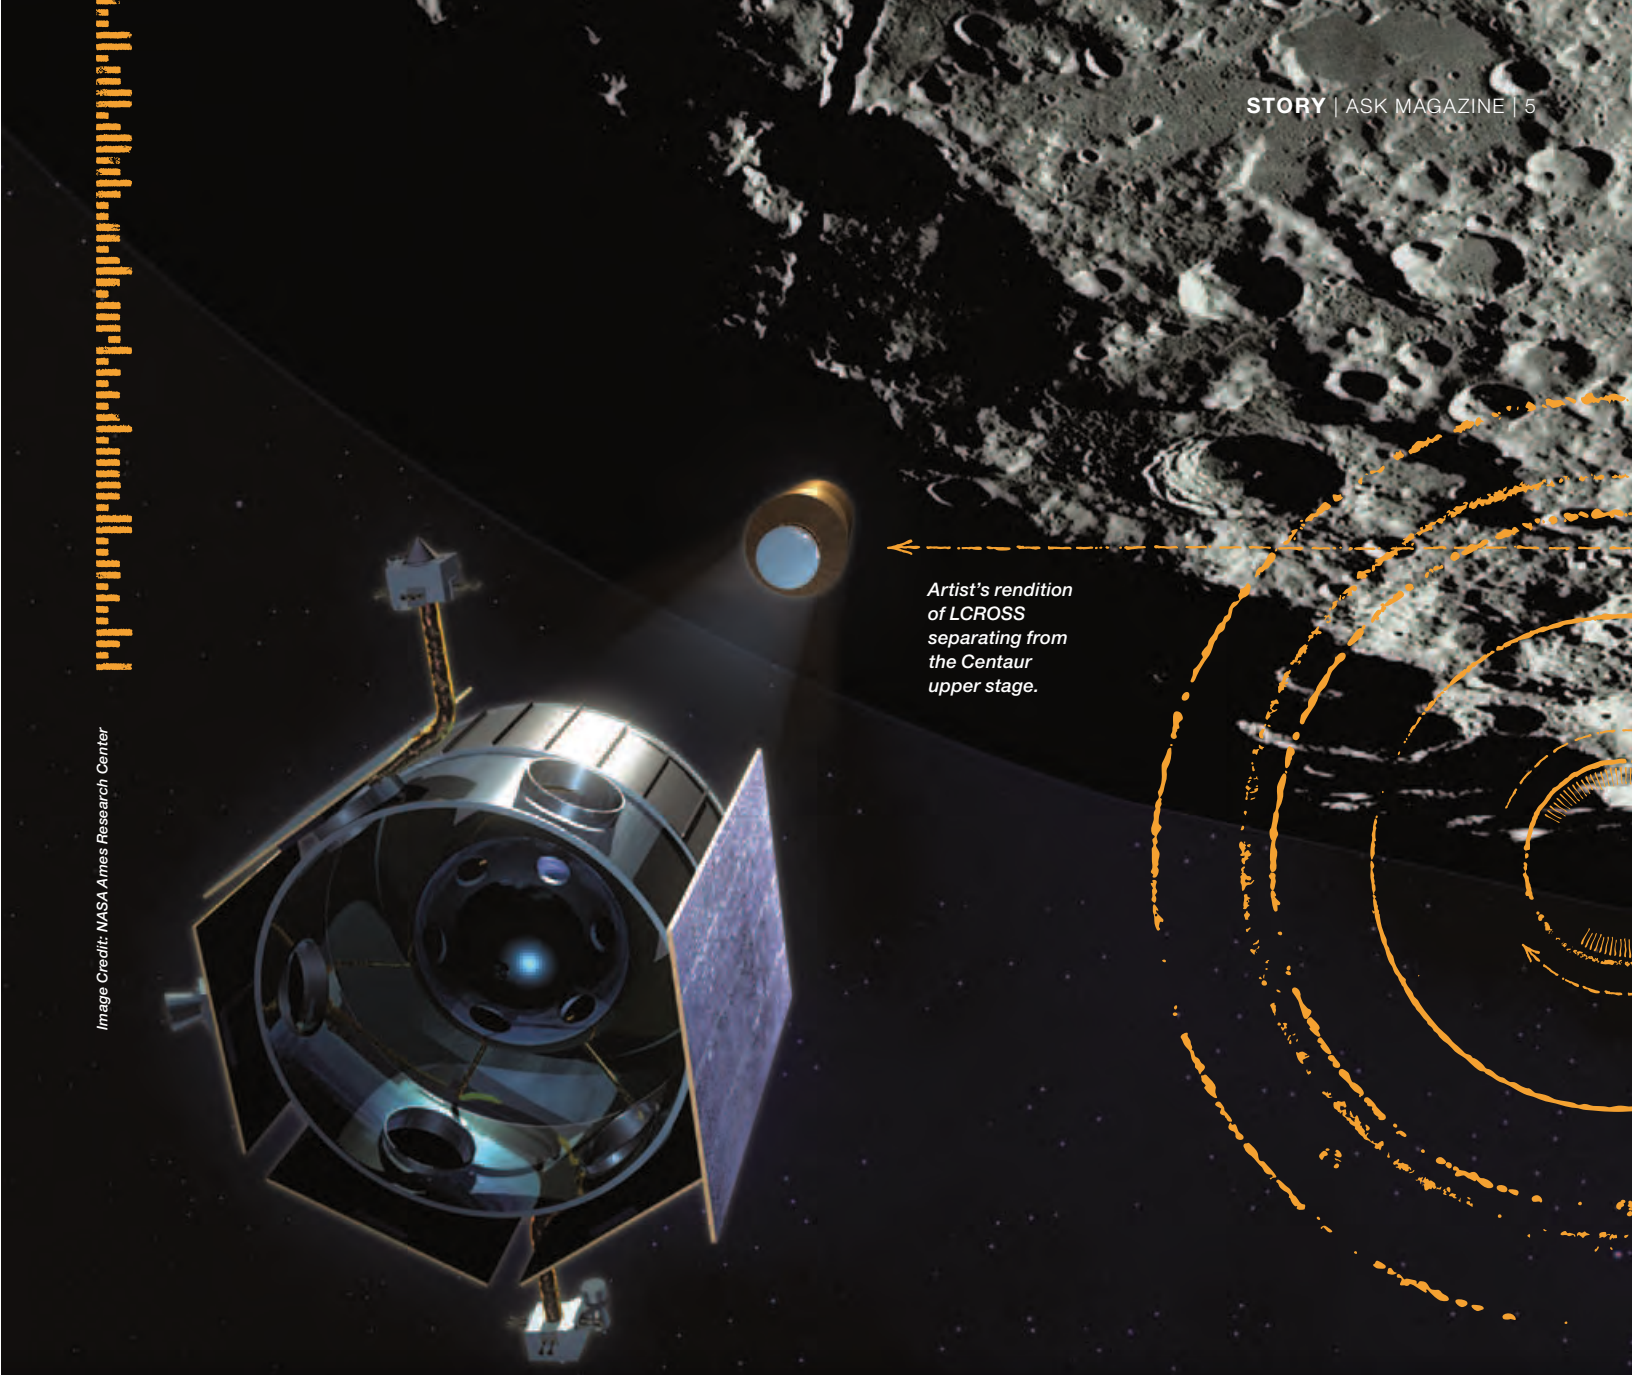

Artist's rendition of LCROSS separating from the Centaur upper stage.

Image Credit: NASA Ames Research Center

MANAGING THE BAD DAY

BY DANIEL ANDREWS

It's 3:30 a.m. on Saturday, August 22, 2009. My cell phone rings. As the project manager for the Lunar Crater Observation and Sensing Satellite (LCROSS), I was used to sleeping with the phone near my bed ever since launch. The LCROSS operations team was preparing to do a spacecraft-orientation maneuver, turning the cold side of the spacecraft to the sun to burn off any residual ice remaining on the Centaur upper stage—what we called a "cold-side bake." I was planning to go in and observe the activities later that morning. The phone had never rung this early before.

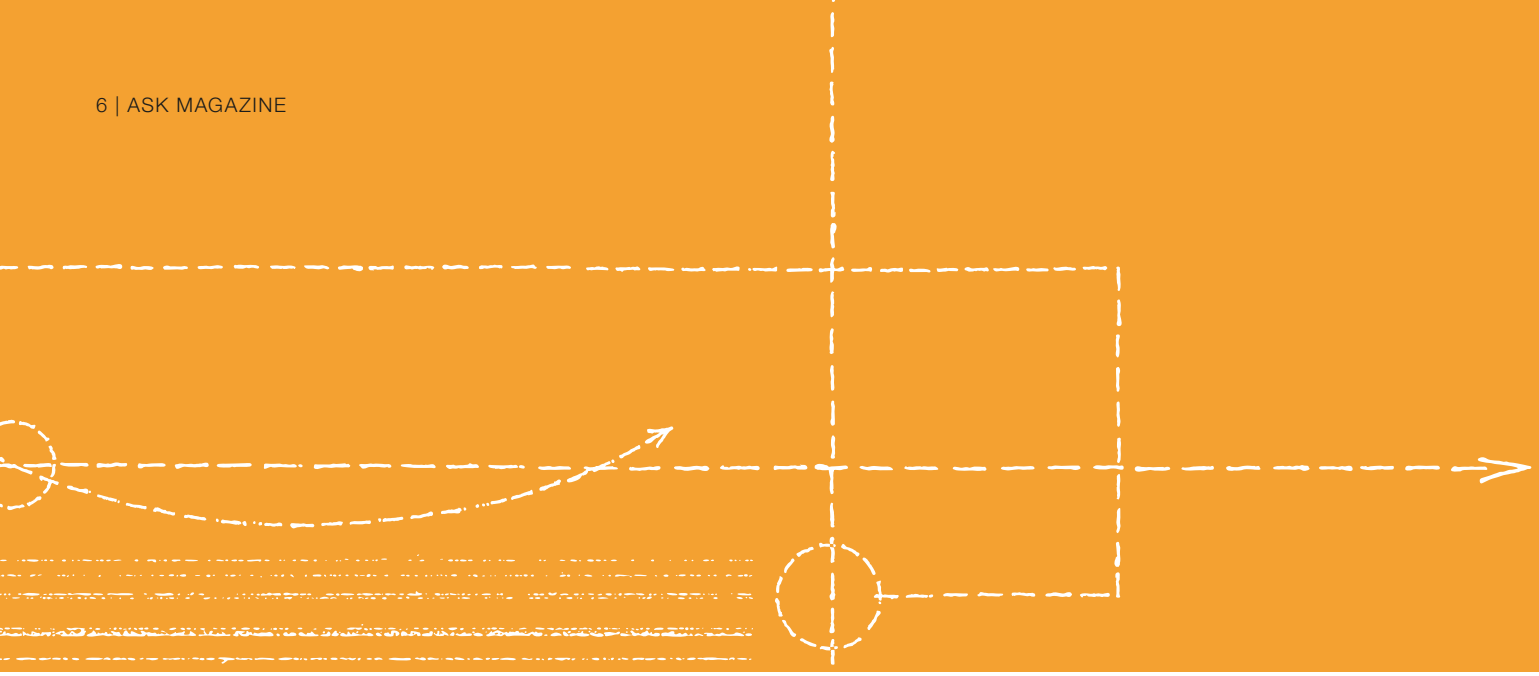

"Project, this is Mission," the LCROSS mission ops manager (MOM) stated.

"Go, Mission," I replied.

MOM indicated the team had just gotten "acquisition of signal," which means the operations crew had reestablished communication with the spacecraft after a planned period of no communication. MOM told me that once spacecraft telemetry began flowing, the ops team discovered that a very large amount of propellant had been mysteriously consumed while the spacecraft was out of view of the ground stations.

MOM explained, "When we acquired the spacecraft, we discovered that the thrusters were firing almost continuously and believe a substantial amount of propellant was consumed." I asked if we knew if we had enough propellant remaining. "We do not yet know if we have enough propellant to finish the mission—working it now," replied MOM. "The thrusters are still firing, and we are trying to get that stopped."

It was clear that if we hadn't scheduled an early-morning activity when we did, we would have consumed all the propellant and lost the mission. Furthermore, if we didn't get it stopped immediately, we'd lose the mission anyhow.

This was LCROSS's bad day.

I got dressed and headed in to the mission ops control room at Ames Research Center and learned the thruster firing had stopped after a commanded power-cycling of the spacecraft's inertial reference unit, or IRU. The IRU is standard spacecraft equipment used to measure the spacecraft's velocities so its attitude can be controlled. The ops team discovered that an IRU fault flag was set. After some consideration, the team issued a reset command, which cleared the fault and halted the thruster firings, returning the spacecraft to its normal condition.

Later analysis revealed that when the IRU fault occurred, the autonomy and fault management system appropriately kicked in, no longer trusting the IRU for velocity feedback and switching to the star tracker's velocity feedback. For (then) unexplained reasons, this changeover drove the attitude control system to fire the spacecraft thrusters at an extraordinary rate. The spacecraft ultimately consumed some 140 kg of propellant, leaving a mere 60 kg to finish the mission.

It eventually turned out that two root causes led to this event and our subsequent challenges:

1. **IRU configuration error:** A spurious, short-lived error on the IRU was interpreted as a more serious fault by the spacecraft fault-management system because the IRU fault-flag update rate and the autonomy and fault management sampling rate were not properly synced, leading the autonomy and fault management system to believe a persistent error was present and to subsequently switch to the star tracker for velocity measurements. This issue alone wouldn't have been a problem.
2. **Star tracker velocity noise:** Since star-tracker measurements compute velocity from the spacecraft position relative to the stars, the computations can be noisy, or jittery, which is why IRUs are employed for velocity measurements. The noise levels were within manufacturing specifications, but our high-performance spacecraft attitude-control system was sufficiently sensitive to think the noise was velocity error and tried to control it when it should have ignored it. This led to the excessive thruster firings and propellant consumption.

LCROSS formally declared a spacecraft emergency with NASA's Deep Space Network, given the spacecraft's precarious condition. With this declaration, all missions using the Deep Space Network have an understanding to yield their communications pass time to a mission in danger. This enabled LCROSS to have near-continuous communication with the ground, limited only by geometric constraints of the spacecraft's position relative to ground stations on Earth.

As it turned out, one of those outages was again coming, so we needed to put some protections in place just ten hours after

Photo Credit: NASA/Jack Pfaller

LCROSS and LRO are installed inside their fairing.

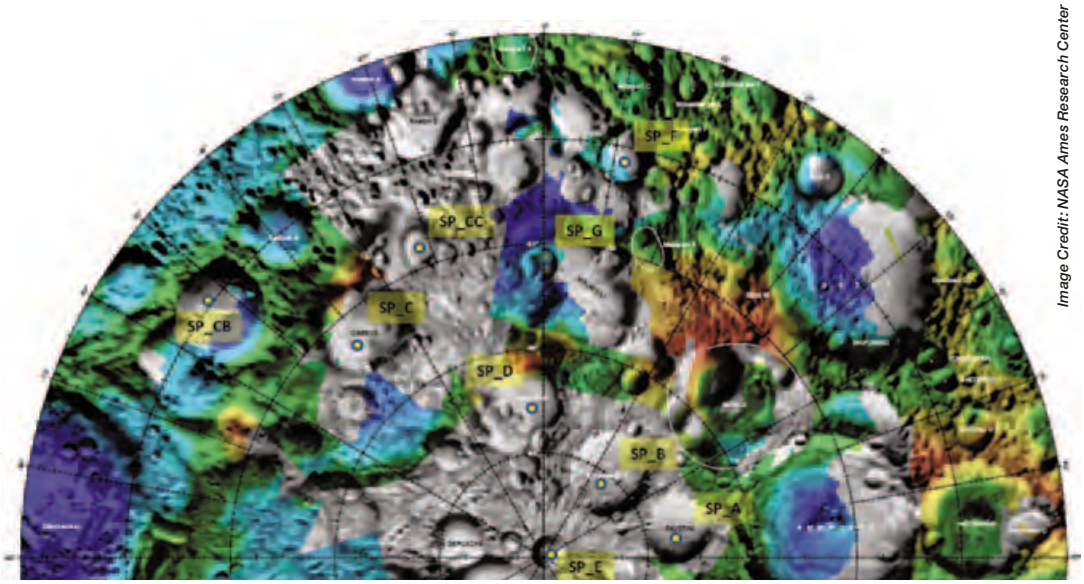

Image Credit: NASA Ames Research Center

LCROSS candidate impact craters.

discovering the anomaly. Our plan was to update the persistency with which the IRU fault was monitored so a spurious fault would not throw us into another costly propellant-consumption situation. Then we went dark again and crossed our fingers.

From Anomaly to Recovery

When communications were reestablished, we discovered there had been no further incident. We had made it through, but this was the beginning of a new operational environment for LCROSS as we moved from anomaly to recovery. This required serious triage. Here were the steps we took:

1. **Stop the bleeding.** The mission is over if you cannot stop the elevated rate of consumption of a finite resource like propellant. Electrical power can be renewed through solar arrays, but there is no mid-air refueling of spacecraft propellant. We needed to stop the propellant consumption ASAP.
2. **Make it through the night.** We needed to survive upcoming known communication outages caused by orbital geometries. We needed a way for the spacecraft to monitor when excessive firing occurred and prevent further consumption *automatically*.
3. **Ensure long-term health.** Once you are out of imminent danger, how do you ensure finishing the mission? What are the tasks remaining and the risks of executing them? How far do you go with analysis, simulations, and other risk-mitigation means? At what point does the risk of human error become greater than the technical risk associated with the spacecraft?
4. **Address the root cause (if you can).** Discover the specific cause for the incident. Is there anything that can be done to prevent this in the future? Is there a way to fix it, or only ways to avoid the circumstances that led to it?

The Project Manager's Role

Along with this triage process, the operations team's most important job, the project manager takes on a new series of responsibilities when a mission has a "bad day."

Inform and Manage the Stakeholders

Understandably, stakeholders get very engaged after an anomaly. They want to help ensure the mission. The morning of the anomaly, I followed established procedures to call the various stakeholders and inform them of what had happened. Shortly after those notifications went out, the Ames center director and most of his directors arrived at the ops control room with bags of breakfast food and drinks, a gesture much appreciated by the team. And we were grateful that leadership understood the team needed to be given room to work.

I provided frequent stakeholder updates on findings and progress, in person and via e-mail for the broader agency audience, with a brief daily status teleconference by the MOM. E-mail updates were nearly hourly in the beginning, dropping to updates at shift changes near the end of our emergency. My deputy project manager and I tag teamed to cover shifts in the mission ops control room, writing a summary and publishing it to the stakeholders at shift changes, keeping the stakeholders informed and comfortable.

Protect the Team from External Distraction

The LCROSS team was of course attempting to get back to more normal operations as soon as feasible after the anomaly. Center management demanded that additional controls be put in place to protect the remainder of the spacecraft's propellant; however, this challenged the team at a time when they were stressed and fatigued—our staffing plan was not designed to support 24-7 operations. It is the project manager's job to try to manage stakeholders to a consistent level of risk tolerance, despite the strong drive to eliminate future risk, which is not

possible. This mission had grown to be very important to many, but reason and balance needed to prevail.

Steer Parties Away from Hunting for the Guilty

Once you stop the bleeding, questions naturally begin to surface about why the anomaly occurred. These queries, while important to understanding your continuing risk, should not distract the team from focusing their attention on continuing the mission. I had to push back on this questioning to prevent the team from getting frustrated or distracted.

Handle the Press

When a spacecraft experiences an anomaly, you have to be available to the press. The traditional media want to know all the details and can turn against you if they suspect you are holding back; openness is important. The blogosphere is different in that their "facts" come from unknown sources and their conclusions are sometimes based on personal agendas. We handled the press with frequent phone interviews and updates to the project web page. I conducted about ten phone interviews in two days.

Watch for Things Getting Complicated

After the anomaly, engineers worked through the data and invented responses, but engineers (like me) are predisposed to solving problems and have a tendency to create complex, multilayer solutions to stomp out the risk of reoccurrence. Discussions would work their way from one incremental fix to another, arriving at complex fixes and patches that would move the team far from its operations training and might not be testable. This complexity growth actually *grows risk* that the system will become so sophisticated it will be prone to operator error or create unforeseen interactions. In the heat of battle, there needs to be someone who keeps an eye on the risk of the solution. There were a couple of times when I would ask, "Do we need to go that far, or can we live with just the first corrective measure?" We would usually agree we could accept the residual risk after addressing the

Photo Credit: NASA/Dimitri Gerondidakis

On Launch Complex 41, the Lunar Reconnaissance Orbiter and LCROSS are moved into the mobile service tower.

(Left to right) John Marmie, Jack Boyd, Lewis Braxton III, Tina Panontin (standing), Pete Worden, and Chuck Duff celebrate LCROSS's separation from the Centaur upper stage.

Photo Credit: NASA/Eric James

principal problem. Missions have been lost because smart people did well-intended things that made problems worse.

Watch Operations Console Staffing

Because the LCROSS team was small, we had the project systems engineer staff the systems engineering console station. The project systems engineer would take one shift, and his deputy would staff the other shift. The idea seemed sensible—why not put your most competent systems engineer right in the middle of the action? I later realized that having your project systems engineer on the console removes him from his normal responsibilities—that you still need. Yes, you benefit from having your lead systems engineer monitoring the spacecraft, but he needs to sleep as well and is less able to participate in important assessment and planning activities, making him unavailable to advise you with his technical assessments and recommendations. I would not organize staff this way again.

Watch for Crew Fatigue

Hardworking, dedicated people get tired. Our cost-capped mission was not designed for post-anomaly staffing demands. A small number of people were covering an extraordinary number of hours. Their work was impressive, but fatigue inevitably sets in. You need to balance attacking technical problems with the growing operational risks associated with fatigue. I saw heads bobbing while on console as people fought back sleep; I saw people struggle to complete thoughts during shift-handover discussions. There was also growing stress at home for many who were working difficult hours. It was essential to remediate the problem as soon as possible.

Meeting the Challenge of the Bad Day

The LCROSS team behaved remarkably through its bad day. The triage process was exactly the right mix of urgency and focus, which comes from many, many operational rehearsals where the team trains for what is supposed to happen and even what is not supposed to happen. Of course, you cannot afford to spend unending money training for a low-cost mission, which means you need to focus not on the specifics of what could go wrong, but on your behavior and process when something goes wrong.

The project manager has many responsibilities when a bad day happens. You will depend on individual and team capabilities, training, and roles in ways that are hard to describe. You know that you must trust the team's abilities and judgment, but also watch for signs, both within the team and outside, of good intentions yielding problematic results. You must be reasonable and evenhanded, understanding that you cannot eliminate risk. The bad day is a time when a mission team shows what it is really made of. The LCROSS team earned its stripes on its bad day and through the end of what became an amazingly successful mission, redefining mankind's understanding of the moon—at a bargain price. ●

DANIEL ANDREWS has managed diverse and eclectic projects at NASA for twenty-four years, including the risk-tolerant pathfinder, LCROSS. Favorite motto: "Take calculated risks. Be willing to change course. Keep moving."