An Altair Overview - Designing a Lunar Lander for 21st Century Human Space Exploration

Kendall K. Brown, PE, PhD
NASA, Marshall Space Flight Center; USA; kendall.brown@nasa.gov

John F. Connolly, PE*

Altair, the lunar lander element of NASA's Constellation program, was conducted in a different design environment than many other NASA projects of similar scope. Because of this relatively unique approach, there are a number of significant success stories that should be considered during the development of any future lunar landers or human spacecraft. This paper is divided into two separate themes; the first is the approach used during the conceptual design studies, including the systematic analysis cycles and the decision making process associated with each: and the second is a summary of the resulting lessons learned that were compiled after looking back at the lifetime of the Project. Altair was terminated before entering Phase B of its design, and was often criticized for being a very heavy and very large vehicle. While there was specific rationale for all of the decisions that led up to that configuration, future design cycles were specifically planned to re-address the mass challenge. Had the project continued, the deliberate, stepwise design process would have converged on an optimized lander design that balanced mass, risk, cost and capabilities. Some of the specific items that will be addressed in this paper include project development strategy, organizational approach and team dynamics, risk-informed design process, mission architecture constraints, mission key driving requirements, model-based systems engineering process, configuration studies, contingency considerations, subsystem overviews and key trade studies. The paper will conclude with a summary of the lessons identified during the Altair project and make suggestions for application to future studies.

I. THE ALTAIR LUNAR LANDER DESIGN PROCESS

Designing a new human lunar lander is a multi-layer systems challenge. The Altair Project created a lander design that responded to the physics of spaceflight and limitations of human performance…while as a project balancing performance, cost, schedule and risk...while working within the integrated architecture performance, cost profile, schedule and integrated risk and reliability targets of NASA’s Constellation program…while fulfilling the policy directives of NASA’s strategic plan, Congress’ NASA Authorization Acts, and Administration/OMB/OSTP policy…and performing within budget guidance. This required a team with a true systems perspective – an understanding of how all the pieces fit together, at all levels.

A lunar lander is a physics machine. The physics of lunar landing demands that the lander perform velocity changes - ~1000 m/sec to decelerate into lunar orbit, ~2000 m/sec to decelerate to a soft landing, and another 2000 m/sec to accelerate back into lunar orbit. Additionally, a lander must include life support to provide the physiological environment for the human crewmembers. Resultantly, much of the lunar lander “design space” is fixed by physics – large tanks of propellant surrounded by structure, an attenuation system for landing, and a pressurized volume for crew habitation. The designers of the Apollo mission understood the physics of the problem perfectly, even though they were inventing much of it for the first time. The Altair team’s challenge was to apply the lessons learned from Apollo, combined with the incremental improvements to technology from the past 4 decades. Not surprisingly, Altair bears some resemblance to the venerable Apollo Lunar Module – the physics of lunar landing is unchanged, the planform of the lander will reflect its functionality, and the technologies that have improved most dramatically (computers, avionics, composite structures) are mostly invisible at the vehicle level. So Altair looks like the big brother of the Apollo lunar Module – not because the Altair team wanted it to, but because the Apollo lunar module designers were pretty smart and understood the physics as well. Figure 1 illustrates the similarities and differences between the Apollo LM and Altair.

---

* John C. Connolly, NASA Johnson Space Center; USA, john.connolly-1@nasa.gov

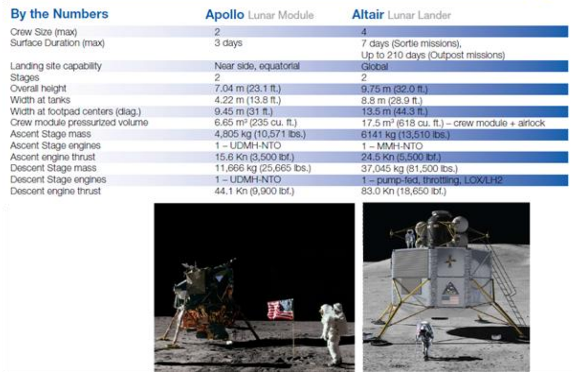

By the Numbers	Apollo Lunar Module	Altair Lunar Lander
Crew Size (max)	2	4
Surface Duration (max)	3 days	7 days (Sortie missions), Up to 210 days (Outpost missions)
Landing site capability	Near side, equatorial	Global
Stages	2	2
Overall height	7.04 m (23.1 ft.)	9.75 m (32.0 ft.)
Width at tanks	4.22 m (13.8 ft.)	8.8 m (28.9 ft.)
Width at footpad centers (diag.)	9.45 m (31 ft.)	13.5 m (44.3 ft.)
Crew module pressurized volume	6.65 m³ (235 cu. ft.)	17.5 m³ (618 cu. ft.) – crew module + airlock
Ascent Stage mass	4,805 kg (10,571 lbs.)	6141 kg (13,510 lbs.)
Ascent Stage engines	1 – UDMH-NTO	1 – MMH-NTO
Ascent engine thrust	15.6 Kn (3,500 lbf.)	24.5 Kn (5,500 lbf.)
Descent Stage mass	11,666 kg (25,665 lbs.)	37,045 kg (81,500 lbs.)
Descent Stage engines	1 – UDMH-NTO	1 – pump-fed, throttling, LOX/LH2
Descent engine thrust	44.1 Kn (9,900 lbf.)	83.0 Kn (18,650 lbf.)

Figure 1 Comparison of Apollo LM and Altair

As the Altair project team began the Concept Studies portion of Project Formulation, what NASA defines as Pre-Phase A, it built upon the lessons and experiences of Apollo, the Space Shuttle, the International Space Station, and the early phases of the Orion project. NASA management challenged the Altair Team to find a way to develop the next human lunar lander more efficiently; to develop a lunar lander with the lowest reasonable mass, with high reliability and safety, at a low total cost and to meet the declared objective of returning humans to the moon by 2020. The Altair Team evolved, combined, and developed a number of different design and team integration practices to meet that challenge. As a result, many of the lessons identified during the Altair Project should be incorporated into the design of any future crewed lunar lander or human spacecraft design.

This paper will provide an overview of the Altair design process and then discuss the primary lessons the Altair Team identified during Pre-Phase A and early Phase A (Concept and Technology Development). The team identified lessons associated with how a multi-discipline, inter-center team should function; how to implement a model-based systems engineering process; how to use a risk-informed design process at the beginning of a project that is based upon detailed engineering studies; and how to incorporate contractor input into the project formulation phase.

II. ALTAIR SYSTEM DESCRIPTION

Altair consisted of four major components: an Ascent Module (AM), a Descent Module (DM), and an Airlock, as shown in Figure 2, and the Ares V Earth Departure Stage/Altair Adapter (EDSA). The AM is built around a crew cabin that serves as the primary habitable volume for the crew during at least the descent and ascent phases of the crewed sortie mission, and provides pressurized access to the Airlock and Orion. Altair's DM main function is to deliver hardware (AM with crew, Airlock, cargo) to the surface of the Moon, and is built around the descent propulsion system, landing gear, and structure necessary to carry loads through all flight phases. The Airlock module provides ingress/egress access to the Altair AM in the Sortie mission mode. The Airlock allows the crew to perform split operations (e.g. two crew members perform EVA while two crew members remain in the Altair) and serves as one of the primary mechanisms used to control the transport of lunar dust into the AM. Different combinations of these modules yield 3 separate configurations of the lander for crewed sortie missions (AM+DM+airlock), crewed outpost missions (AM+DM) and dedicated cargo delivery (DM only).

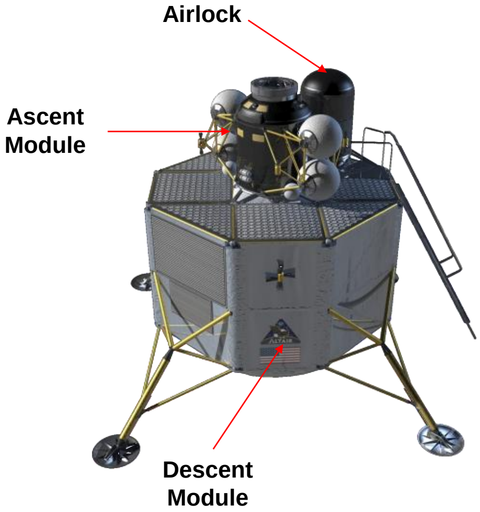

Figure 2 Altair p905-A Configuration

This initial design was chosen with some thought. Over the past decades, NASA has well over 100 discrete lander designs that explored the trade space variables of staging, delta-V split, propellant types, crew size, surface duration, launch configuration, landing configuration, launch shroud dimensions, abort capabilities, crew access, cargo accessibility and offloading, C.G. control, and a myriad of other, often competing, design drivers. Based upon this history of studies, the Altair Project selected a configuration to begin its risk-informed design process – a two stage vehicle using a large, efficient LOX/LH2 stage for both landing and a piggyback LOI burn with the Orion vehicle attached, a small, lightweight ascent stage for crew habitation and "flight deck" functions, and a separate airlock.

This choice of initial configuration was a necessary starting point to initialize the risk-informed design process, and has been held constant to facilitate the design process. Freezing a design early in the process does create a risk that this initial (likely non-optimal) design choice becomes confused with THE ultimate design. However, an important step that was been inserted into the Altair design process is to periodically "step back" and re-evaluate the vehicle configuration to assess if the team is pursuing the most optimum design.

III. RISK INFORMED DESIGN

The Altair project is using a design approach that is unique to NASA's human spacecraft. A typical NASA project first begins with a set of requirements that describe the entirety of the functions and performance a spacecraft must possess, and a vehicle is then designed to satisfy all of these requirements. This process results in a design that initially attempts to meet all requirements equally, and from which it is difficult to extract capability if the vehicle is found to exceed mass or cost limitations. Altair's approach was to first design a vehicle that meets only a minimum set of requirements, and then incrementally add functions and performance back into the design. This approach allows the decision to accept each additional requirement to be informed by its individual impact to cost, performance and risk. This process was derived in part from NASA Engineering Safety Center Report NESC PR-06-108[1], "Design, Development, Test and Evaluation (DDT&E) Considerations for Safe and Reliable Human-Rated Spacecraft Systems".

After defining the "minimum functional" vehicle in the first Lander Design Analysis Cycle (LDAC-1), subsequent design cycles identified major risks that affected the safety of the crew (LDAC-2), and the success of the mission (LDAC-3). By using risks to inform these early design cycles, the Altair Project was able to identify the specific performance "cost" of each increment of crew safety and mission reliability that was added to the minimum spacecraft design. Residual spacecraft risks continued to be re-evaluated as subsequent design cycles assessed the performance, cost and risk impacts of adding additional vehicle functionality, and other factors such as manufacturability and maintainability.

The first step of the process is to establish a "minimum functionality" baseline design. This requires that the design team scrub the vehicle requirements back to a small number that described the essential functions and constraints of the lander. For the Altair lander these "core" requirements were to carry 4 crew to the surface for 7 days with 500 kg of payload, to loiter for up to 210 days at a polar outpost, to deliver 14,500 kg of dedicated cargo, to package within the Ares V shroud, to perform the lunar orbit insertion burn with the Orion spacecraft attached, to carry an airlock, and to work within the Constellation EOR-LOR architecture. Key constraints were control masses of 45,000 kg for crewed missions and 53,600 kg for cargo missions.

From this minimum set of requirements, the design process was begun. "Minimum Functionality" is a design philosophy that begins with a vehicle that will perform the mission, and no more than that. It does not consider contingencies nor provide any added redundancy, and is approximately equal to a "single string" design approach. A "minimum functionality" vehicle is NOT a design that would ever be contemplated as a "flyable" design! What this design philosophy did provide was early, critical insight into the overall viability of the end-to-end Constellation transportation architecture - if a transportation architecture cannot "close" with a minimum functional lander, it will certainly not close when all the additional functionality is added back into the lander design. More importantly, the minimum functional lander design provides a starting point to make informed cost/risk trades and consciously buy down risk.

Design standards were also scrutinized in formulating the minimum functional design. Existing NASA standards on redundant systems were put aside for the initial design, and were used only as one possible risk mitigation option in later design cycles. Technical standards were individually scrutinized – for example, the initial design used the nominal design standard structural factors of safety of 1.5 (and 2.0 for pressure vessels), but left these open to trade during later design cycles.

This initial design cycle was completed in two months using a "collocated" team of engineers. The LDAC-1 design that resulted from the initial design cycle is shown in Figure 3. Though not a "flyable" vehicle, this design provides a starting point for informed risk reduction design cycles that were to follow.

III.1 Altair Lander LDAC-1 ("Minimum Functional") Design

The primary LDAC-1 design figure of merit was to maximize the residual payload to the surface of the Moon with a crewed Lander. Large payloads landed with crewed missions were being investigated as an option to incrementally building lunar surface capabilities, and Constellation studies sought to know the maximum payload that could be delivered with lunar crews. One of LDAC-1 results was that a "minimum functional" vehicle (illustrating the extreme of maximizing delivered payload) could deliver less than 4 mt to the lunar surface. From this, the Constellation program concluded that small payloads could be delivered with crewed landers, but a cargo variant of the lander would be needed to build up a lunar outpost.

The result of the initial design cycle was a bottoms-up design of a "single string" vehicle that met all the fundamental design reference missions and requirements, but no more. Each subsystem provided detailed engineering analysis and bottoms-up design. Each then provided equipment lists, schematics and CAD models to Altair's Integrated Vehicle Performance team, who assembled the products that describes the overall lander's performance characteristics: A Master Equipment List listing over 2000 individual components, a Powered Equipment List, an integrated vehicle schematic, an integrated vehicle consumable and resource utilization analysis, and a detailed CAD model.

III.2 Design Analysis Cycle 2: Buying back crew safety.

The LDAC-1 design provided the baseline from which to identify vehicle risks in order to mature the design from one that was "minimum functional" to one that was "safety enhanced". Risks that contributed most directly to the Loss of Crew (LOC) were first identified, and mitigation options then studied for these risks. Decision processes were developed for both selecting the LOC risks to be studied, and evaluating the mitigation options that would be incorporated into the LDAC-2 design. For this initial risk reduction cycle, the primary measures were mass and change to risk.

Altair's risk analysis team was key to the success of this cycle of risk-informed design. NASA safety personnel first identified crew safety risks and prioritized them so the design could be matured to increase the likelihood of crew survival. Beginning with the minimum functional LDAC-1 design, two lists of risks were developed, one using a top-down reliability model and the other using bottoms-up subsystem and vehicle fault trees and hazard analyses.

All risk inputs were referred to a Risk Prioritization Team (RPT), comprised of Safety and Vehicle Engineering personnel, and Altair's Chief Engineer. This team took on the complex task of synthesizing the results from the Lander Hazard Analysis, Lander Reliability Tool, and Subsystems Single Point Failure Assessments. From these inputs, the team compiled top composite risks, and created task sheets detailing 28 individual studies that were assigned to the Altair team.

Subsystem risks, vehicle-wide risks such as micrometeorites radiation threats, and trajectory dispersions were all identified and mitigation measures studied. Additionally, aborts were "bought back" into the design as an additional mitigation against LOC. The end result was an "expenditure" of almost 2000 kg of mass (including both dry mass and propellant) for a 1.5 order of magnitude decrease in LOC probability.

LDAC-2 results are plotted in Figure 3. The probability of Loss of Crew is read from the stacked bars with the scale on the left Y axis, and the change of mass "available for payload" is plotted by a blue line using the Y axis on the right. The stacked bars are further broken down to show the contribution of individual subsystems to the overall LOC metric. The composite of all decisions made in LDAC-2 to

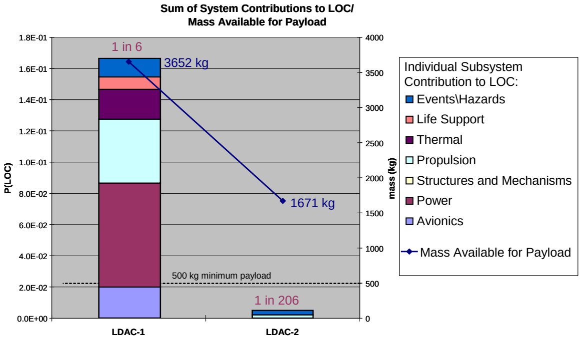

Figure 3. LDAC-2 Summary Metrics – Probability of Loss of Crew, Mass Available for Payload

reduce Altair's LOC probability resulted in the mass available for payload being reduced from 3652 kg (in the minimum functional, single-string LDAC-1 lander design) to 1671 kg. This still exceeds the 500 kg of payload required for the lander to deliver, but does not yet include the "buy back" of Loss of Mission (LOM) risks or the addition of additional capabilities. LOC risk was improved from approximately 1 in 6 (LDAC-1) to 1 in 206, which begins to approach the 1 in 250 requirement for Altair lander LOC.

Analyzing Loss of Crew risks, identifying mitigation options, and choosing design solutions that optimized risk buydown and mass performance gave rise to a number of useful observations. Most notably, full redundancy was usually the most massive and frequently NOT the most effective option for improving LOC. Analyzing options other than full redundancy, however, adds technical challenge and consumes a greater amount of effort than applying simple design "rules of thumb". A bonus is that the design team ends up much more intelligent on risk and design drivers. Another lesson learned is that one or more quantitative risk tools are necessary to inform good design decisions. The Altair team combined PRA-based lander reliability model with tops-down Fault Trees and bottoms-up Single Point Failure identification to assess the breadth of risks. Finally, it will always be necessary to correlate engineering judgment with the results that risk tools produce. The Altair team did not rely solely on tool results, but used the results to focus technical discussion of specific risks. During the analysis of specific risk mitigations, designers also sought to understand the analytical risk modeling when a result did not correlate with their engineering experience. The risk tools ultimately become an aid that the designers interacted with, and each design cycle improves both the tools' calibration and the designers' understanding of design and risks. Ultimately, design for Minimum Risk proved valuable in building a smart design team.

III.3 LDAC-3: Buying back mission reliability

For the third Analysis Cycle analyzed safety and reliability design changes that target the highest Loss of Mission (LOM) risks residual in the LDAC-2 lander design. A Risk Prioritization Team, similar to that used in LDAC-2, was tasked with synthesizing the outputs of an updated Fault Tree, Lander Reliability Model, and subsystem-identified Single Point Failures as shown in Figure Y. From these sources, the RPT identified the fundamental LOM risks and created analysis tasks that were assigned to either Altair subsystems or Integrated Vehicle leads. Each of these tasks will result in decision packages that present options for "buying back" each LOM risk using redundant systems, dissimilar redundancy, highly reliable components, increased testing, and other methods to decrease risk.

In addition to LOM risks, LDAC-3 also incorporated "global access" functionality decisions made in collaboration with Constellation Lunar Architecture Team (CxAT-Lunar) transportation architecture studies. Extensive sensitivity studies were performed to determine the combined effect of launch vehicle capability, lander LOI delta-V, LLO loiter duration, and lunar global access coverage. These studies concluded that Constellation program global access requirements could be satisfied with a combination of 4 additional days of LLO loiter, and by sizing Altair's tanks for an LOI maneuver of 1000 m/sec (though the tanks would be filled only to 950 m/sec LOI for the majority of missions).

Loss of Mission risk mitigations and global access capabilities will be incorporated into the vehicle closure segment of the LDAC-3 design cycle, along with improved subsystem and spacecraft design maturity.

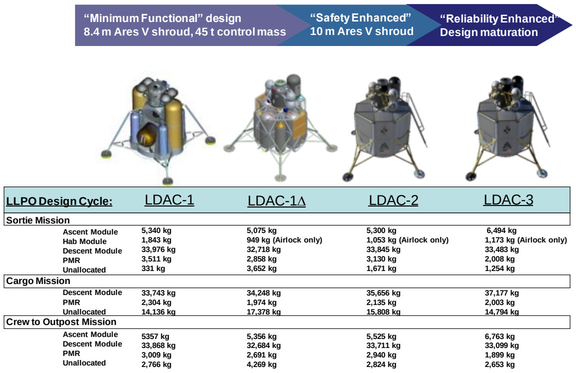

LLPO Design Cycle:	LDAC-1	LDAC-1Δ	LDAC-2	LDAC-3
Sortie Mission				
Ascent Module	5,340 kg	5,075 kg	5,300 kg	6,494 kg
Hab Module	1,843 kg	949 kg (Airlock only)	1,053 kg (Airlock only)	1,173 kg (Airlock only)
Descent Module	33,976 kg	32,718 kg	33,845 kg	33,483 kg
PMR	3,511 kg	2,858 kg	3,130 kg	2,008 kg
Unallocated	331 kg	3,652 kg	1,671 kg	1,254 kg
Cargo Mission				
Descent Module	33,743 kg	34,248 kg	35,656 kg	37,177 kg
PMR	2,304 kg	1,974 kg	2,135 kg	2,003 kg
Unallocated	14,136 kg	17,378 kg	15,808 kg	14,794 kg
Crew to Outpost Mission				
Ascent Module	5357 kg	5,356 kg	5,525 kg	6,763 kg
Descent Module	33,868 kg	32,684 kg	33,711 kg	33,099 kg
PMR	3,009 kg	2,691 kg	2,940 kg	1,899 kg
Unallocated	2,766 kg	4,269 kg	2,824 kg	2,653 kg

Figure 4. Altair Project Lander Configuration and Performance Maturation thru LDAC-3

III.4 Requirements Analysis Cycles (RAC) 1 and RAC 2.

Following the series of analysis cycles that stepwise followed NESC PR-06-108[2], the Altair Project Office undertook two RACs in order to mature the Altair requirement set and assess which ones were truly necessary versus those that required further scrutiny. Although no major changes were made to the vehicle during these cycles, a variety of maturation studies were pursued in order to achieve a better understanding of vehicle performance. These resulted in potential changes that were then decided upon in LDAC-4 to begin reassessment of the vehicle's overall configuration.

This period of time was also used to develop design detail behind sixteen major variants of the DM main propellant tank configurations. Of these sixteen tank options, the team selected three tank configurations for which to develop structural designs, due to favorable assessments of tank mass, propulsion system operability, and tank manufacturability. Using these three tank options, the structures team developed five Descent Module structural options. At the end of RAC-2, two of these options were selected for further consideration as a part of future vehicle-level trade studies, in addition to the baseline configuration, with one option being put on hold pending further discussions with the Ares V team.

Finally, the safety and risk assessment team continued to refine the vehicle's quantitative risk model during this period. As a result, while no material changes were made to the design, the vehicle's loss of crew risk posture decreased from a value of 1 in 256 to a value of 1 in 277 and the loss of mission risk posture remained constant at 1 in 22.

III.5 LDAC-4: Incorporating Requirements and System System Definition Trades

The focus of LDAC-4 was to include requirements in the vehicle design that had not previously been satisfied, but were ready to be implemented as a result of the maturation efforts of RAC-1 and RAC-2. These requirements spanned the following areas:

- Global Access Coverage
- Contingency EVA Capabilities
- Life Support Cabin Air Monitoring, Contingency Supply, and Vestibule Pressurizations
- Crew Personal Protective Equipment
- High Definition Video transmittal, Data Storage, and C3I Protocols
- Increased geometric antenna coverage and simultaneous links
- Ascent Module Disposal
- Potable Water (Hot and Cold) Performance

In addition to new requirements, vehicle maturation changes were worked into the vehicle design. Some of the most significant changes were as follows:

- Command and data handling system architecture was switched from a centralized design to a distributed system design.
- The landing loads and dynamics were assessed with higher fidelity models
- A line-item budget was established for the DM propellant.
- Boil-off calculations for the liquid hydrogen in the DM main propellant tanks were refined.
- The propellant scavenging hardware used to scavenge gases from the DM propellant tanks for use by the fuel cell and the crew was reassessed.
- Improved modeling and line sizing of the DM pressurization system led to a 110 kg decrease in the threat that was being held for propellant tank imbalance.

As a result of all of these changes, the Sortie vehicle's total expected vehicle mass decreased by 2,311 kg in the bottoms-up lander design. The addition of new requirements also had a negative effect on the baseline vehicle's safety and reliability.

IV. SPACECRAFT DESIGN LESSONS

In 2010, the Constellation program was cancelled, and a "lessons learned" identification process was conducted as part of the Program's closeout activities. Inputs were solicited in the categories of Management, Systems Engineering & Requirements, Organization, Communication, Resources, Technical Authority, Planning, Manufacturing, Test and Verification, Design and Development. The following is a summary of the Altair Project's inputs to the lessons learned activities. These lessons form an excellent basis for any new human spacecraft program beginning Pre-Phase A and Phase A conceptual design activities.

IV.1 Risk-informed design should be started during conceptual design

Risk-informed design provides early, critical insight into the overall viability of the end-to-end architecture, and provides a starting point to make informed cost/risk trades so that risks can consciously be bought down. The Altair team has used the education afforded by risk-informed design to look at risk reduction in its many forms and not to blindly apply fault tolerance rules or preconceived risk reduction solutions. The process inherently produces risk metrics for each added capability, and cost analysis can easily be added to facilitate evaluation of the true cost and risk changes that accompany each added capability. Perhaps most importantly, risk-informed design creates a true "Smart Buyer" team that inherently understands the balance of risk drivers and mass performance within the design.

The initial design analysis cycle (DAC-1) for the Altair Lunar Lander provided a spacecraft with "minimum functionality." Minimum function was defined to mean that the lander was designed to meet the primary key driving requirements, not the complete set of Level II Lander requirements. Primarily, it meant that the vehicle and it's subsystems were designed for zero fault tolerance and no contingency operations. It was understood that this design did not represent a flyable vehicle, rather this provided a theoretical vehicle that could perform a lander mission if everything worked with 100% reliability, obviously an unrealistic design point. However, during DAC-2 the risks that would result in Loss of Crew (LOC) failures were identified and specific vehicle, subsystem, and component trade studies were performed to identify the reliability improvement as a function of the mass of the alternatives. The resulting vehicle concept was referred to as the Increased Safety vehicle. This process was repeated during DAC3 for the Loss of Mission (LOM) risks to provide the Enhanced Mission Success vehicle. A tool was developed using probabalistic risk assessment component failure data, including data from the Space Shuttle and ISS PRA, to quantify the risk associated with the subsystem and component alternatives. The term "risk-informed design" is used because the design decisions were literally" informed" by the quantified risk information instead of using a rule-based decision process. By using this process, the Project Manager was able to systematically add back safety, reliability, and capability with a more complete understand of the integrated effect on the spacecraft. Implementing this type of process early in the design process is crucuial because changing the design later to reduce mass is difficult and expensive.

Risk-informed design works best when the configuration of the spacecraft is held (steady), so as not to introduce additional variables into the design. It is also a time consuming process (the first 3 design analysis cycles took the Altair team approximately 24 months to complete), and during that time requirements may change, interfaces with other elements may become better defined, and the lander design itself will mature. To (focus/ best suit/optimize) the risk-based design effort, the Altair team chose to hold the vehicle design constant throughout the design cycles, with a plan to revisit vehicle configuration once LOC and LOM "buyback" cycles were complete. With the completion of the risk and reliability design cycles, the next step is to prioritize the configuration and maturation studies that would have the greatest impact on the vehicle design. Altair considered a list of over 200 potential configuration/maturation trades, and from that list chose the following studies as the basis for a special Trade Analysis Cycle (TAC) that was inserted into the vehicle's development schedule.

IV.2 A multi-center in-house Skunkworks® team is a good way to initialize new projects

In late 2006 a study was performed to identify options for initiating development of the lunar lander to meet the Human Lunar Return by 2020. That study determined that insufficient time and budget were available to execute a lunar lander development project using a standard NPR 7120.5 process with contractors performing pre-phase A and phase A/B studies. The Lunar Lander Project Office was stood up to implement more streamlined and efficient path through the project formulation phase. The key tenets of the approach included: 1) NASA Administrator buy-in to the approach and broad lattitude to deviate from NASA policy, as needed; 2) began with a small, hand-picked team made up of spacecraft and subsystem design experts from across the agency; 3) developed and communicated a set of guiding principles; and 4) began the project by co-locating the mult-center team in a single facility for the first 2 months of the project, then maintained that cohesiveness with regular short-term co-located meetings.

Altair's experienced showed that a small, focused multi-disciplinary in-house NASA team is a very effective way to initiate the formulation phases of new projects.

IV.3 Model-based systems engineering should be used to provide mission functional modeling for requirement decomposition.

The Constellation program initiated the requirements development process by determining the capabilities needed from each of the systems in the architecture to accomplish the mission, and levied the requirements through various documents, including the Constellation Architecture Requirements Document (CARD). As a project within a large program, it is the project's responsibility to implement a more detailed assessment of the mission to validate the requirements levied on it from the program and determine if they are necessary and achievable, as well as to scrutinize the mission for latent requirements. The Altair Project utilized Models-based Systems Engineering (MBSE) as the approach to requirements development. MBSE focuses on building data models that clearly depict the operational flow of the mission. These hierarchical models manage the complexity of the mission and functional allocations, and make excellent integration products for a group review between different organizations for consistency in assumptions. The operational models can then be analyzed to determine the functionality the vehicle must possess to enable these operations. The models are kept in the requirements database for the program, and cross-references between the vehicle functions and the mission phases the functions are used in serve as validation of the function. Further, linkages between the functions and requirements that enforce them on the design provide complete traceability from requirements to the concept of operations. Finally, by assigning durations to the activities within the operational models, they can be simulated to provide a complete timeline of the mission from the same models that are being used to establish vehicle functionality. By developing these common products, a single model set becomes the authoritative source for the requirements, functionality and operations, thereby improving the quality of the requirements, integrating the various organizations within the projects and minimizing disconnects.

Models-based systems engineering streamlines requirements development and validation by developing integrated products between the Requirements, Design and Operations communities, and should be used to provide an integrated set of products upon which architecture, element, system and subsystem requirements can be sequentially decomposed.

IV.4 Regular co-locations are essential to using geographically distributed teams

The diversity provided by a geographically distributed (in the case of Altair, a multi-NASA-center) team is worth the extra effort it requires. While teleconferences and internet assisted virtual meetings are required, the key to operating with a multi-center team is to periodically bring the team together to work in a single location. When Altair began as the Lunar Lander Project Office, it co-located approximately 60 people in a single small building at Johnson Space Center for a period of 8 weeks. This period allowed the team to get through most of the forming, storming, norming, and began performing as a high performance team before returning to their home centers and organizations. The project maintained the team cohesiveness by planning a robust travel budget that allowed a significant part of the team to get back together for three or four days of jointly working together. These periodic co-locations were rotated amongst the centers; and to the extent possible, the team stayed in the same hotels, went to dinner together, and spent time socializing. This approach ensured that the personal relationships that were initially formed were maintained, and that as new members joined the team, they were more quickly assimilated. The interpersonal bond towards the Lunar Lander team was tested as inter-center institutional competition for roles and responsibilities emerged, however the strength of the team prevented it from getting in the way of the work.

IV.5 Detailed design during pre-phase-A and early phase-A can identify important issues

The cost and schedule associated with performing significant engineering design during the project formulation phase provides a tremendous return on the investment. One of the key initial tasks for the Altair Project was to develop a preliminary in-house design within six to nine months, and the emphasis was to focus on performing significant design, i.e., focus on the "D" in LDAC. The detailed design work allowed us to validate and identify weaknesses in the parametric modeling. A few notable examples can illustrate this point. The mass of the landing legs could not be easily scaled from the Apollo Lunar Module due to the dramatic size difference between it and the hydrogen-fueled Altair. By performing detailed analysis to determine the required size for stability at the bounding landing site terrains, and then developing and analyzing a detailed design, we were able to improve the mass confidence. As the team used the detailed design and began identifying assembly, integration, and test during both development and operations, we identified that the Constellation Program decision made early in LDAC-2 to increase the descent module diameter to take advantage of the 10m Ares V shroud had significant implications for transportation and thermal-vacuum propulsion system testing. That assessment found that the Altair descent module could not be transported in the NASA Super Guppy aircraft and would require barge transportation. To perform an acceptance test of each descent module propulsion system in the Plumbrook Station B2 test facility, the barge would require an ice breaker to clear a path through Lake Erie during the winter. These examples represent some of the more dramatic items that were revealed by allowing the NASA in-house team to perform significant detailed design work during the project formulation phase.

Detailed design during pre-phase-A and early phase-A can identify performance characteristics that cannot be parametrically modeled and other important issues that cannot be determined until a design concept is of sufficient maturity to evaluate a more complete set of design, development, test, and evaluation considerations.

As early as possible in the project formulation phase, a detailed design concept should be developed to allow a more complete understand of the performance and programmatic implications of the design. The concept needs to be more than an artist rendering based upon parametric design characteristics, it needs to have engineering design analysis substantiating it so that a design team is made of a subject matter experts with the experience to foresee the potential issues during DDT&E.

IV.6 Industry should be brought into the process early.

One of the early tenets of the Altair Project Office was to solicit input from industry throughout the project formulation phase. In June 2008, a little over a year after Altair was created, the office awarded Broad Area Announcement (BAA) contracts to five companies; three major traditional aerospace companies and two small, independent companies. The BAA had two primary objectives. The first asked the companies to review the government's Lander Design Analysis Cycle (LDAC-1) minimum function lunar lander conceptual design, the plans for implementing risk-informed design, and to provide suggested alternative design concepts. The second asked the companies to provide recommendations for the roles and responsibilities of government and industry for the development of the lunar lander. A process was then set up whereby the companies that participated in the BAA and other entities meeting (International Traffic in Arms (ITAR) export control regulations could continually obtain updates on the government design. This two-way interaction was extremely valuable to both the government and industry – it provided the government some important alternative perspectives and it provided the industry information which helped them focus their internal investment funding. The input from industry also helped shape the roles of the government and industry that were incorporated into the Altair acquisition strategy. As the Altair project proceeded into Phase A concept exploration and refinement, and while preparing for the System Requirements Review (SRR), it developed a process whereby multiple contractors would be integrated into the Altair Team to support the work. This process was being executed using a fixed cost procurement titled Altair Conceptual Design Contract (ACDC); this contract was ready for award in the Spring of 2009. This process would have maintained government leadership of the design until a prime contractor was selected for the flight vehicle development somewhere in the post-SRR timeframe, enabling the government team to develop a better Request for Proposals solicitation, and allowed all contractors greater in-sight into the information that the government was using to shape the requirements. In fact, during the Heavy Lift and Propulsion Technology Request for Information (HLPT RFI), two contractors specifically cited the Altair approach to NASA and industry working together during the early phase of a project as a good model.

Although the ACDC contract was never executed, it was held up as a model of how government and industry could work cooperatively in the formulation phases of a large project. Industry should be invited into the process early to ensure the NASA project team in aware of alternative ideas and perspectives, and to help the help the contractor community better understand the basis of the requirements. So while the acquisition model that Altair was developing may not be appropriate for wide-spread application, the primary recommendation is that the project team should explore unique and creative ways for incorporating the industry community into the project as early as possible.

V. CONCLUSIONS

During its 5 year existence as a critical element of NASA's Constellation architecture, the Altair Lunar Lander project set out to change the way the Agency designs large-scale human spacecraft. Through the initial use of small, hand-selected "Skunkworks" team, the project was able to greatly reduce the size of a typical human spacecraft project office. By utilizing risk-informed design, the Altair team was able to approach the vehicle's design from the bottoms-up with the knowledge of how every component contributes to the vehicle's overall performance, cost, safety and reliability. The use of lander Design Analysis Cycles (LDACs) proved to be an effective method to implement the risk-based design process, and showed the importance of safety and risk analysis personnel very early in the design process. Altair also made creative use of small contracts with both traditional and startup aerospace companies to allow them to participate early in the lander design process. At the time of Constellation's cancellation, the Altair project was preparing to further integrate multiple contractors into Phase B of the lander design.

Though it is unfortunate that the Altair team did not see its lander become flight hardware, it is fortunate that the lessons learned from this unique design experience were captured in a succinct set of recommendations that may benefit future human spacecraft designers.

REFERENCES

1. "Design, Development, Test and Evaluation (DDT&E) Considerations for Safe and Reliable Human-Rated Spacecraft Systems", NASA NESC PR-06-108, 2006
2. International Astronautics Congress (2008), Hansen, Lauri and Connolly, John, "The Altair Design and Minimum Functionality Approach", IAC-08-D2.9.-D1.6, Glasgow, Scotland, October 2008.
3. International Astronautics Congress (2009), Laurini, Kathleen and Connolly, John, "Altair Lunar Lander Development Status: Enabling Lunar Exploration", IAC-09.D2.3.2, Dejeon, South Korea, September 2009.
4. International Astronautics Congress (2010), Connolly, John F., Bayt, Robert L., and McMichael, James H., "Human Planetary Spacecraft Design Lessons", IAC-10.D2.3.5, Prague, Czech Republic, September 2010.
5. Constellation Program Lessons Learned; Volume I: Executive Summary, NASA SP-2011-6127-VOL-1, 2011
6. Constellation Program Lessons Learned; Volume 2: Detailed Lessons Learned, NASA SP-2011-6127-VOL-2, 2011