Atlas V LRO/LCROSS Mission Overview: A study in Lunar Mission Design Evolution

John G. Reed [1] and Dirk Schreier[2]
United Launch Alliance, Littleton, CO 80127

Brian Lathrop [3]
United Launch Alliance, Littleton, CO 80127

Jeff DeKruif [4]
United Launch Alliance, Littleton, CO 80127

and

Pamela Alstott [5]
United Launch Alliance, Littleton, CO 80127

Not since the Apollo era has a team worked with such creativity to achieve a successful lunar mission. While the Atlas team has had a number of novel and complex missions to build upon, this mission required a new level of integrated focus. The team developed and executed an ingenious dance to deliver both the LRO and LCROSS spacecraft into their respective trans-lunar orbits with the Centaur upper stage conditioned to properly function as the impactor for the LCROSS science mission. Years of unique and complex mission design analysis, described herein, have resulted in the unprecedented use of the Atlas/Centaur Launch Vehicle to the launch of the LRO/LCROSS mission on June 18, 2009. The team is taking the lessons learned from this experience to develop concepts to facilitate commercial use for future missions.

I. Nomenclature

FOM = Figure of Merit
KSC = Kennedy Space Center
LCROSS = Lunar Crater Observation and Sensing Satellite
LGALRO = Lunar Gravity Assist Lunar Return Orbit
LRO = Lunar Reconnaissance Observer
LSP = Launch Services Program
MMS = Magnetospheric MultiScale
RBSP = Radiation Belt Storm Probes
RCS = Reaction Control System
SOHO = Solar and Heliospheric Observatory
TIP = Trajectory Insertion Point

[1] Technical Fellow, Guidance and Navigation, P. O. Box 277005 MS B3600, AIAA Senior Member
[2] Mission Manager, Customer Program Office, P.O. Box 277005, MS 3600
[3] Staff Engineer, Trajectory and Performance, P. O. Box 277005 MS B3600
[4] Staff Engineer, Upper Stage Propulsion & Fluids Analysis, P. O. Box 277005 MS B3600
[5] Staff Engineer, Control Dynamics, P. O. Box 277005 MS B3600, AIAA Senior Member

II. Historical Missions

9th Atlas Mission to the Moon

The Atlas program has a long history of space flight. In fact, among its earliest missions were a series of lunar missions. In the early 1960's Ranger 7, 8 and 9 were successfully launched on the Atlas/Agena-B vehicle. Five Surveyor missions were successfully launched as precursors to the Apollo landings on the Atlas SLV-3C/Centaur D-1A vehicles. So NASA and lunar missions were key elements of the Atlas/Centaur development.

Increasing mission complexity

More recently the Atlas program has evolved through a series of missions that have expanded the capabilities of the Atlas system. The SOHO mission was an early foray into polynomial targeting with direct insertion into the L-1 Halo orbit. The MRO mission was the first interplanetary mission for the new Atlas V vehicle and built on the targeting techniques developed for the Titan Cassini mission. The Pluto New Horizons mission took the Atlas V one step farther with the inauguration of the Block II avionics package and its Fault Tolerant Inertial Navigation Unit providing the Atlas with a faster and more powerful brain. The Air Force STP1 program provided a final enhancement to the system with the incorporation of the Generalized Guidance algorithm which allowed for more flexible mission design, including guided out-of-plane burns. The new, more capable system was fully brought to bear to achieve the Lunar Reconnaissance Observer (LRO) and Lunar CRater Observation and Sensing Satellite (LCROSS) mission.

III. LRO Overview

Challenge of lunar rendezvous

The return to the moon resulted in new lessons for most of the team. Part of the complexity with lunar missions rests in the uniqueness of a "lunar rendezvous." With the outer planets, given the length of the cruise, the launch vehicle is simply concerned with providing the spacecraft the energy to be on its way.

For lunar missions, there is a significant increase in the trajectory design complexity. Given the sensitivity of the trans-lunar trajectory to the velocity and gravity of the moon (Figure 1), it is quite easy to miss the arrival

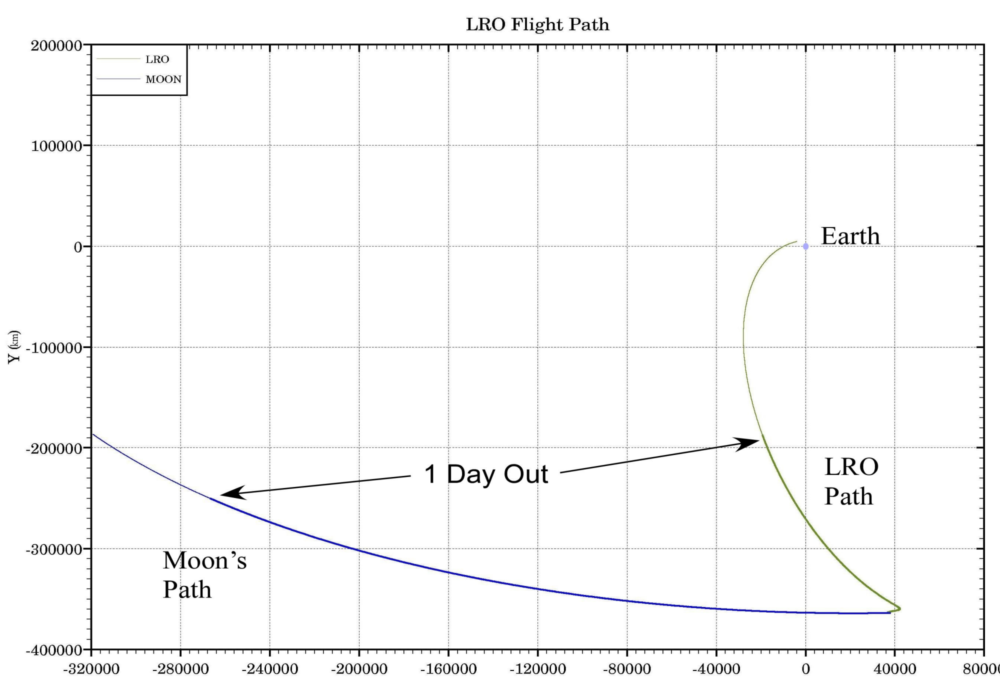

Figure 1. LRO Lunar Encounter.

conditions. In fact, given the higher order terms involved in attraction and conjunction timing, the problem quickly becomes non-linear. This becomes problematic, given the linear assumptions in many optimization schemes.

Finite Burn Effects

In fact, given the non-linear effects, it was quickly found that the finite burn effects were significant enough to prevent guidance system target generation from impulsive burn modeling. With that discovery, it became critical to the successful integration of the missions to develop a teamed, spiral targeting process where successive steps in the targeting process integrated all the way to lunar encounter. This meant that, while the arrival targets could be generated with low fidelity or impulsive burns, the Centaur targets developed using finite burn models achieved the same lunar arrival conditions. However, this also meant that all the LV simulations integrated the full trajectory to lunar periapsis.

Launch Availability

Another aspect of the lunar mission was the limited number of days every two weeks when the earth-moon-sun alignment met the LRO arrival conditions. One of the capabilities of the Atlas V system is polynomial targeting. With this capability, the target can vary smoothly across the launch window. It was expected that this capability would be used to enhance the likelihood that launch would occur sometime during a one hour window. Once the first set of targets was received, it was recognized that there were discontinuities in some of the targets across the window.

In order to maximize the likelihood of launch, the LRO team had developed targets with nearly identical performance requirements across the window. This was achieved by rotating the line of apsides as the launch time moved. However, this meant that there were significant shifts in arrival time across the one hour window. This can easily be understood by visualizing the earth rotation of 360 degrees in 24 hours while it takes the moon 28 days to sweep through 360 degrees. Thus it takes more than one day for the moon to sweep the same 15 degrees covered by a one hour launch delay.

IV. LCROSS Overview

Dual Manifest Impacts

But the complexity of the lunar trajectory was dwarfed by the complexity of the LCROSS targeting effort. LCROSS was proposed as a way to use the excess performance that the Atlas 401 configuration possessed after placing the 2000 kg LRO spacecraft on course for the moon. It was proposed as a secondary mission, but it presented the team with a series of challenges.

The LCROSS mission was constrained to be launched when the LRO wanted to launch. Initially this was not seen as significant limitation, since both spacecraft wanted to fly close to the moon on their initial pass. The initial proposal identified an additional 50 m/s delta V requirement to move the LCROSS from the LRO orbit to its own optimized orbit, termed a Lunar Gravity Assist Lunar Return Orbit (LGALRO). Since the initial flyby of the LGALRO sets up the eventual impact, LCROSS needed a specific flyby time each day. With LRO shifting the arrival time by 4.6 hours every ten minutes through the window, LCROSS needed to shift the arrival to its optimal time each day, which might also warrant a jump forward a day across the window.

This unexpected variation across the window also added a wrinkle to the sequence requirements. The initial proposal identified a requirement of up to 50 m/s of delta V. But, with the possibility of the arrival times becoming close between the spacecraft, the low-side delta V now went to zero. In fact, given the possible shift of 24 hours in arrival time from between two targets on a single day, there was the possibility of large swings in the burn attitude through the window.

Two fully constrained missions and an “empty” vehicle

In fact, since both spacecraft also had independent sets of arrival conditions across the window, the Atlas guidance system had to control the system to achieve tight tolerances in the presence of normal vehicle performance variations. The Atlas system is designed to accurately place the LRO spacecraft on its orbital path to the moon, but the point of injection (or the true anomaly) is allowed to vary under dispersed conditions. This complicates the transfer from the LRO trajectory to the LCROSS. Since the orbits may not intersect at the point of the injection burn, the guidance design needs to be constrained to preclude excessive steering to shift the position of the vehicle over to the desired orbit.

Additionally, the mission design effort was constrained by the LCROSS science requirements. LCROSS was designed to look for water in the ejecta created by the impact of Centaur. Thus there were tight requirements on how much hydrogen and oxygen could be carried to the moon by Centaur. Since liquid hydrogen and oxygen are Centaur's main propulsion source, extraordinary measures were required to fully deplete the Centaur tanks without perturbing the final orbit.

V. Mission Design Overview

Considerations

A rather elegant solution was developed to achieve the LCROSS requirements (Fig. 2). Some 1500 seconds after LRO Separation a partial blowdown of the propellant tanks was performed to provide 25 m/s in a targeted direction. Once the desired delta V was achieved, the Centaur was turned normal to the velocity vector and placed into a transverse spin, expelling the remaining liquids and depressurizing the tanks to unprecedented levels. In addition, a hydrazine burn off was performed to reduce the remaining control system fluids down to the level required for the remainder of the mission. At a fixed time after LRO separation, the hydrazine settling motors were activated to perform a low thrust, 15 minute, guided burn to achieve the LCROSS orbit. The controls system helium pressurant was vented down to minimum levels while the vehicle was oriented to allow charging with the LCROSS solar panel, to ensure successful LCROSS activation. Finally the Centaur control system was depleted, and the hydrogen tank further vented to minimize the potential impact of slow leaks of the remaining helium and hydrogen after control was passed over to LCROSS.

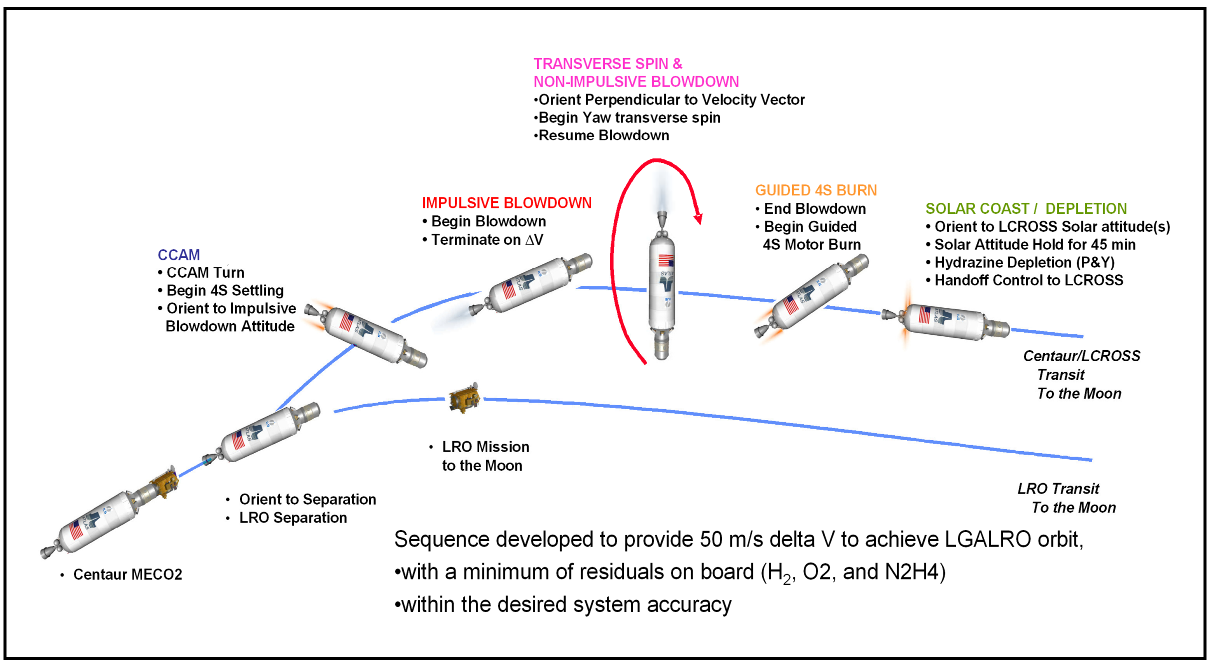

Figure 2. LCROSS Sequence Overview.

Modeling updates

To target this sequence required knowledge, not just of the nominal systems behaviors but, of the 3 sigma uncertainties in the predicted behavior and the probability density functions which best describe the system behaviors. Efforts were undertaken in a number of areas to characterize the behavior based on flight history. New models of the propulsion system behavior during blowdown and refined control system impingements were developed and implemented. Higher fidelity models of the guidance and control system interactions were also implemented in the vehicle targeting simulation. The net result was the highest fidelity simulation, post primary spacecraft separation, utilized to date and the effectiveness of this effort was demonstrated on flight day.

Test like you fly & Mission Success focus

One of the primary tools ULA uses to manage efforts like this is our test like you fly methodology. Every effort is made to ensure that our models are anchored in how we've used our systems, how we plan to use them is based on our models, and the systems are tested with how we expect to use them. As the mission design developed, any concerns with any of the tested ranges of the vehicle sub-systems were tracked. Operational ranges for time and thermal were reviewed to ensure that limitations either on the design or in testing were identified. In some cases these limitations were accommodated by refining the analysis to ensure component temperatures were adequate. In other cases additional testing was required. In all cases, it was focus on overall mission success that drove the decision process.

Verification/Qualification of systems

One prime example is the low pressure use of the Reaction Control System (RCS). Since helium is used to pressurize the propulsion tanks and the RCS system, it was critical to overall mission success to lower the He bottle pressure. This meant that the control system would have to control the attitude of the vehicle as the pressure levels dropped to untested levels. It should be noted that the flight based models did predict that the system would perform with plenty of margin. However, in order to ensure mission success a qualification test plan was developed and implemented.

Validation of Hardware

Similarly, given the attitude requirements for LCROSS charging, the vehicle was exposed to an unprecedented solar profile. Standard operation for Atlas is to roll the upper stage periodically to ensure that thermal requirements on all components are maintained. Given the mission requirements, special techniques were developed to ensure that bounding combinations of exposure profiles were analyzed and demonstrated sufficient margin for all components. While attention is always given to operational components, special consideration was also given to the other components to ensure their survival during the 120 day cruise before lunar impact.

VI. Implementation Overview

Given the mission complexities previously described, it was essential that an integrated team approach was used to achieve the mission objectives. While the need for an internal team is clear from the integrated design description, it may not be obvious that an integrated customer team was also critical to mission success. Given the number of agencies involved, it was in fact essential that all parties worked together to develop the plans, schedules, deliveries and services to achieve this launch.

Teaming with Customers – 4 way meetings

Another complexity for this mission was the integration of the four agencies in a single flight design working group. While ULA provided the basic trajectory design and flight software targeting, there were three NASA agencies which each had a hand in the process. Kennedy Space Center (KSC) Launch Services Program (LSP) had integration responsibility and as such also performed independent verification and validation of key analysis efforts. They acted as the official conduit for data flow. NASA Goddard had responsibility for the LRO mission and provided targets based on their understanding of the ULA Centaur Upper Stage capability. NASA Ames had responsibility for the LCROSS mission and provided targets derived from the LRO targets and the ICD requirement of an additional 50 m/s of delta V to be eked out of the upper stage. Key to success was the implementation of weekly teleconferences to develop the rapport to ensure issues were identified and raised quickly, rather than waiting for the periodic face to face working group meetings. Many issues, like those resulting in discrete launch points across the window, were uncovered and solutions worked off through these weekly discussions.

Multiple Targeting cycles/target timelines

Given the iterative nature of the targeting problem for these lunar trajectories, another key was laying in a schedule that allowed evolution of the targeting process. Providing three (and eventually 3 and ½ with the shift in the launch date) targeting cycles the quality and fidelity of the data products was greatly improved from simply ensuring that coast modeling for the trans-lunar trip used equivalent ephemeris data (think Jovian gravitational influences) to providing higher fidelity upper stage modeling of the finite burn effects and delaying the application point of the net impulse of the LCROSS insertion maneuver(s). The challenge in targeting was one of defining sufficient accuracy in the NASA target generation processes to ensure that the as-targeted high fidelity booster vehicle would achieve the desired arrival conditions. But building in higher fidelity adds cost both in execution time and in model certification to ensure that the model produces behaviors consistent with the on-board Atlas V flight software program. What was normally a challenge was compounded by the extremely long duration of some of the required maneuvers, like the 15 minute, 25 m/s injection "burn" using the Centaur settling motors. This maneuver is, in and of itself, hardly instantaneous but when combined with the impulsive blowdown the better part of an hour earlier makes placement of the NASA impulsive burn even harder. Similarly, developing the pointing for the impulsive blowdown and the 4S burn to achieve the targets provided from the single impulsive "burn" proved problematic. Given the long duration between the two maneuvers, simple propagation of the gravitational effects made it impossible to achieve both the position and velocity at the end of the 4S burn. With each cycle process improvements were made that simplified the targeting of the two delta-Vs so that by the final cycle, the process was down to a matter of hours for each days set of targets thus allowing time to refine the process after each targeting cycle was key to achieving both the timeline and fidelity required to successfully target the mission.

Use of Super computers

Probably the biggest oversight in the evolution of the mission was the complete independence (and in fact inverting dependence) of the targets for LRO and LCROSS. The concept of operation that was developed by ULA for targeting the mission made use of our capability for polynomial target variation through the launch window. This seemed like a reasonable approach given the fact that as time into the window advances the moon is advancing through its orbit in a very smooth and consistent manner. What was not anticipated was the influence of other mission constraints on the arrival time and the fact that each mission had its own set of unique constraints. In the end, this resulted in seven unique instantaneous target points across the window for each day. This then meant that the combined effect of the targets on things like burn duration across the window was now a discontinuous function. This greatly complicated the dispersion analysis which is used to bound the impact of target effects on a whole host of rocket systems (as well as to verify some customer duration and timing requirements). Instead of assessing a handful of points for each target block, now every target point had to be independently assessed.

Fortunately, we were able to work closely with NASA to solve this problem. Just as the magnitude of the problem (needing to execute some 1.6 million simulations) was being realized, NASA had procured a new super computer cluster system and we were granted access. With the dedicated efforts of two key engineers we were able to re-host our high fidelity flight software simulation onto the Linux based cluster. As a result, by the final cycle where new targets were being provided every two weeks (water-falling at L-120 days), the entire set of 56,000 simulations for each block were executed and analyzed before the next set started. This was a phenomenal leap forward in our ability to execute and analyze dispersed effects.

Joint review process w/customers

One final key to our success was the level of insight and participation. We had representation from all agencies at each of our program-level reviews and in fact ended up with critical side discussions to refine and improve the overall mission design at every one of them. The ability to have everyone in the room and poring over the mission plan, talking through the implications and recognizing subtle implications was paramount for a combined mission of this complexity. Another aspect of this was our ability to review the spacecraft PDR and CDR materials. We had unprecedented access and communication with our counterpart on the SV team. This meant that we could work together to develop solutions that met the desires of both systems (e.g., with telemetry coverage) instead of being constrained to the literal content of the Interface Control Document.

Thus there were many aspects of the integrated team approach that were necessary to achieve the resulting mission success. This was a great example of what a well-integrated team can actually achieve, but it did require a great deal of open communication, weekly teleconferences, bouncing issues from group to group, and relying on others to help find a better solution to optimize the team's effectiveness, not just the individual agency's bottom line.

VII. Design Details

Key to the design effort was the integrated mission success perspective. Rather than allocating requirements to specific sub-systems, the requirements were pooled and an integrated sequence developed and refined to achieve the combined requirements. Frequent discussions were held to ensure the integrated focus was maintained and issues were raised as quickly as possible.

Integrated stack/mission perspective

This integrated focus extended to the overall mission design, with the target definition extended to the lunar arrival. This shift in focus was essential to ensure that all decisions reflected the net impact on the lunar arrival conditions critical to the customer mission success. This focus also enabled early identification of system issues/impacts as well as catching assumptions and allowed rapid refinement of the mission sequence.

One key development from this focus was an early effort to compare the simulations from the various parties. As a result of this effort, the team was able to recognize some of the impacts of the customer impulsive burn assumptions. This resulted in a decision to place simulation constraints at the TCM state rather than the traditional Separation (or Handoff) plus some few number of minutes. That decision resulted in future simplification in the targeting methodology described below.

One of the early successes of this approach was the detection of conflicting window assumptions. The primary payload mission, LRO, streamlined their targeting process to allow use of the Atlas polynomial targeting which increased the launch availability and by extending the launch window. LCROSS required specific arrival times so some targets had a discontinuity when the target switched from shifting to an earlier time to shifting to a later arrival. This resulted in the decision to target seven discrete opportunities each day, to facilitate the discontinuity while maintaining the maximum launch window and saving effort to assess intermediate launch opportunities which would not have been viable to the LCROSS customer.

Ingenious design

Of course, the success of the design approach relied on detail design refinements in each of the major systems. The effort required a number of enhancements or evolutionary refinements to work. The success of this effort really rests on the basic evolutionary approach of the Atlas vehicle. By taking incremental improvements and providing significant test programs for major changes, a significant amount of data is available to define the expected behavior of each system. By leveraging this experience base, this success was achievable.

Adaptation of Development Spiral Model to mission targeting

One novel leveraging example was in the guidance targeting effort. The historical approach is to develop a simplified open loop model of the system performance for use in optimization simulations. Guidance system targets are generated from the open loop simulation and fed into a closed loop system simulation. Target biases are then refined with the closed loop model to achieve the desired separation conditions.

One of the major concerns coming out of the first targeting cycle was the time required to develop and test the guidance system targets. It was decided that improvements in the open loop simulation were required to reduce the targeting timeline to ensure that the final cycle would produce sufficient targets should the launch date slip.

Thus an additional effort was added to refine the open loop models based on the closed loop simulation and targeting bias effort. This allowed significant contributors to the bias to be included in the open loop effort. One example of this was the determination of the blowdown attitude. Initially, a handful of possible attitudes were evaluated in the closed loop and the one with the least 4S burn impact selected. By the last cycle, the open loop included sufficient fidelity to model the impact of the transverse spin attitudes and the starting attitude was included in the optimization. With this effort, order of magnitude improvements to the targeting timeline were achieved each cycle.

Blowdown modeling

Another key aspect was the refinement of the model of the blowdown of the Centaur propellant tanks. The initial effort was focused on developing a flight based model of the thrust and flowrate effect from each tank. The result was a model that bounded the flight experience. This was combined with the knowledge of the control system behavior during the second burn to define a probabilistic model of the system behavior during the blowdown (both impulsive and non-impulsive).

The result of integration with this refined model was the identification of coupling issues with the attitude control system. The combined effect of the two systems is that during the transverse spin, the vehicle will actually perform a cone maneuver, riding the switch-line, resulting in a net acceleration along the velocity vector.

Identification of this effect allowed its inclusion in the targeting process, removing a significant bias from the injection accuracy.

Guided RCS "Burn"

Another of the novel developments in support of this mission was the application of the guidance burn logic to the 4S settling thrusters. One of the early studies in support of this mission involved studying the use of such a low thrust system. Of key concern was the impact of control system firings on the net "acceleration" vector. For the Centaur system the four settling thrusters only produce 24 lbf of thrust. The pitch/yaw thrusters are capable of providing 25 lbf thrust combined in the lateral direction. While the lateral control firings are for a relatively short period of time this still results in significant "pings" to the sensed acceleration vector.

A number of minor, parameter changes to the nominal guidance system operation were made to accommodate the low thrust system. They were all based on techniques developed for other uses of the Atlas V system. These changes allowed a larger integrated effect to build up, so that the system response to small deviations in the velocity vector was reduced. They also reduced the impact of the control system firings and reduced the burn variability for effects which provided little benefit to LCROSS (a fact covered in the injection accuracy section below).

Thus by building on the techniques commonly applied in the Atlas V family, the low thrust aspects of the 4S burn were successfully characterized and mitigated. This allowed the guidance system to effectively compensate for variations in initial state and 4S thruster performance to accurately achieve the desired LCROSS inject conditions.

Controlled Depletion and Handoff

One of the hardest aspects of the design was the depletion of the hydrazine control system fluid. Given the conflicting constraints of guidance systems need for a variable 4S burn and the avionics battery life, hydrazine management took on a new level of importance. The implementation was further complicated by the LRO need for potentially three different park orbit coast lengths. What ensued was a tiered approach to hydrazine management.

The nature of the different park orbits resulted in significantly different requirements for hydrazine. For the short park orbit, we simply turn to the second burn attitude and hold that for 15-20 minutes; for the long park orbit (given Centaur's cryogenic fuel) we turn roughly normal to the sun and slowly spin the stack to balance the solar heating, reversing the spin periodically for a variety of reasons. Since no constraints were placed on the burns that could occur for a long versus short coast, the full 50 m/s capability had to be provided under the long coast usages. Since the park orbit duration could vary from day to day, the hydrazine bottle had to be fully loaded for each launch attempt, regardless of the planned consumption between the first and second burns. Thus the first tier in the management approach was to drive the hydrazine level to be consistent by the end of the second burn.

To that end a hydrazine burn-off capability was. This utilized the on-board hydrazine monitor function to target a desired cumulative consumption and fire the motors until that consumption was achieved, either due to park orbit usage or 4S burn-off during the 2nd burn. One side benefit of this was more consistent mass properties during the LRO separation event.

The second tier of the management approach was based on the guidance targeting process. Since the targets vary through the window there can be significant differences in the guided 4S burn duration from point to point in the window (basically from 200 to 900 seconds). Given this variation and the battery limit constraining the duration of the final depletion phase, it became apparent that a second burn-off period was required. To enhance mission success, this was placed late in the transverse spin (toward the end of blowing down the gas pressure in the Centaur tanks). The target was developed by removing the remaining fluid required, based on the 4S burn duration, from the full bottle. Thus the second tier allowed more control fluid to be available during the beginning of the transverse spin in case un-expected behaviors occurred and expelled whatever wasn't needed for the remainder of the missions, based on the nominal target for each point in the window.

The final tier of the management approach was to perform a balanced, hydrazine depletion after the LCROSS was ready to assume attitude control. LCROSS did require a series of attitude maneuvers and discrete commands to be performed after the 4S burn was completed. This provided a smooth power up of the systems and ensured that the control system cat bed was ready for operation before the Centaur hydrazine depletion began. The depletion logic was configured to minimize the disturbances imparted on the stack as the hydrazine was depleted.

Based on historical flight data, it was expected that the resulting torques would be very small. But even small torques can result in significant attitude errors over time, so two techniques were implemented to provide the handoff of control to LCROSS. Once an attitude error built to a parameterized level the discrete would be sent, but if the threshold had not be tripped by a certain time (well before a battery failure would result in power loss, but after the latest depletion of fluids) then the discrete would be issued. This ensured that control would be passed over to LCROSS without any loss of power due to rotation of their solar panels away from the sun.

Injection Accuracy

It was apparent from the beginning of the mission design effort that injection accuracy and especially mission variability was going to be a key challenge. Traditionally, there are two effects that impact injection accuracy: the ability of the system to achieve the target and the ability of the system to know where it is. Much of the mission design efforts described above were developed to reduce the impact of performance dispersions on the handoff or separation conditions. Typically the vehicle dispersion effects are approximately ten percent of the navigation instrument errors and the result of the design efforts is that this relationship was maintained for the LCROSS injection as well.

There were actually two separate aspects to the issue of injection accuracy for the LRO/LCROSS missions, how to describe the errors and how to compute them. Generally system accuracy is simply a matter of determining or defining how well the system can place the spacecraft into a desired orbit. For most satellites, precisely where on that orbit is of little consequence since there is some need to either phase or precess to the final earth or constellation relative location during a checkout period prior to activation. There is usually a significant period of time to trim out any injection errors at the optimal time or location for each orbital element. Thus for most satellites, three or four orbital element requirements suffice to constrain the errors that will impact their maneuvering system.

For interplanetary missions, the effort is taken a step further. Again the goal is to quantify the system error impact on the spacecraft fuel budget. Unfortunately, errors in traditional orbital elements don't really describe the impact on the spacecraft. Instead, a Figure of Merit (FOM) technique is employed. In a nutshell, the figure of merit process develops a linear transformation matrix by simulating unit changes to position and velocity to determine the impact at the planet encounter, both at separation and the correction point (called TIP) some number of days away from earth. Then by relatively straightforward matrix math, the uncertainties at separation can be transformed into the delta velocity corrections required for compensation at TIP. Since delta velocity has a linear relationship to the fuel requirements (as a function of the spacecraft thrusters) this provides an excellent representation of the spacecraft impact from launch vehicle insertion errors.

For LRO/LCROSS the fuel to correct the launch vehicle injection errors would need to be consumed before they arrived at the moon (i.e. in the first half of the delivered orbit). Mission constraints dictated that most of the correction be accomplished at 24-25 hours out. Additionally, since the moons gravity actually rips the spacecraft out of the earth orbit it was injected into, apogee timing variations can have significant effects. Finally, these effects can be non-linear, since for instance, overshooting the moon could cause the spacecraft to pass behind the moon instead of in front of it. Thus neither the typical earth orbital element nor the Figure of Merit approach allows good characterization of the impact on the spacecraft fuel.

The new methodology developed for LRO/LCROSS is a hybrid method again reflecting the evolutionary methods. The approach involved diagonalizing the traditional state covariance matrix. The six eigenvectors are then individually simulated and the dispersed TCM optimized for the arrival conditions. These values are root sum squared to define the launch vehicle impacts on the TCM.

VIII. Results/Summary

While there was still an incredible amount of work and analysis to develop the final targets, with the incremental, integrated approach to the development the team was able to waterfall targeted closed loop trajectories back to NASA as quickly as LCROSS was able to provide final targets (basically a block every two weeks).

On June 18th 2009, LRO was successfully injected into LT orbit. LRO only used 1.3 m/s of 20 m/s allocation to correct for LV injection errors. Similarly, LCROSS was successfully injected into its LGALR Orbit. LCROSS only required 8 m/s out of the 30 m/s allocation. This value also reflects a decision by the LCROSS team to change to more optimal fly by condition after final targets were delivered to ULA. Thus three quarters of the correction was due to this target change, and only 2.5 m/s would have been used out of the 30 m/s requirement. It is also worth noting that the ability to launch on the initial block opportunity also preserved more than 75 m/s of correction LCROSS had reserved for the later opportunities (i.e. the arrival condition in the next block of days required more than the 50 m/s of adjustment from the LRO orbit provided by Centaur). This provided more than enough LCROSS control fluid to accommodate the issues encountered in flight.

Model estimates predicted vehicle behavior

ULA post-flight analysis found excellent performance compared to pre-flight predictions. The delta-V provided during the LCROSS mission phase was 20.57 m/s versus 20.64 m/s targeted. The blowdown duration was 205.4 seconds versus the preflight prediction of 210 seconds. The pointing behavior during the transverse spin was exactly as predicted. The delta-V imparted during the transverse spin (exclusive of the impulsive blowdown) was 5.5 m/s versus the nominal prediction of 4.2 m/s. As a result the 4-S burn only needed to impart 7.7 m/s versus the preflight prediction of 8.7 m/s.

As a result of late LCROSS analysis of possible Centaur leak rates, an additional balanced vent of the hydrogen tank after handoff reduced the tank pressure to less than 0.9 psia. This represented less than 0.1 lbm of hydrogen left at the beginning of the LCROSS cruise and matched preflight prediction to less than 0.1 psia. The one unexpected behavior was in the torque from the low pressure flow through the oxygen tank fill and drain valve after handoff. While this resulted in higher than predicted disturbance torques, it also decayed much more quickly than predicted.

IX. Future Opportunities

The knowledge gained in the execution of the LRO/LCROSS mission directly enhances the capabilities and preparedness of ULA to execute future missions. The enhanced understanding of the blowdown forces is being used to refine the final orbit of all Centaur missions to mitigate conjunction concerns with all active satellites. The lessons learned from the mission unique Flight Software enhancements are being pursued in support of future multi-spacecraft missions. The delta-V capabilities are being considered for new and challenging missions like NASA's Magnetospheric MultiScale (MMS) and Radiation Belt Storm Probes (RBSP) missions. With the flight experience and detailed system knowledge ULA possesses, we are well positioned to extend our reliability record and flight expertise into a myriad of commercial ventures in support of NASA's future missions.