before providing final approval to proceed. The minimal team supporting this phase comprised a single shift of activity, immediately following Planning Phase.

3) Execution: In this phase, the MOS enacted the Command Plan. All execution was conducted by the team in near real-time (with speed-of-light delays of seconds). The Flight Controller coordinated DSN ground stations and performed all commanding. Systems and Payload Engineers monitored spacecraft and payload telemetry, and made real-time recommendations. The Telemetry Data Engineer coordinated, packaged, and archived telemetry. A Flight Director orchestrated the team, often via procedure walkthrough, and was the lead authority for nominal operations. During 24-hour operations, the Execution team was divided into two shifts, A and B, each staffed with operators for every core role. Often one shift actively monitored the spacecraft during Planning and Command Generation phases, and the other performed Execution. Single Execution shifts conducted the operations for isolated DSN passes. Science team operators augmented the standard Execution team for events focused on science data collection.
4) Assessment: This phase ran in parallel with Execution phase and persisted afterwards while the team estimated the trajectory, characterized orbit perturbations and evaluated burn performance (Navigation), analyzed engineering telemetry (spacecraft and payload engineers), and analyzed science data (Science Team) to infer spacecraft health status and to determine the degree of success for maneuvers and science activities. This information was fed back into the next Planning phase, and to status briefings for project stakeholders.

Day-to-day operations were led by the Flight Director on duty. The Mission Operations Manager ensured that the flight team adhered to accepted operational practices, helped coordinate anomaly responses, and was the primary communicator between the flight team and project management and LCROSS stakeholders.

B. Facilities, Physical Distribution, and Team Communications

The organization of LCROSS facilities was strongly influenced by the workflow design and team distribution. Operations were centered at ARC, and utilized three primary rooms, all of which were built up and equipped in stages during LCROSS development.

The Mission Operations Control Room (MOCR; see Fig. 1) hosted the core MOS team that lead Execution Phase and supported Assessment Phase. The MOCR was a self-contained facility that accommodated seven operators with telemetry and command workstations and primary and backup telemetry and commanding data systems.

The Science Operations Center (SOC) was the room from which the Science Team evaluated science instrument data streaming in real-time from the spacecraft during science events, coordinated the ground telescope observation campaign at lunar impact, and performed some post-event science data analysis. The SOC was physically separate from the MOCR to allow science and engineering-oriented activities to proceed in parallel without disturbance to either team. However, the Payload Engineer acted as a representative of the Science Team within the MOCR for engineering-related concerns.

The Mission Support Room (MSR) was the venue for the Planning and Command Generation and Verification Phases, and portions of Assessment Phase. It was the center for off-line activities: maneuver planning, development and test of command products, and discussion of engineering issues (including anomaly resolution). The MSR was separated from the MOCR and SOC to allow discussions and coordination in parallel with real-time and science activities.

NG assembled and maintained two Remote Operations Centers (ROCs), one at the NGAS facility in Redondo Beach, CA, and another at NGTS in Lanham, MD. These supported two roles. First, the NG ROCs acted as an extension of the MOCR, providing real-time oversight of subsystem and system behavior during Execution phases. Second, they supported Assessment as engineering extensions of the MSR for more detailed review and discussion of spacecraft behavior (both nominal and anomalous), with easy access to NG design and test documentation and factory support. Most events were supported by at least one NG Systems Engineer. Critical events were also supported by subsystems engineers, up to the full set of disciplines, depending on need. Most of the subsystem engineering disciplines and the Systems Engineers operated from NGAS, while the subsystems engineers overseeing Avionics and Flight Software operated from NGTS.

The JPL Navigation and Link Scheduling teams operated from their own fully-equipped facilities, while GSFC Mission Design team partners supported the critical initial and final weeks of the mission from the MSR at ARC, but worked from their home facilities during the months of Cruise Phase.

Achieving effective communications and situational awareness over this distributed team was a major challenge. Much of the solution was standard: communication was over a multi-channel voice loop system that linked all operations consoles, launch site operators and the Deep Space Network. Off-loop communications were minimized to maintain situational awareness. All operators had access to real-time telemetry at their workstations, and all workstations also displayed the command log of the single workstation used by the Flight Controller to issue

commands. A video distribution system enabled operators to mirror their console graphics on overhead monitors distributed throughout the operations rooms. At least one of these often displayed a real-time animation of the LCROSS spacecraft attitude, driven by real-time telemetry, depicting sensor fields of view and the spacecraft trajectory in relation to the Earth and moon. A secure Mission Data Product Server acted as the central repository and transfer point for all file-based electronic data, including daily plans, command products, and logs.

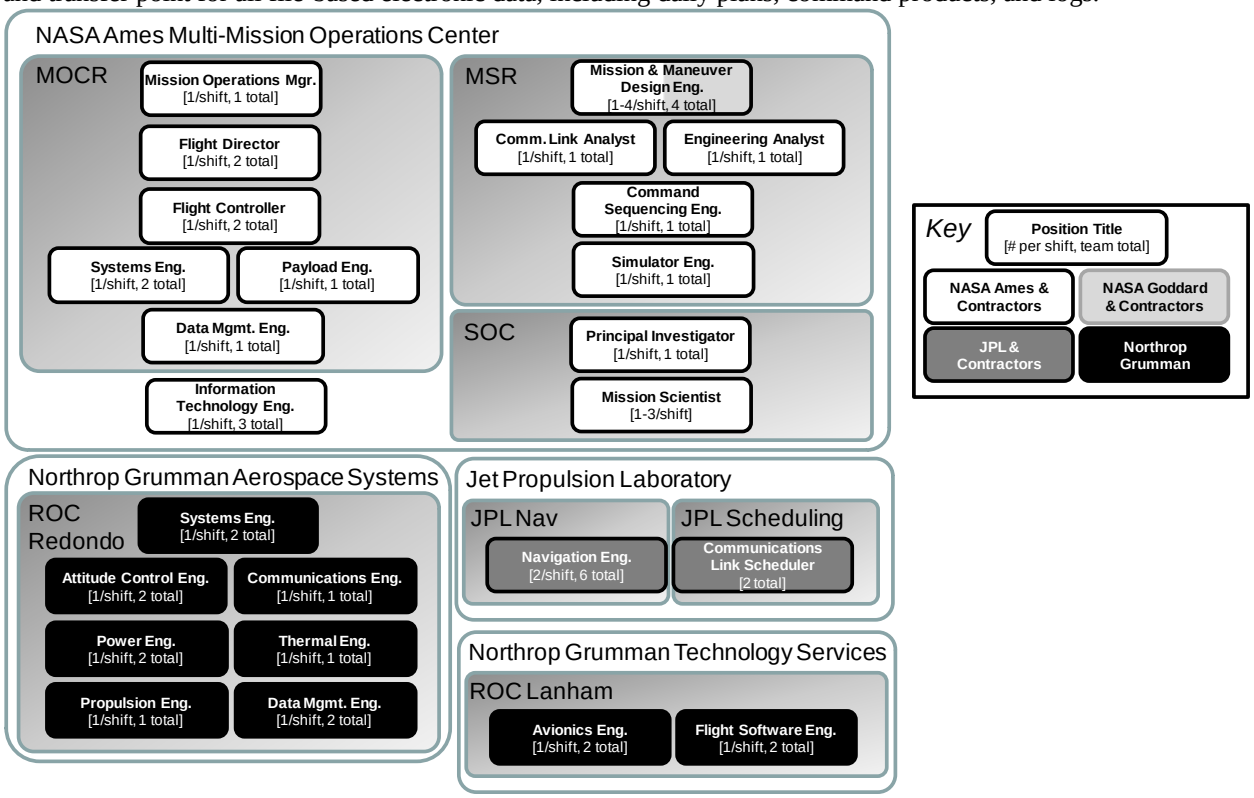

Fig. 2. LCROSS Flight Team and Physical Distribution. *Encompassing boxes represent the facilities from which team members typically supported operations.*

Early tests proved that these measures were insufficient to keep the distributed team informed and coordinated, and several additional tools were added. An event timeline containing DSN tracks, procedure timespans, payload data collection times, shift schedules and other key events was added to the overhead displays. The timeline display updated in real-time to highlight current activities and to give time from the present to past and future events. Another overhead monitor displayed the procedure being followed by the Flight Director and Flight Controller, highlighting the current step. Subsystem engineers had personal copies of all procedures, but the overhead copy always showed exactly where the team was, lessening the need to keep continuous track of this via the voice loop. This particularly helped keep remote subsystem engineers in sync, but it also allowed all engineers to disengage and re-engage with the procedure depending on their subsystem's relevance and their monitoring and analysis task load. Third, webcams cameras provided overviews of all rooms allowing anyone to see whether a distant team member was attending their console, in discussion or otherwise unavailable. For instance, the Payload Engineer kept the SOC webcam on his screen to see the PI and payload scientist as he spoke to them over the science voice loop.

C. Tight Integration of Science Team and Payload Engineers with MOS

LCROSS science operations were focused activities lasting up to one hour[***] (but separated by hours, days or weeks), during which the science team collected images and spectra, assessed instrument performance, and adjusted instrument parameters on the fly. These were some of the most complex operations of the mission. Early in development, MOS leadership recognized the importance of coordinating closely with the science and payload teams before and during flight. Without coordination, there was a significant risk that MOS would not meet the

[***] Thermal constraints dictated that the LCROSS payload could not be powered on for more than one hour at a time.

needs of the science team during these events and that the science team-derived operations concept would be inconsistent with the rest of operations.

Our effort to fully integrate the scientists and Payload Engineers into operations began through months of command sequence and procedure development covering all science events, and including both nominal and contingency scenarios. Cycles of tabletop discussions, procedure drafting, and simulation-based testing were an effective means of readying the team for these events. Joint discussions between spacecraft- and science-oriented operators were very informative, and team exercises with the newly-drafted procedures were quick to identify the strengths and shortcomings from the engineering, team communications, and timeline perspectives.

The team also jointly formulated an operational process for science activities. For Execution Phase, the Payload Engineer was assigned one of seven consoles in the MOCR and led the assessment of payload engineering performance during all payload activities. The SOC was equipped with science workstations networked to the MOS-wide telemetry data server, enabling scientists to view and assess science data seconds after capture on the spacecraft. Prior to execution, a science or payload representative became a required signatory for the approval of command products involving the payload. To keep real-time commanding authority consistent with other operations (approved by the Flight Director, and sent by the Flight Controller), Science Team members and the Payload Engineer made verbal command requests (from a pre-arranged set) to the Flight Director and Flight Controller over the voice loop. The Flight Director quickly evaluated the criticality of science requests relative to other immediate ground-commanded tasks; if approved, the Flight Director tasked the Flight Controller with sending the requested commands for radiation. This preserved a notion of centralized authority at the "big picture" level, albeit with some additional delay. In cases where timing was especially critical (e.g. real-time adjustment of camera exposure settings), the MOS devised a page of virtual buttons that enabled the Flight Controller to issue specific instrument commands via a mouse click.[†††]. As flight approached, science and payload team members were integral participants in all science-oriented MOS tests and rehearsals.

D. Nominal and Contingency Procedure Development

One of the dilemmas in implementing the Class D philosophy was in establishing the proper balance between preparations for nominal and off-nominal scenarios. Under the highly-constrained LCROSS budget and schedule, LCROSS management and stakeholders recognized that the MOS could not develop contingency procedures with the depth and breadth of coverage possible with greater funding and staffing levels. Yet under a more risk-tolerant spacecraft development program, with a largely single-string design and a streamlined test program, anomalies were perhaps more likely to occur.

The spacecraft design factored into the MOS strategy. In the absence of redundancy in many subsystems, LCROSS spacecraft engineers placed great emphasis in the design of a robust safe mode and fault detection algorithms. Rather than provide many optimized fault responses in an effort to preserve operations whenever practical, the spacecraft enacted relatively few responses to address broad classes of faults more simply. In response to a few faults, the spacecraft was designed to switch to redundant units, or simply to disable units and maintain operations under degraded performance. But for many serious fault conditions, the response was the same – to transition to a Survival State which could preserve spacecraft health and maintain communications with Earth indefinitely.

Through discussions with LCROSS advisors, the MOS adopted a strategy that strongly emphasized the development of nominal processes and procedures, and prepared off-nominal contingency responses as time and resources permitted. The rationale was that serious spacecraft fault conditions would result in transitions to Survival State. In many scenarios, with the spacecraft in this safe state, the flight team could simply delay nominally-scheduled events while it focused on specific solutions to the anomaly, and then re-engage the mission plan after the anomaly had been mitigated.

ARC led procedure and command sequence development, with input from NG and other partners. Where highly-important mission events could not be postponed, for example Lunar Swingby, Separation, and Lunar Impact, the MOS spent months developing contingency procedures and specialized command sequences to maximize the chances for success.

Just months prior to launch, the LCROSS project learned that once expended and decommissioned, the attached Centaur upper stage might exhibit a propellant leak behavior that could jeopardize the mission. Leaking propellants would induce a torque on LCROSS, which would be counteracted by the Shepherding Spacecraft attitude control system. Under the most severe (but short-lived) projected leak rates, LCROSS might expend a large fraction of its

[†††] These were instituted during the mission, as a result of lessons learned during the Lunar Swingby event. The commands were used extensively during the Impact event and were critical to mission success.

propellant reserve fighting this torque in the first week of the mission. However, under more moderate, prolonged leaks, LCROSS was projected to expend its entire propellant load, leaving nothing for other mission events. Under ARC orders, the MOS developed a comprehensive strategy to detect, measure, and minimize the effects of a Centaur leak. Though a contingency development effort, some form of the propellant leak was expected to occur. Over weeks of intense development, simulations, and rehearsals, the result was a specialized set of command sequences and procedures to measure external torques, reconfigure the attitude control system to achieve greater efficiency (at some performance loss), and to implement a passive, spin-stabilized control configuration (a significant modification of Survival State) that would drastically reduce propellant usage under the worst-case leak conditions.

E. Anomaly Resolution Process

Complementing the development of tactical procedures to address contingencies, the MOS developed a high-level plan for how the team would operate in response to anomaly discoveries. Significant anomalies might be met initially using the tactics codified in contingency procedures. But to determine the root causes of an anomaly, and to design workarounds and protections against future repetitions, the MOS recognized that it would have to strategically reconfigure to allow some of its members to dedicate themselves to anomaly resolution, while keeping others focused on maintaining spacecraft health and performing nominal mission events.

The plan described notionally how the team would react to anomalies of varying severity and time-criticality, spanning from immediate actions directed by a Flight Director (with input from other operators), to more deliberative responses requiring input from the Project Manager, Mission Operations Manager, the mission Principal Investigator, or Systems Engineers, depending on the nature of the fault and its potential effect on mission objectives. Owing to the emphasis on preparing for nominal operations, and the intensity of the development schedule to complete other high-priority tasks, the plan did not include some important details, for example designating specific anomaly resolution leadership; specific team assignments to anomaly resolution versus spacecraft health management; the division of labor between various shifts; or the process for anomaly procedure and command sequence development and approval.

V. Operational Training and Test Program

The LCROSS MOS training program included a combination of conventional and less conventional training methods. This section describes all of the more conventional approaches we used to ready ourselves for flight. System-wide tests and rehearsals were not only the prime means of requirement validation, but also were the cornerstone of operational training, and provided a basis for informal operator certification through MOS peer review.

A. Ground Data System Training

Operators learned how to use GDS tools (telemetry and commanding software, voice loops, link analysis, etc) through a combination of formal classes taught by software vendors, classes taught by MOS team members to the broader team, and via informal on-the-job practice. Some GDS software tools were custom-designed and implemented by the end-user, obviating the need for training. Partial-team and system-wide tests presented many opportunities to practice GDS tool usage.

B. Reviews, Workshops, and Seminars

An important part of training was to familiarize team members with the spacecraft design. The MOS recognized the value of spacecraft design reviews as training opportunities. LCROSS encouraged MOS attendance at early reviews, and required attendance at detailed subsystem Critical Design Audits presented by NG.

To complement training via spacecraft design reviews, the MOS held a series of Spacecraft Subsystem Workshops: Attitude Control, Avionics, Flight Software, Power, Propulsion, Communications, Autonomy and Fault Management (each conducted by the NG lead subsystem designer), and Payload (conducted by the Principal Investigator and lead payload software engineer). Each workshop lasted from four to eight hours and reviewed subsystem design, but emphasized operation and upkeep, including the use of primary relevant commands and telemetry, typical on-orbit behavior, troubleshooting and operational constraints. All ARC MOS members were required to attend these workshops.

A series of Operational Focus Workshops, led by the Lead Flight Director, provided in-depth reviews of operations for specific mission events at the system level. They reviewed the sequences of events, DSN utilization, geometric considerations, command sequences, operational procedures and constraints, the workflow covering the events, timelines, and staffing. Workshops were held for Activation and Checkout; Cruise Phase Housekeeping;

Cruise Phase Trajectory Correction Maneuvers; Star Field Calibration and Lunar Swingby; Earth Look Calibrations; and Separation and Lunar Impact. All MOS team members were required to attend these workshops.

In conjunction with the Operational Focus Workshops, the MOS conducted many workshops to develop and review command sequences for the same events. These meetings gathered spacecraft subsystem and systems-level expertise, Flight Directors and Flight Controllers, and for science-related activities, members of the science and payload team. During meetings, team members communicated design requirements to the Activity Planning & Sequencing Lead, who implemented the sequences and, for parameterized activities, the tools used to generate sequences during flight.

C. MOS Operational Tests

With the concurrent development of the MOS, facilities, and GDS, MOS training was limited by the degree of availability of facilities and tools, as governed by the GDS development schedule. The LCROSS team recognized two things: first, that MOS training could not be delayed until the facilities and GDS were fully deployed; and second, that the GDS team could not confidently deploy a quality product without intermediate validation testing by the MOS team. Hence, the MOS coordinated with the GDS development team to create a schedule of interleaved GDS releases and MOS operational tests that served as both training exercises and GDS requirement validation tests (see Fig. 3). The GDS release schedule was designed to provide increasingly capable sets of hardware, software and command products that built logically upon previous releases. MOS tests grew in sophistication and realism, starting with tests of contiguous "threads" of the operational workflow, and growing to system-wide tests of the team and GDS to support the operations for full mission events. The contents of each GDS "build" was driven in part by the needs of the MOS tests that it would serve, for example software or command sequences to support a specific mission event. MOS tests doubled as GDS requirements validation tests. Results from MOS tests occasionally prompted GDS requirements updates to capture needs that were discovered under realistic simulations.

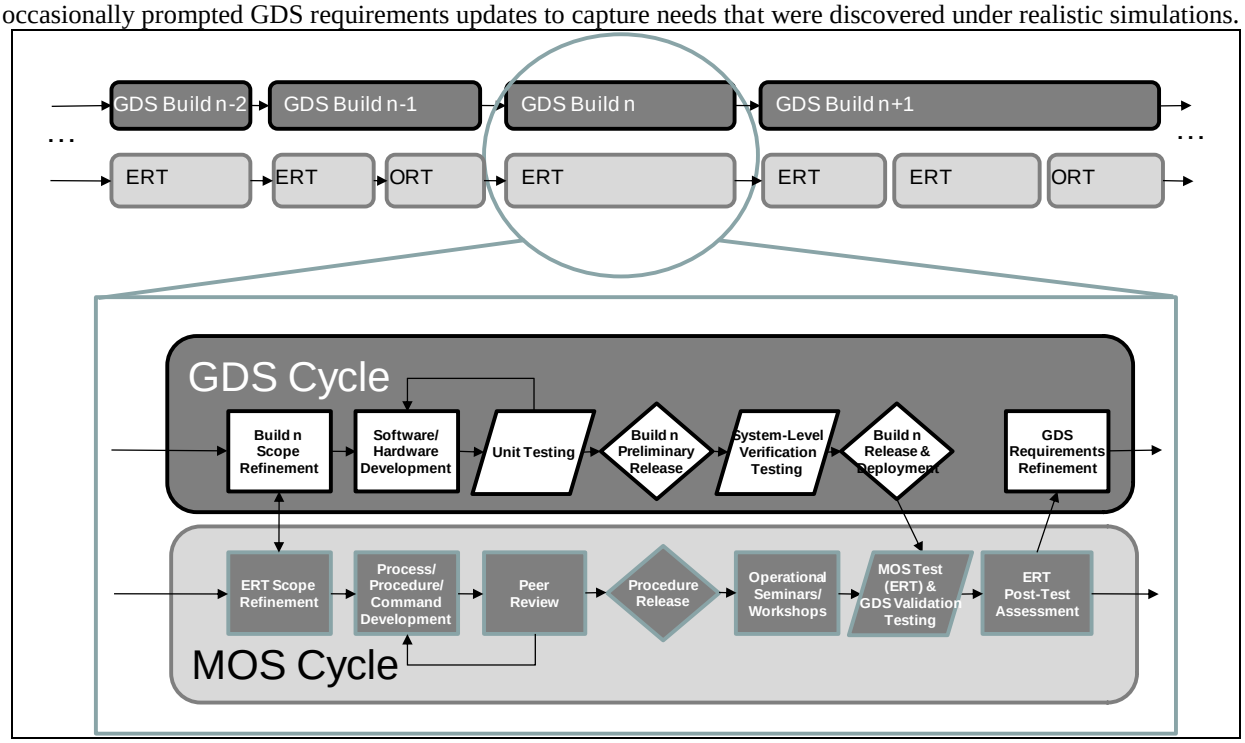

Fig. 3. Interleaved GDS and MOS Development Cycles to Support Parallel Development.

An LCROSS spacecraft simulator was a key GDS delivery in support of MOS tests. It combined a partial copy of LCROSS avionics hardware, a full copy of flight software, and a software-based dynamic simulator to simulate vehicle flight dynamics and the behavior of other spacecraft systems not represented in hardware (e.g. IRU, star tracker, power electronics and solar array, thruster modules). Importantly, the simulator was built partially from engineering test units used by NGTS to develop LCROSS flight software. Therefore, MOS simulation-based validation tests could not be conducted until the primary flight software delivery to NGAS. However, once transitioned to ARC, the LCROSS spacecraft simulator became a dedicated resource for GDS development (for

command products) and for MOS tests (during which it represented the spacecraft). The simulator was critical to the success of LCROSS.

Early MOS tests were designed and conducted by the Flight Team Lead, who also held an operational role as Lead Flight Director. As tests became more complex, the MOS recognized that it needed a dedicated Test Conductor to design test scenarios and to orchestrate tests behind the scenes. The Test Conductor was selected for his extensive prior experience with mission operations and spacecraft systems engineering. Once trained in the specifics of the LCROSS mission and simulator operation, the Test Conductor could lead system-wide tests and independently judge the performance of the MOS team.

MOS tests of operational "threads" were called Thread Tests (TTs), and began with tests of telemetry receipt (telemetry decoding, distribution and display), commanding, and data product routing and archival, to name a few. Later TTs tested more significant threads, for example planning and command generation for trajectory correction maneuvers. In addition to supporting GDS validation tests, TTs helped train operators in the use of GDS tools, in the use of realistic basis data to create and deliver mission data products, and also to accustom them to segments of the MOS workflow.

With the availability of all ARC operations rooms and sufficient GDS readiness, the MOS introduced system-level Engineering Readiness Tests (ERTs) that assembled a significant portion of the team and tested the complete workflow in support of specific mission events. The primary goal of ERTs was to prove that all GDS elements and MOS procedures were ready to support specific mission events. Team readiness was secondary. Hence, ERTs allowed intermediate stops and starts, timeline stretching and compression, and event re-ordering to focus on the most important aspects of the test. ERTs sometimes exercised contingency execution paths, but always with advance warning. The MOS conducted 14 ERTs of one to three days each.

Once an ERT had proven out the GDS and MOS procedures for a particular event, the MOS conducted an Operational Readiness Test (ORT) for the same event. ORTs were like ERTs, except that they adhered more strictly to the mission timeline (duration and ordering), and focused on validating MOS processes and team readiness over procedures and GDS tools. ORTs also tested contingency paths without warning the MOS team. Being able to correctly plan, execute and assess a mission event within the constraints of the mission timeline was a strong indication of operational readiness. The MOS conducted three ORTs, one for each of the three most critical periods.

D. Rehearsals

Rehearsals were conducted as final training events before launch and also during flight. Unlike ORTs, rehearsals typically exercised two or more days of continuous operation (active and inactive periods), mimicking a segment of the mission timeline to the minute. The Test Conductor ran rehearsals with the highest possible fidelity, and the team was expected to exercise every MOS process exactly how it would in flight. Rehearsals also exercised ancillary support activities like planning catered meals for the team during critical periods of 24-hour operations, and securing and using on-site lodging for any operator living beyond a maximum range to avoid the dangers of driving after long shifts. The MOS conducted a First Week Rehearsal prior to launch covering launch through Lunar Swingby (six days in duration), and three rehearsals during flight covering the final trajectory correction maneuver, Centaur separation and lunar impact (1-2 days in duration). A significant result from rehearsals is that team members grew to fully appreciate the rigors of the mission timeline, and discovered their own strengths and weaknesses in adapting to unusual sleep and waking hours while supporting critical events. Furthermore, in 24-hour operations, with shifts on opposing sleep schedules, team members came to understand the importance of good communications during shift handovers, the only times in the schedule when the full team could interact.

VI. Opportunistic Training and Staff Acquisition

In addition to the more conventional training approaches, many of the MOS team enhanced their training through participation in spacecraft, payload, and GDS development and testing (see Fig. 4). Many of these assignments were set up explicitly as training engagements, while in other cases, the MOS acquired staff to take full advantage of the training they had already received in their roles as GDS and spacecraft developers and testers. Active participation in the design, test, and review of the spacecraft and GDS was a far better training mechanism than classes and rehearsals alone. Dual responsibilities also had their share of disadvantages.

A. MOS Team in Spacecraft Development

The ARC Project Systems Engineer (PSE) and Deputy PSE were both assigned as lead Systems Engineers on the MOS. Furthermore, the MOS enlisted many NG spacecraft engineers to augment the core ARC team. The MOS benefitted from the inherent training each of these engineers acquired over more than two years of deep involvement

and leadership of LCROSS design, construction and test. Maintaining the engineering team through the entire flight saved the project significantly in training time and coordination effort. The MOS trained these engineers in operational practices far more easily than it could have trained operators in the intricacies of LCROSS systems design and test results.

B. MOS Team in LCROSS/LRO Testing

To build greater ARC-internal expertise in LCROSS spacecraft designs and operation, the LCROSS project negotiated with Northrop Grumman to embed two team members into the spacecraft test flow. One served on-site at NASA GSFC as a technical liaison between the LRO and LCROSS projects, focusing on common hardware development and test. He also served at NGTS as an avionics and flight software test engineer during avionics unit testing and integrated "flatsat" testing. The second served on-site at NGAS as a liaison and interface between NASA ARC and NG systems engineers for science payload integration and S/C integration and test (I&T). He also acted as an interface between different groups within the I&T team, splitting his time between NG and the launch processing facility for final spacecraft tests, and at NASA ARC during the conductance of S/C end-to-end testing.

As a result of their assignments, these operators became two of the most knowledgeable on the team in the detailed operation of the spacecraft. Both were assigned as primary Flight Controllers, the operators directly responsible for sending commands to, and acquiring telemetry from, the spacecraft. However, during their remote assignment over months during MOS development, the remainder of the team could not regularly call on their expertise or their support for overflow work. Furthermore, these volunteers spent 11 and 16 months, respectively, displaced from home.

C. MOS Team Members as Developers of Onboard and Ground Command Sequences

In another collaboration between the MOS and NG, the MOS Command Sequencing Engineer (the ARC operator in charge of implementing command sequences before and during flight) was assigned the responsibility of implementing the command sequences used by the spacecraft autonomy and fault management (A&FM) system for critical events, including initial spacecraft power-up and onboard fault responses. Working closely with NG systems engineers who designed the command sequences, this team member became expert in using the command and telemetry databases, onboard command sequence authoring and compilation, and the basic operation of onboard subsystems.

ARC was also responsible for designing and implementing all command product generation tools (e.g. parameterized templates, custom command generation scripts) and command products (onboard command sequences, ground command procedures, etc) to support the LCROSS mission. The same MOS member who developed the A&FM command sequences also led this task, with input from other MOS team members. While this was an enormous effort, there was no more effective way for the MOS to acquire such deep knowledge of spacecraft commanding, telemetry, and operation.

D. MOS Team in LCROSS Payload Development

The Payload Flight Software Lead for the science payload evolved naturally into the primary Payload Engineer for flight. In his development role, he designed and implemented all supplementary flight code for the operation of the payload, implemented all instrument command sequences for science activities, and also designed and implemented software used by the science team to analyze imagery and to assess payload throughput. He also supported payload testing. No other experience could have prepared this person so thoroughly for the position of Payload Engineer.

One of the science team members served as a Payload Test Engineer during development, and took a key role in science operations from the SOC during flight. Her detailed knowledge of payload test results benefitted her role as a mission scientist, and equipped her to help evaluate payload performance during science activities.

The only negative aspect of these arrangements was that when the timing of MOS test exercises conflicted with payload development and test activities, these two operators often had to prioritize their payload responsibilities.

E. MOS Team in GDS Development

Several MOS team members either oversaw GDS development or developed software tools for their respective MOS subsystems. For example, the GDS Development Lead, in charge of all LCROSS ground system development, became one of two Flight Directors. While larger teams might benefit from having dedicated software developers, the LCROSS MOS found distinct advantages in having its software end users also develop code. Most importantly, the problem of accurately communicating detailed requirements is obviated when the customer and the developer are one. This also goes for software training – the software developer is the most familiar with a tool's

capabilities, limitations and idiosyncrasies. This is especially important in flight, when deep software knowledge could enable an operator to work around a bug that might otherwise interfere with the support of an important mission event. Other MOS team members served as GDS team members, but in capacities distinct from their operational roles, with little training advantage.

Despite these advantages, having dually-tasked operators had its share of disadvantages. For one, when developer and operator are distinct individuals, the developer can perform unit-level testing, then pass the software to the operator to perform independent requirements verification testing. When one person serves in both roles, another operator, often less well-trained or even from a different discipline, must perform independent verification testing. Furthermore, peak pre-launch workloads often volleyed between the GDS and MOS teams according to the interleaved GDS and MOS test cycles. For those working on both teams, the workload was extremely difficult and afforded them essentially no rest time. Also, because most of the MOS and GDS work was concurrent (e.g. software development and MOS rehearsal preparations), one task or the other often suffered.

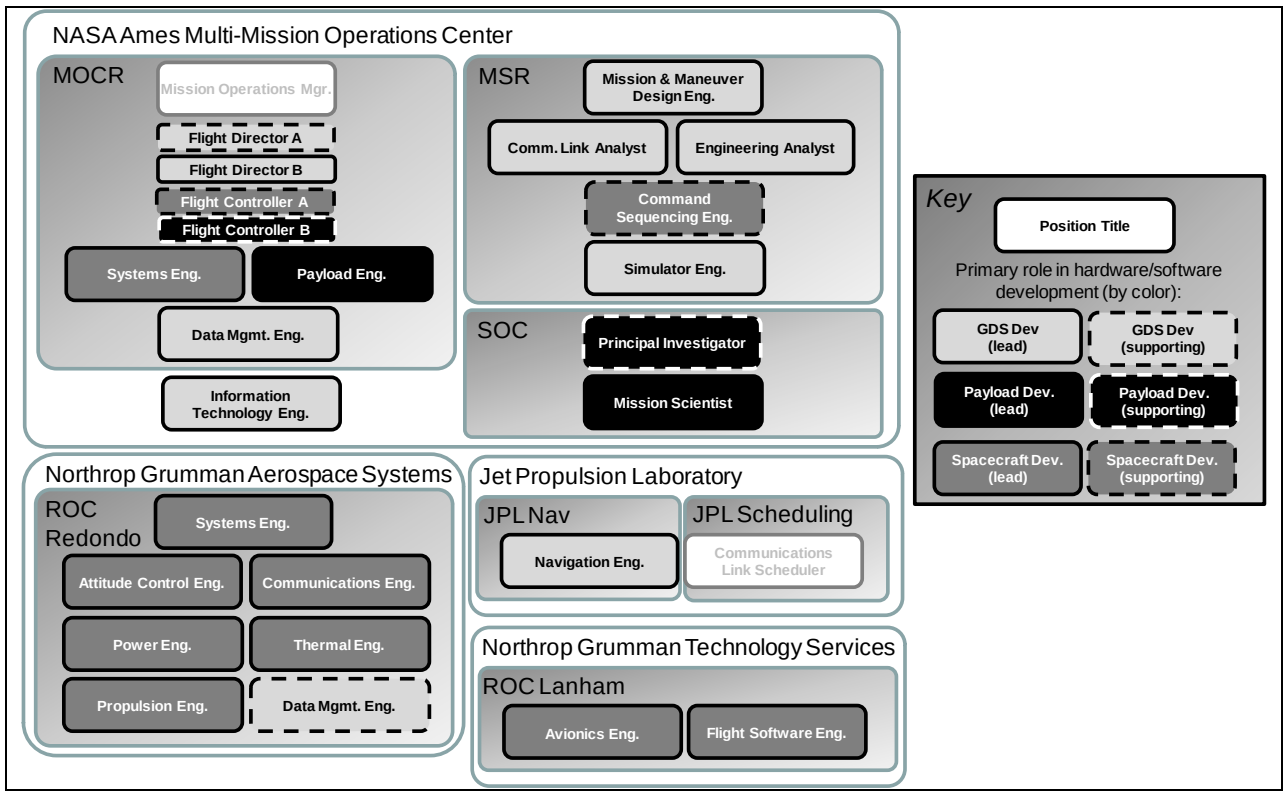

Fig. 4. MOS Operator Roles in Hardware and Software Development. *Most MOS staff played key roles in spacecraft, payload, or GDS development, including design, implementation, and testing. This simplified training, but contributed to a heavy workload and occasionally hindered simultaneous MOS and GDS development.*

These disadvantages continued in flight. LCROSS operators were, on average, busier than anticipated (see section VII). Inevitably, GDS bug fixes and enhancements were developed during flight, but developer/operators were too consumed with flight duties to perform global GDS deployments. The team had to resort to less-formal point deployments for specific tools, complicating the GDS configuration management task.

VII. Team Performance in Flight

Arguably, the best measure of flight team preparation is its level of performance during its flight mission. This section provides a brief synopsis of LCROSS MOS performance during 112 days of flight.

A. Summary

Transfer Phase was six days of 24-hour operations, covering launch through Lunar Swingby, and the most challenging nominal segment of flight, save the pre-impact sequence. The MOS operated in two overlapping shifts of 13 hours, synchronized with major events. The MOS successfully performed all nominal events and responded to

several in-flight spacecraft anomalies that forced significant operational changes from the baseline to maintain spacecraft health. On the first day of flight, the MOS detected a significant torque on the LCROSS stack, presumably induced by propellants leaking from the Centaur. It exercised the detection and some of the less severe mitigation tactics devised prior to launch to minimize the propellant penalty. However, the torques decayed significantly in the first days of the mission, and none of the more significant responses were needed. On Days 3 and 4, during an experiment to test a thruster thermal anomaly mitigation strategy (see below), flight rules were inadvertently violated, in the first case causing the spacecraft to transition to its Survival State, and in the second case causing a star tracker to transition to a standby mode (later recovered). The errors were attributed to a loss of situational awareness stemming from the departure from nominal operations. The processes for shift handovers and executing off-nominal commands were improved as a result, and the mistake had no lasting negative effect. By the end of Transfer Phase, the team was noticeably fatigued, particularly Shift B which had worked graveyard hours.

Cruise Phase, the period of the mission between Lunar Swingby and Impact Phase and occupying the bulk of time in flight, was more difficult than anticipated. Before launch, planning and training had focused on isolated Cruise Phase events. It had not sufficiently considered how the superposition of activities and unexpected changes in flight would affect the timeline and workload. The early part of Cruise was spent developing sustainable workarounds to anomalies discovered in Transfer Phase. Due to late changes in the LCROSS launch date, without the advance notice typically needed to establish a regular pattern of DSN tracks, DSN contacts were at highly variable times of day, and at variable intervals, making it difficult for the team to establish an operational rhythm. Results from earlier science events prompted the Science Team to modify and add to later payload calibrations, requiring in-flight re-designs of command products and procedures. Furthermore, efforts to remove ice from the Centaur upper stage were less effective than expected, prompting the team to design and execute two additional maneuvers for that purpose.

On the second Earth revolution, a substantial spacecraft anomaly (described later in this section) caused the MOS to divert from nominal operations for two weeks. The resulting effort to save remaining propellant and develop safeguards, all while supporting extra DSN passes, stretched the MOS team to its limit. Upon emerging from this taxing recovery period, the team finally progressed towards a more sustainable operational cadence. In the final weeks of the mission, rehearsals for Centaur separation and impact were interleaved with nominal operations, and the frequency of trajectory maneuvering increased in preparation for impact. Despite these challenges, the MOS performed with very few operational errors, none of which had measurable negative influence on mission outcome.

Impact Phase was executed nearly flawlessly. Over 27 hours, the MOS performed two full operational cycles to plan and execute Centaur separation and lunar impact. There were some voice loop communications problems in the final minutes of flight, attributed to shortcomings in training for real-time instrument commanding. These had some effect on data collected at the time of the Centaur impact, but with a negligible influence on overall science goals. LCROSS met all of its mission objectives.

B. Team Response to its First Significant Spacecraft Anomaly: Thruster Thermal Control

During flight, the LCROSS flight team responded to several off-nominal spacecraft and ground system events. Two anomalies in particular challenged the team's lightweight approach. In both cases, the team re-configured itself to balance the sometimes conflicting goals of maintaining daily spacecraft health and safety, continuing with critical mission events, diagnosing the cause of the anomalous behavior, and developing and testing workarounds and risk reduction strategies to regain nominal operations.

Within the first 24 hours of flight, the team noted that specific thruster valve heaters were not activating upon reaching the intended low threshold temperature condition. This problem risked freezing hydrazine in the propellant lines, which might cause their rupture; hence it was deemed a serious issue that had to be addressed immediately. The details of the problem are discussed elsewhere[‡‡‡]. Because thruster valve heaters were controlled by thermostats with fixed hardware heater setpoints, the team could not simply re-configure the thermostats to achieve better behavior. A team of NG engineers consulted the thruster thermal design and test data and quickly isolated the cause of the behavior. The only two effective means of warming these cold thrusters was to either fire the thrusters or to expose them to direct sunlight via a spacecraft attitude change. Without refinement, both techniques required frequent manual intervention by the flight team. Valve temperatures returned to the cold state just tens of minutes after firing a cold thruster (and consumed propellant), and most attitudes that warmed the cold thrusters also caused other parts of the spacecraft to get too cold or too hot.

Managing the thruster temperatures distracted the team from the already very busy schedule during Transfer Phase. With operations cycles (described in IV.A.) running around-the-clock in support of daily trajectory correction

[‡‡‡] Reference [4] provides a summary of flight, including a description of the most significant anomalies.

maneuvers, science instrument checkouts, and Lunar Swingby, the core team could not devote sufficient time to solve the problem, but could only fight its effects. To minimize the impact to critical nominal events, the Test Conductor, who for flight had been assigned to cover as a backup for several operational positions, was assigned to lead anomaly investigations. He tasked a small team from NG to perform analyses of spacecraft attitudes that might warm the thrusters while maintaining overall thermal balance. Based on these analyses, in between daily planning activities, he and the Command Sequencing Engineer created and simulated short ground command sequences to bias the spacecraft attitude by specific angles. In successive shifts, the MOCR team frequently commanded thruster firings and, as biasing command sequences became available, executed these sequences to experimentally determine the effect of the biases on thermal control. Meanwhile, the Mission Operations Manager reported daily on the status of anomaly operations efforts to ARC, LPRP, and NASA Headquarters stakeholders, insulating the team from further responsibilities during this difficult time.

The anomaly response process developed prior to launch had not sufficiently fleshed out the details of impromptu command sequence generation and approval, in-flight procedure development, and the communication of off-nominal spacecraft configurations between shifts. These shortfalls contributed to two human errors on successive days.

On Day 3, the team inadvertently caused the spacecraft to transition to Survival State. In an effort to warm the thrusters, one operational shift had placed the spacecraft in an attitude control mode with a narrow deadband to naturally induce more thruster firings. This was a deviation from the standard Cruise state configuration. Under flight rules and onboard fault detection software, attitude changes from this control mode were illegal. A breakdown in shift handover communications meant that the oncoming shift was not aware of this change. This following shift intended to command a spacecraft attitude change to test a biased orientation that might maintain thruster thermal balance. Before the attitude change was commanded, with insufficient time to develop detailed procedures for this off-nominal activity, team members neglected to check the current control mode in telemetry. The flight-generated ground command sequence did not have some of the precondition checks that were more common in sequences developed before flight, and hence did not check the attitude control state either. Upon receiving the command, onboard fault management triggered a spacecraft reset and transition to Survival State. The team regained the nominal Cruise State two hours later, with no negative effect on the mission.

On Day 4, a similar attitude bias maneuver inadvertently pointed the spacecraft star tracker toward the sun, causing the sensor to transition to a standby mode. For command sequences conceived before launch, sensor interference conditions were commonly checked. No such check was instituted for these attitude bias command sequences. Due to time constraints in the mission schedule, the team was unable to restore nominal operation of the star tracker prior to Trajectory Correction Maneuver (TCM) 3. Though not illegal, this was an off-nominal configuration for such a maneuver, and temporarily left the spacecraft dependent on inertial navigation for absolute attitude tracking. The burn was successful, and following that maneuver, the team restored the star tracker to its nominal operational state, with no lasting effects on the mission.

Immediately following these two errors, the flight team significantly improved its shift handover process, and established stronger required preconditions for anomaly commanding, including command sequence testing, analysis, and approval, and accompanying procedure development. This enhanced process served the team successfully for the remainder of the mission. Within the first two days, after many trials, the team discovered attitudes that kept the spacecraft thermally stable for hours (but not indefinitely). The team continued in this way through Lunar Swingby.

Following Lunar Swingby, with process lessons learned from the first week, and far fewer nominal activities to attend to, the team re-configured to focus on enabling a long-term operations strategy. The more relaxed operational schedule allowed the Lead Systems Engineer to assume leadership of the remaining anomaly resolution tasks , and he received dedicated support from the core team at ARC, including the Command Sequencing Engineer, the other Systems Engineers, and Flight Directors. All of these operators attended design meetings for new onboard command sequences to augment spacecraft fault protection. The team designated a minimum set of analyses and simulations to be conducted before new command sequences could be approved. It also held Command Approval Meetings for each revision prior to execution on the spacecraft. Furthermore, the team drafted increasingly more formal procedures to accompany each successive onboard command sequence delivery. These ultimately resulted in several changes to nominal procedures to accommodate the new workaround strategy. Despite these changes, the workaround strategy allowed the team to fly LCROSS in nearly the same way as it had been designed. This first significant spacecraft anomaly was an important formative experience that the team would call on again later in the mission.

C. Team Response to its Second Significant Anomaly: IRU Fault and Excess Propellant Usage

On Day 64 of the mission, an unexplained, short-lived error condition on the spacecraft's inertial reference unit (IRU) triggered nearly continuous attitude control thruster firings that remained undetected for nearly 24 hours during a scheduled DSN coverage outage. Upon detecting the faulty state, MOCR operators sent commands to safe the spacecraft and restore nominal attitude control. The team quickly recognized the severity of the anomaly – pressure readings indicated that LCROSS had expended 152 kg of 190 kg of propellant remaining. After consultation with other operators, the Mission Operations Manager contacted the DSN to declare a spacecraft emergency, which enabled the DSN to secure maximum additional antenna time to permit the LCROSS team to closely monitor the spacecraft and to continue diagnosing the problem. Within hours of the anomaly discovery, the combined ARC and NG team had correctly identified the coarse chain of events leading to the loss of propellant.

Following this initial response, the team spent roughly two weeks compiling detailed evidence of how the anomaly occurred, designing and testing a succession of protective measures, and analyzing whether remaining propellant would enable mission success. Fortunately, Cruise Phase was far more tolerant to delayed execution of trajectory correction maneuvers, and the team opted to suspend all nominal operations to focus on the anomaly. Freed from the pressure of daily time-critical events, the Lead and Deputy Systems Engineers took the lead in anomaly resolution, and a greater segment of the team involved in various aspects of the resolution effort. Given the serious threat this anomaly posed to mission success, NASA ARC center management and the ARC Office of Chief Engineer worked with the LCROSS Project Manager to provide strategic oversight of the effort.

The flight team transitioned to 24-hour operations, with breaks in DSN coverage only where geometric line-of-sight was impossible[§§§]. It settled into a schedule of three shifts, led by the Flight Directors, each focused on a different aspect of the recovery. The morning shift designed, implemented, and simulated new preventative fault protection measures. Team members held a coordination meeting each morning to determine the daily objectives. The majority of the shift was spent in the MSR in design discussions, command sequence or software implementation, and spacecraft simulator testing. After design and test cycles of one or more days, the team would hold a design review in which the broader team would inspect the designs and evaluate test results, and critique the proposed flight test plan. Upon Flight Director and Mission Operations Manager approval, new command sequences and table loads were loaded onto LCROSS from the MOCR and tested according to the test procedure. The morning shift was 10-12 hours long. The afternoon shift, also 10-12 hours long, overlapped the morning shift by several hours to permit better team-wide coordination, discussion and test continuity. The afternoon shift predominantly performed and monitored extended tests of any newly-augmented software and command sequences, and analyzed the initial anomaly in greater depth to support reporting to stakeholders. ARC management and other LCROSS stakeholders mandated the spacecraft be monitored 24 hours daily while fixes were being put in place to prevent a potentially mission-ending loss of propellant. A third overnight shift, added to accommodate this request, was instituted to make sure the spacecraft remained under constant surveillance.

Given the intensity of activities in the first and second shifts, staffing the overnight monitoring shift continuously for two weeks was difficult. It was decided that a minimum of two operators should perform spacecraft monitoring overnight, both to ensure the operators would remain awake through the night, and also for operator safety. To offload the core team, some nights were staffed by one ARC and one NG operator rather than two ARC operators. More creatively, ARC management helped arrange volunteers from the greater ARC operations and research community to serve as second operators with a dedicated LCROSS operator for overnight shifts. This was a very successful program – it satisfied the alertness and safety needs for the shifts, and it was a useful and exciting educational experience for a number of ARC employees interested in future spacecraft operations.

Focused entirely on anomaly recovery, the flight team was very productive during this time period. It developed a comprehensive set of improvements to spacecraft fault management, control, and operational procedures, in three efforts:

- IRU Fault Management Improvements: The first priority was to prevent a repeat of the same spacecraft fault response and propellant loss. The team re-designed the IRU fault infrastructure to ignore spurious IRU faults, and to power-cycle the IRU and retry IRU-based control before resorting to control using star tracker rate measurements.
- General Propellant Fault Management Improvements: Given the precarious state of LCROSS propellant reserves, NASA ARC center management pushed the LCROSS team to augment LCROSS fault management to prevent any future fault from expending excess propellant. Leveraging the strategy originally developed to counteract a Centaur gas leak (see Section IV.D), the team augmented spacecraft

[§§§] During line-of-sight periods, continuous coverage was made possible largely by other mission teams who yielded their DSN time to LCROSS

fault management to detect greater-than-expected propellant usage, and to respond by dropping into a passive, spin-stabilized control mode that could maintain a sun-pointed attitude for many hours without thruster firings. Without a direct means of detecting and recording excessive propellant usage, the team designed an improvised, indirect approach. The onboard fault response built upon the fault infrastructure that performed safe transitions to Survival State. All serious spacecraft faults resulted in transitions to Survival State; however, these new spacecraft upgrades resorted to the more severe spin-stabilized response only in the case of excessive thruster firings.

- Attitude Control Efficiency Improvements: An outcome of the anomaly investigation was that the attitude controllers were not designed to perform efficiently under the noisier body rate measurements produced by the star tracker. In a parallel effort, NG team members re-designed one of the primary attitude control algorithms to drastically reduce propellant usage using star tracker rate measurements (as required after a real IRU failure). As an added benefit, this new controller was also far more efficient under nominal conditions, allowing LCROSS to preserve its dwindling propellant supply.

All of these efforts were tailored to allow LCROSS to eventually emerge from emergency DSN operations and return to a schedule with less frequent spacecraft contacts, per the nominal mission plan. All changes required several stages of development, with repeated design iterations (to protect the spacecraft with partial fixes as soon as possible), simulator-based testing, team reviews, and ultimately spacecraft-based testing before being baselined for flight. The operations team performed a final series of spacecraft tests of key changes to prove to LCROSS stakeholders that the spacecraft was sufficiently well-protected to be removed from emergency operational status. Finally, on Day 77, LCROSS released itself from the emergency.

With the emergency over, ARC management pushed hard to maintain more frequent contact with LCROSS than was originally planned. The flight team was asked to limit any communications outage to no more than eight hours, to ensure that LCROSS would not fall into a serious anomaly without its knowledge. DSN schedulers were extremely helpful in negotiating this time with the help of other missions. However, this continued heavy staffing represented a serious burden to the flight team, since it had been working around the clock for two full weeks. To meet the eight-hour rule, it was often necessary to add fragments of DSN contact time (1-3 hours) in the middle of the night. For any pass involving commands, one of two Flight Directors and one of two Flight Controllers[****] were required to be present. These team members, in particular, were becoming gradually more fatigued. To reduce the staffing pressures, the LCROSS project pursued two avenues. First, it devised a coordination plan with the DSN to designate and perform "Unattended" tracking passes, where DSN antenna operators would acquire spacecraft telemetry and perform ranging with the spacecraft, but during which LCROSS operator support would be optional. In the event that the DSN could not successfully establish contact with LCROSS, it would immediately notify the Mission Operations Manager or one of the Flight Directors. Under agreement with ARC management, LCROSS operators were obligated to attend at least one DSN contact each 24 hour period. Under the new policy, the team could designate the overnight contacts as "Unattended", and thereby establish a more regular daytime schedule. The second move was to negotiate with ARC security and Multi-Mission Operations Center information technology support to approve secure remote login to the mission network and to assemble several secure mission laptop computers, so that operators would be able to monitor the spacecraft from home. Once this was instated, Flight Directors and Flight Controllers could avoid a long drive from home to ARC on weekends, and could also monitor overnight contacts with LCROSS as necessary or desired without coming in to work. The "Unattended" contacts and mission laptops were critical in offloading the team after such a long emergency period.

There were distinct places in the timeline, however, where the DSN was not able to accommodate the maximum outage, particularly when the LCROSS orbit took it below the equator where contact would be limited largely to a single ground station in Australia. The LCROSS project was asked to augment its DSN contacts with other ground station networks. LCROSS investigated the possibility of utilizing the NASA Goddard Space Flight Center (GSFC) Ground Network (GN), and engaged in several exploratory and negotiation meetings with GN representatives, and even conducted two in-flight tests to prove the concept. This was another effort that occupied the flight team, even after the anomaly had been officially resolved.

VIII. Conclusions

What were the keys to the success of the LCROSS flight team? Perhaps most importantly, LCROSS benefitted tremendously from a pervasive spirit of cooperation and trust that crossed organizational boundaries. This improved communications at all levels, spawned additional collaborations, and caused people to devote extra time throughout

[****] The former Test Conductor was also later qualified as a backup Flight Controller to offload the other operators.

the project to ensure mission success. From the MOS development perspective, securing full-time access to the LCROSS simulator was critically important. It provided the team with an accurate test platform in creating the hundreds of command products to support various mission events, and allowed the team to gain early, risk-free, and frequent experience in the operation of LCROSS.

Not surprisingly, frequent, repeated testing of the team under realistic conditions was invaluable. These tests exposed weaknesses throughout MOS development, enabling the team to refine its equipment, processes, procedures, and staffing schedules before launch. By launch, the team had "flown" each major event multiple times, and this significantly improved team confidence in flight. As an unintended consequence, MOS simulator-based tests were partially responsible for exposing two significant spacecraft bugs prior to launch (subsequently corrected). MOS testing, using MOS-developed flight command sequences, provided a level of realism that could not easily be achieved in standard spacecraft system-level verification tests.

One aspect of operations that was not well-practiced was extended anomaly resolution. Through the above-mentioned tests, the team was proficient at identifying and quickly responding to off-nominal spacecraft behaviors to maintain safety. However, none of the tests or rehearsals forced the team to exercise the process of long-term anomaly resolution. As we discovered in flight, splitting the team to simultaneously manage an ailing spacecraft, determine the root cause of a problem, and design workarounds was a major challenge, and deserved greater attention before launch. This training shortfall was, in part, due to the difficulty in designing and conducting a realistic, engaging simulation of off-nominal spacecraft conditions. To exercise the process fully, a simulated anomaly has to be severe enough to prevent nominal operations, and convincing and subtle or complicated enough to consume the team's attention for multiple shifts, It should also demand a significant design effort to mitigate. A decent anomaly exercise should also involve higher levels of management in daily reporting and decision-making. Beyond familiarizing the team with the anomaly process, exercising the anomaly process in simulation would likely have exposed its weaknesses before launch. This would have pushed us to improve the process, presumably resulting in improvements to our flight performance. That stated, one could argue that this level of preparation is a natural candidate for de-scoping under a Class D approach. Our relative lack of practice here may have resulted in team inefficiency, but did not present a strong risk to the mission.

Managing workload with a limited staff was a continual challenge for LCROSS. The development phase, entailing vast overtime hours, remained challenging up until launch. The flight schedule was busier and more irregular than anticipated, contributing further to team fatigue, particularly during anomaly responses. For the most part, the team was getting sufficient sleep to avoid human error, but had little time off to tend to personal matters. Due to a lack of time and a small team, the MOS did not train a full set of backups for critical operational positions. Fortunately, the team did not falter on attendance or performance due to illness or accident. Having a floating backup operator with extensive operational experience was an efficient use of manpower – the former Test Conductor served as a backup Flight Controller, Systems Engineer, and Anomaly Resolution Lead during flight, and was trained as a backup Command Sequencing Engineer. There were also an insufficient number of personnel to compose a separate anomaly response team, and therefore day-to-day operations competed for time with anomaly investigations, particularly during Transfer Phase. However, the MOS responded effectively to a limited (but significant) number of anomalous events concurrently with routine operations. Future missions must realize that the Class D approach will likely entail greater crew stress.

Distributed operations were largely successful. The MOS made substantial improvements to communications prior to launch, and operators from all disciplines made strong contributions to the team, despite being remotely situated. However, communications and situational awareness could have improved even more. During anomaly investigations, NG took full advantage of its co-location with the full spacecraft test team and documentation resources in determining root causes for problems, confirming test results, and even designing and testing flight attitude control parameters to improve spacecraft propellant performance. Some basic problems (e.g. poor room acoustics) hampered communications between ARC and NG ROCs. Furthermore, we observed that the two roles played by the NG ROCs had mutually exclusive communications models. Supporting closely-coordinated procedure execution demanded quiet focus and attentive voice loop participation, whereas providing assessments of spacecraft performance benefitted from offline group discussion. Recognizing this conflict, we revised ROC protocols mid-flight to separate the two activities, with noticeable improvement.

Staff members that were dually-tasked as MOS operators and GDS developers were unable to fully satisfy the demands of both roles in flight. High-priority MOS tasks consumed most of operators' time, leaving no time to perform rigorous testing in support of a global GDS deployment. Instead, critical fixes were introduced on isolated workstations. A suggestion for future missions is to employ distinct GDS developers that double as backups for key operational roles.

Prior to launch, telemetry data management, interpretation and easy access were known but underappreciated challenges. Experience gained during the mission prompted some changes that improved team performance in all three areas. NG and ARC both utilized a dedicated Telemetry Data Manager. The ARC operator supported the Flight Controllers in managing telemetry and command data connections to the DSN during more complicated antenna transitions, and managed and archived real-time and stored telemetry data streams for team consumption. The NG operator composed plots of key engineering data after each spacecraft event in support of Assessment Phase. Both operators helped avoid duplicate work of telemetry data organization and display, and allowed the engineering team to focus on data interpretation. In a separate effort, the ARC Payload Engineer developed a very useful tool for quickly plotting multiple telemetry data over arbitrary time intervals[††††]. This made it far easier to scan data for extreme values, find transitions in spacecraft behavior, and establish time correlations between measurements. Though it began as a tool to assess the instrument payload performance, it was used universally by both the ARC and NG engineering team. Finally, as a by-product of the IRU/propellant anomaly, the LCROSS project pushed ARC to allow its operators to access spacecraft telemetry via remote network connection. Security concerns barred this type of a capability in the early part of the mission. However, when the team was required to support a greater frequency of spacecraft monitoring, ARC assembled a set of secure laptop computers, assigned to senior operators, that enabled them to assess spacecraft health from home. The move averted further crew fatigue, and, because it made telemetry access so easy, actually resulted in increased oversight of the spacecraft.

Interpreting the intent of Class D continued to be a challenge in flight. Through several anomaly responses, the flight team made significant changes to the Shepherding Spacecraft fault management and attitude control system, under stressful conditions. Some changes were essential to mission success. Others could be viewed as going beyond the scope of the Class D philosophy. The investigation into the IRU/propellant anomaly showed that, aside from the propellant consumption, the spacecraft remained completely healthy. LCROSS design upgrades were not only designed to prevent similar anomalies from happening again, but also to protect the spacecraft from a broad class of propellant anomalies. This increased spacecraft complexity, but had far less test heritage than the pre-launch design. Once the fixes were in place, the team maintained a monitoring schedule that was far more intensive than originally conceived, adding to team fatigue. Were these actions inappropriate for a Class D mission? The answer is not clear. However, future Class D missions should carefully consider whether in-flight changes to the spacecraft design or operational plan, particularly those that exceed the protection levels set prior to launch, are likely to reduce or add risk to mission success.

It is important to recognize that the Class D approach enables cost and schedule savings by accepting greater risk. Most of the shortcomings in flight team preparation were not surprises discovered in flight, but gaps that had been fully recognized and accepted prior to launch. Still, some of these could have caused the mission far greater harm had more severe anomalies occurred during less forgiving parts of the mission. Some part of the LCROSS success must be attributed to good fortune.

Despite its success in supporting a single mission, it is debatable whether the LCROSS MOS development approach is sustainable over multiple missions. LCROSS succeeded, but this could not have happened without the extraordinary hours the team contributed from start to finish – a level of effort that is not easily repeated.

Acknowledgments

This work was funded wholly by NASA under the LCROSS project. The authors wish to thank the other members of the LCROSS flight team for their operational contributions to this work; the other members of the LCROSS project, who initiated, built, and carried the program to launch; the management teams from NASA ARC, Northrop Grumman, JPL, NASA GSFC who established our cooperative work atmosphere, who often pushed us hard but kept us on track throughout; our independent reviewers, who found our weaknesses early, and helped us shape our approach to Class D; the LPRP program office at NASA MSFC, and NASA Headquarters, for their intelligent, supportive oversight; and all the others who contributed to the success of the LCROSS mission.

References

[1] Andrews, D., "LCROSS Lunar Impactor: Pioneering Risk-Tolerant Exploration in a Search for Water on the Moon", *Proceedings of the 7th International Planetary Probe Workshop*, Barcelona, Spain, June 2010 (submitted for publication).

[2] Binder, A., "Lunar Prospector: Overview", *Science*, Vol. 281, No. 5382, 4 September 1998, pp. 1475-1476.

[††††] The original LCROSS ground data system did not provide a time-efficient means of archive data visualization and analysis.

[3] Borucki, W., et al., "Kepler Mission: Design, Expected Results, Opportunities to Participate", *A Decade of Extrasolar Planets Around Normal Stars*, edited by M. Livio, K. Sahu, and J. Valenti, Space Telescope Science Institute Symposium Series, No. 19, Cambridge University Press, 2005, pp. 36-49.

[4] Tompkins, P., et al., "Flight Operations for the LCROSS Lunar Impactor Mission", *AIAA Space Ops 2010 Conference, Delivering on the Dream*, AIAA, Washington, DC, 2010, Paper 2010-1986 (submitted for publication).

[5] Strong, J., et al., "Transport and Use of a Centaur Second Stage in Space", *AIAA Space Ops 2010 Conference, Delivering on the Dream*, AIAA, Washington, DC, 2010, Paper 2010-2197 (submitted for publication).

[6] Alena, R., D'Ortenzio, M., Vho, N., "Communication Link Prediction and Assessment for LCROSS Mission Operations", AIAA Space Ops 2010 Conference, Delivering on the Dream, AIAA, Washington D.C., 2010, Paper 2010-1933 (submitted for publication).