Manned Lunar Missions: An Exercise in Propulsion Trades and Sensitivities

Thomas K. Percy* and Michael PJ Benfield, Ph.D.†
Science Applications International Corporation, Huntsville, AL, 35806

A recent study performed for the In-Space Propulsion Technology Office at the Marshall Space Flight Center investigated the effects of using different propellant types on the different stages required to perform a manned lunar mission. The original study included investigations into propellant type, propellant storage technology options and sensitivities to specific impulse variations for a lunar orbit rendezvous mission. The initial mission characteristics were based on previous work led by Langley Research Center. Outlined in this paper are the results of that study and the work that followed. A lunar direct return architecture was added to the analysis. Since both architectures required assembly of the various propulsive stages in low Earth orbit and multiple launches to deliver those stages, investigations of launch sequence and scheduling sensitivities were also included. Results show that lunar direct return architectures require more mass to complete missions when compared to lunar orbit rendezvous missions. Within the given architectures, trends in the results tended to be very similar with the architectures indicating very little sensitivity to launch sequence and specific impulse variations and indicating more sensitivity to the propellant choice made for each stage and the time between launches. Even though this study investigates a small subset of the possible lunar architecture trade space, it does begin to outline some of the issues that must be investigated and the characteristics of the mission and the mission elements that are of most importance to a full architecture assessment.

Nomenclature

MSFC = *Marshall Space Flight Center*
LDR = *Lunar Direct Return*
LOR = *Lunar Orbit Rendezvous*
TLI = *Trans-Lunar Injection*
LOI = *Lunar Orbit Insertion*
TEI = *Trans-Earth Injection*
ZBO = *Zero Boil-Off*

I. Introduction

During the summer of 2004, a study was conducted within the In-Space Propulsion Technology Office at the Marshall Space Flight Center (MSFC) investigating the effects of propellant choice on the performance of manned lunar mission architectures. Leveraging work from a previous study performed at the Langley Research Center, a lunar orbit rendezvous mission was chosen as the baseline mission for investigation[1]. Spacecraft habitat element masses from the Langley study were used as the baseline masses for the current study and propulsive elements were sized based on payload and delta-v requirements for the baseline mission architecture. The propellant choices traded were liquid oxygen / liquid hydrogen (Lox/LH2), liquid oxygen / liquid methane (Lox/CH4), liquid oxygen / hydrazine (Lox/N2H4) and nitrogen-tetroxide / hydrazine (NTO/N2H4). In addition to propellant choice, propellant storage technologies were traded for the cryogenic propellants to investigate the possible benefits of employing zero boil-off systems on the propulsive stages.

---

* Systems Engineer, Technology Decisions Division, SAIC, 675 Discovery Dr. Suite 300, Huntsville, AL 35806, Member, AIAA.
† Systems Engineer, Technology Decisions Division, SAIC, 675 Discovery Dr. Suite, 300, Huntsville, AL 35806, Member, AIAA.

As a follow up to the MSFC study, a second architecture was added and more sensitivities were investigated. The lunar direct return architecture was assessed. In addition to propellant choice and propellant storage technologies, sensitivities to specific impulse variations were also investigated. Since both architectures required multiple launches and assembly in low Earth orbit, investigations into launch sequence and the time between launches were also performed on both architectures. The next few sections outline in detail the specific characteristics of each architecture as well as the general ground rules and assumptions applied to the study.

It is recognized that this investigation considered only two of the many different lunar mission architectures currently under consideration. While the primary goal of the study was to assess the two specific architectures, a secondary goal was to establish some of the general questions and areas of interest that should be considered when performing top level trades of lunar architectures. The trades cover the high level performance variations inherent in propulsion systems, the effects of launch scheduling and the effects of propellant storage issues. These issues together cover a broad spectrum of possible architecture drivers that may distinguish an architecture apart from the other options. Assessing the entire lunar architecture trade space with the same analysis used for this study could aid in the down selection of the optimum way to perform this mission.

II. Mission Overview

Two different mission architectures were investigated during the course of this study. The Lunar Direct Return (LDR) architecture involves carrying the crew return vehicle to the surface of the moon and returning directly to Earth from the lunar surface. The Lunar Orbit Rendezvous (LOR) architecture places the crew return vehicle into a lunar orbit and requires the crew to rendezvous with this vehicle after the lunar ascent burn and before the trans-Earth injection burn. Both architectures have their advantages and disadvantages both in performance and crew safety. For the purposes of this study, the only comparison performed between the two architectures was on a higher level performance basis.

A. General Assumptions and Ground Rules

Several overarching assumptions were made that apply to both architectures. While no specific launch vehicle was selected, a general guideline was followed to minimize the launch requirements. To this end, a four launch scenario was selected as the standard. Given the high energy requirement of the trans-Lunar injection (TLI) burn, this burn was split between two equally sized injection stages. This reduced the individual stage sizes to more manageable levels and took slight advantage of staging the TLI burn. Also taking into account the high energy requirements of this burn, the two TLI stages were limited to Lox/LH2 propellant only. The study team realized early that using any propellant with a lower specific impulse than Lox/LH2 would lead to significantly larger TLI stages and would be so unreasonable that no propellant but Lox/LH2 could be considered. These two stages comprised two of the four launches.

The third launch was assumed to be made up of the descent and ascent stages as well as the surface habitat and any other habitats required by the architecture excluding the crew entry vehicle. The crew was launched on the fourth launch in the entry vehicle and a requirement was set that the trans-Earth injection (TEI) stage must always be launched with the crew to provide propulsive abort capabilities from low Earth orbit.

A four member crew was used as the baseline to size the crew habitats and the assumption was made that all elements required for the mission would be injected towards the moon with the crew. In order to facilitate a trade of propellant storage techniques, boil-off rates were assumed for the hydrogen propellant (4% mass per month) and all other cryogens (1% mass per month). Data for cryo-cooler assembly masses was obtained from previous studies and curve-fit to provide equations for assessing the additional mass requirement inherent in including a zero boil-off system.[2] Trip times of three days to and from the moon were assumed and drove the calculations of the delta-v requirements for each burn and a seven day surface stay was assumed, driving the surface habitat size. In both architectures a horizontal lander configuration was assumed given that it is slightly more conservative from a mass estimation standpoint and that it is a likely candidate for actual lunar lander designs.

B. Lunar Direct Return Architecture

The Lunar Direct Return (LDR) architecture involves using the crew entry vehicle as part of the surface habitat and using the lunar ascent stage to provide the delta-v required for the TEI burn. The mission begins with assembly of the in-space elements in low Earth orbit. The two-burn TLI maneuver sends the stack towards the moon. After three days of transit, the second TLI stage provides the lunar orbit insertion (LOI) delta-v and the remaining stages loiter in low lunar orbit for a checkout period. Upon completion of the checkout procedures, the descent stage fires and the crew arrives on the surface of the moon. After living and working on the moon for seven days, the ascent/TEI stage performs the ascent and TEI burns and the crew spends three days transferring back to Earth for direct entry upon arrival. An overview of this architecture is provided in Figure 1.

In this architecture, two habitats are used. The crew entry vehicle provides crew living space for the ascent to Earth orbit, the transfer to and from the moon and partial living space on the lunar surface. The lunar surface habitat provides additional work and living space during the seven day lunar surface stay and is left on the surface during the ascent and TEI maneuvers. Due to the fact that there is a higher volume requirement for the surface stay than for the transfer legs of the mission, splitting these habitats relieves some lift requirements on the ascent/TEI stage.

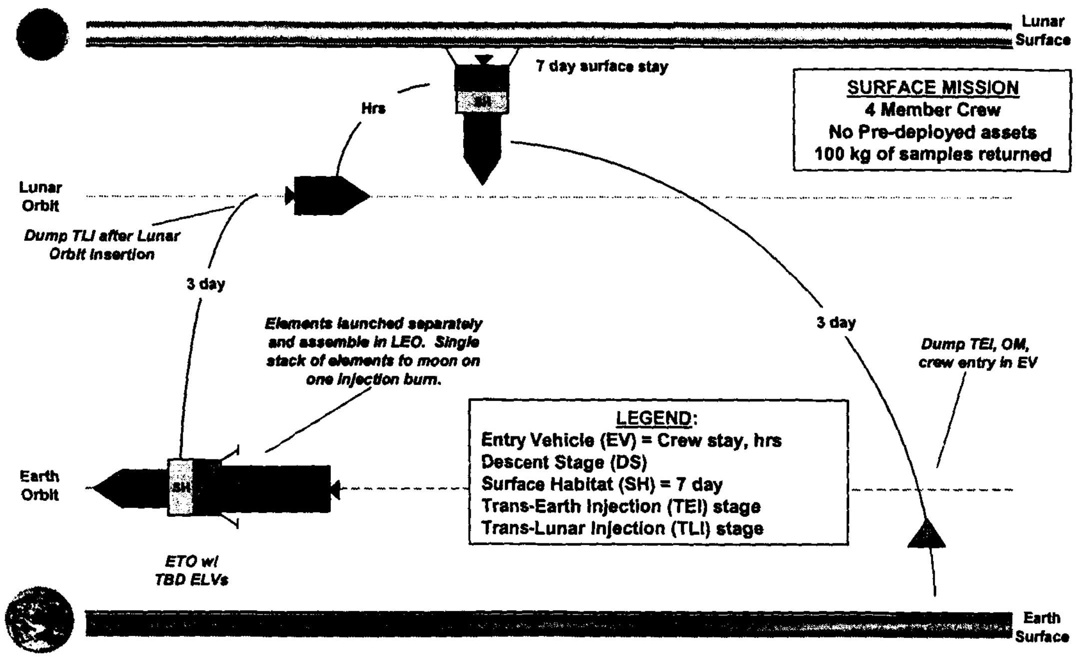

Figure 1: Lunar Direct Return Architecture

C. Lunar Orbit Rendezvous Architecture

The Lunar Orbit Rendezvous (LOR) architecture varies in some significant ways from the LDR architecture. The transit times and stay times are the same, as is the general set up of the mission elements. A single stack carrying all necessary assets is assembled in LEO after multiple launches and is injected towards the moon. The stack is captured into a low lunar orbit for checkout. It is at this point that the architectures begin to differ. The crew transfers from the entry vehicle to a transfer habitat which provides crew support for the descent and ascent maneuvers. The entry vehicle and the TEI stage remain in low lunar orbit until the crew ascends from the lunar surface after a seven day stay. The descent stage and surface habitat are left on the lunar surface and the crew rendezvous with the entry vehicle and TEI stage in orbit. The crew then returns to Earth for a direct entry arrival.

This architecture employs two TLI stages, just as in the LDR architecture. Instead of two habitats, the LOR architecture has three; an entry vehicle, a transfer habitat and a surface habitat. The surface habitat is again left on the lunar surface to relieve some of the ascent stage requirements. However, by using a transfer habitat for crew support during descent and ascent, the ascent stage payload becomes significantly smaller than in the LDR architecture. This transfer habitat also provides added habitable volume for the transfer stages of the mission, therefore allowing for the design of a smaller entry vehicle and a smaller heat shield for re-entry due to the reduced weight of the system.

Requiring a rendezvous in lunar orbit does have some disadvantages. For landing sites that are non-equatorial, the TEI stage is required to perform a second burn autonomously while the crew is on the lunar surface. This burn, referred to as the re-orientation (REO) burn, properly aligns the entry vehicle and return stage with the return trajectory and the ascent trajectory to account for the rotation of the moon under the orbit and the revolution of the moon around the Earth during the surface stay time. Considering that all four crew members descend to the lunar surface, this REO burn and rendezvous maneuver must be performed autonomously, possibly adding more risk to the maneuver. This risk was not assessed as part of this study. An overview of the LOR architecture can be seen in Figure 2.

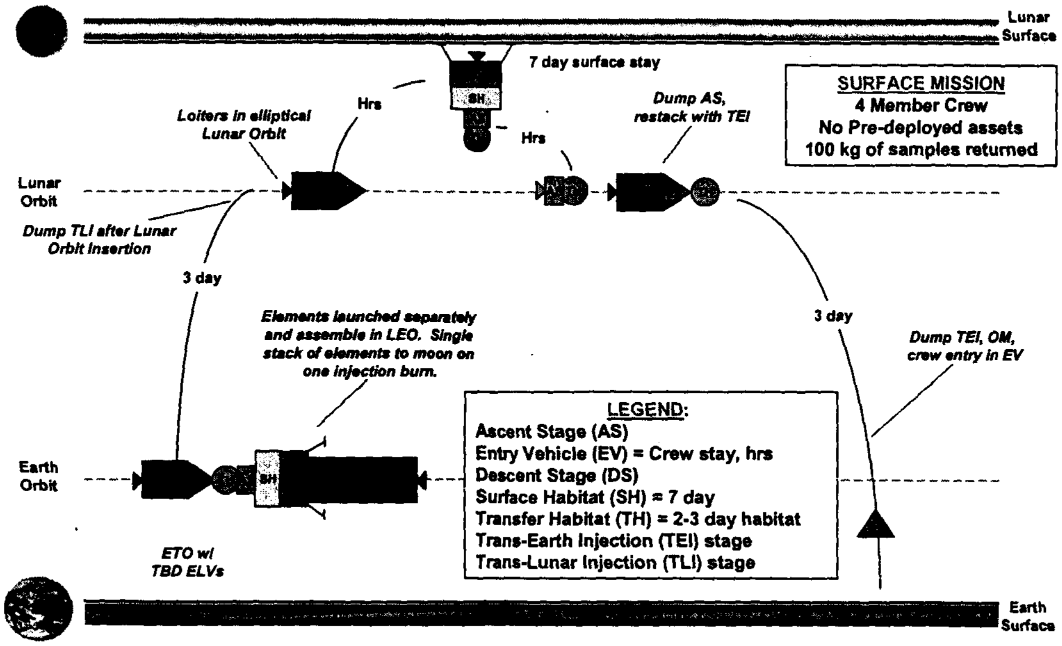

Figure 2: Lunar Orbit Rendezvous Architecture

III. Data Sources

As was previously mentioned, the Langley Research Center Lunar Architecture Study and original MSFC study contributed a significant portion of the data for this analysis[1]. Analysis performed with the Two-Dimensional Kinetics (TDK) tool was used to establish the baseline specific impulses used during the MSFC study[2]. Data used to establish the sizing relationships for the cryo-cooler assemblies was found in papers written by Dave Plachta[3,4]. Delta-v and habitat mass requirements were established during the LaRC study and were slightly modified from their original values to reasonably reflect the change from the LOR architecture to the LDR architecture.

IV. Trades & Sensitivities

Many factors must be considered to give a full evaluation of any lunar mission architecture. The intent of this study was to evaluate what were believed to be some of the major contributors to architecture performance variations. Each trade is unique and has some unique ground rules and assumptions that are outlined at the beginning of each section. The results of these trades are assessed based on the percent variation of the Initial Mass in Low Earth Orbit (IMLEO) required by a given architecture given a set of traded values. For these types of high-level trade and sensitivity analyses, the variation in IMLEO provides reasonable insight into the variations of the architectures from not only a performance standpoint, but also from the standpoint cost and launch vehicle requirements.

A. Architecture Trade: Lunar Orbit Rendezvous vs. Lunar Direct Return

The initial assessments of the two architectures involved evaluating each possible propellant and stage combination and assessing the IMLEO required for each. Given a selection of only Lox/LH2 for the TLI stages the total number of possible combinations of propellant, propellant storage technique and stages was 686 and 98, for LOR and LDR, respectively. This is due to the fact that the LDR architecture has one less stage than the LOR architecture. To assess the relative overall performance of the two architectures, an average of the IMLEO requirements for each trade was calculated and compared between the two architectures. The results of this comparison can be seen in Figure 3. Normalized to the LOR architecture, these results show that on average the LDR architectures are 22% heavier than the LOR architectures.

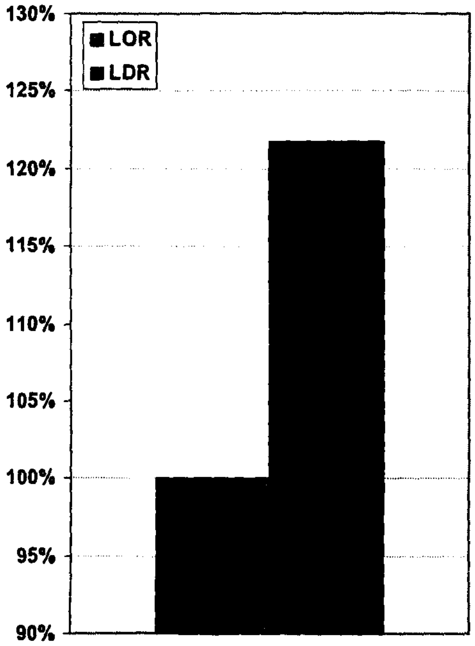

Figure 3: Comparison of Lunar Orbit Rendezvous and Lunar Direct Return Architecture Performance

B. Propellant Trade

A closer look into the effect of trading propellant on specific stages in the two architectures reveals which stages are dominating the architecture performance. Figure 4 shows the variation in IMLEO requirement for the LDR architecture when different propellants are used on different stages. These numbers were assessed by averaging all the cases run in which the stage in question used the propellant in question. For example, the first bar represents the average IMLEO requirement calculated of every architecture variation in which the descent stage used Lox/LH2 propellant. Figure 4 shows the average IMLEO requirement for each of the four propellant choices on both the descent and ascent/TEI stages.

With the baseline value for each of the stages being the Lox/LH2 trade results, the results show that using a lower performing propellant on either of the stages results in an increase in the IMLEO requirement for that architecture. Larger gaps between the baseline and the worst performing propellant option indicate stages where stage performance is more of a driver in the architecture performance. This means that factors such as total payload mass, propellant choice and specific impulse affecting individual stages with larger gaps will also affect the overall architecture performance more significantly. For the LDR architecture, there is a 10% increase in IMLEO requirement for the descent stage propellant choice variation and an 11% increase for the ascent/TEI stage. This indicates that the choice of ascent/TEI propellant and the choice of descent propellant should carry equal weight. These increases in IMLEO are also considered statistically significant leading to the conclusion that choosing propellant for either of these stages can play a major role in determining the overall performance of the LDR architecture.

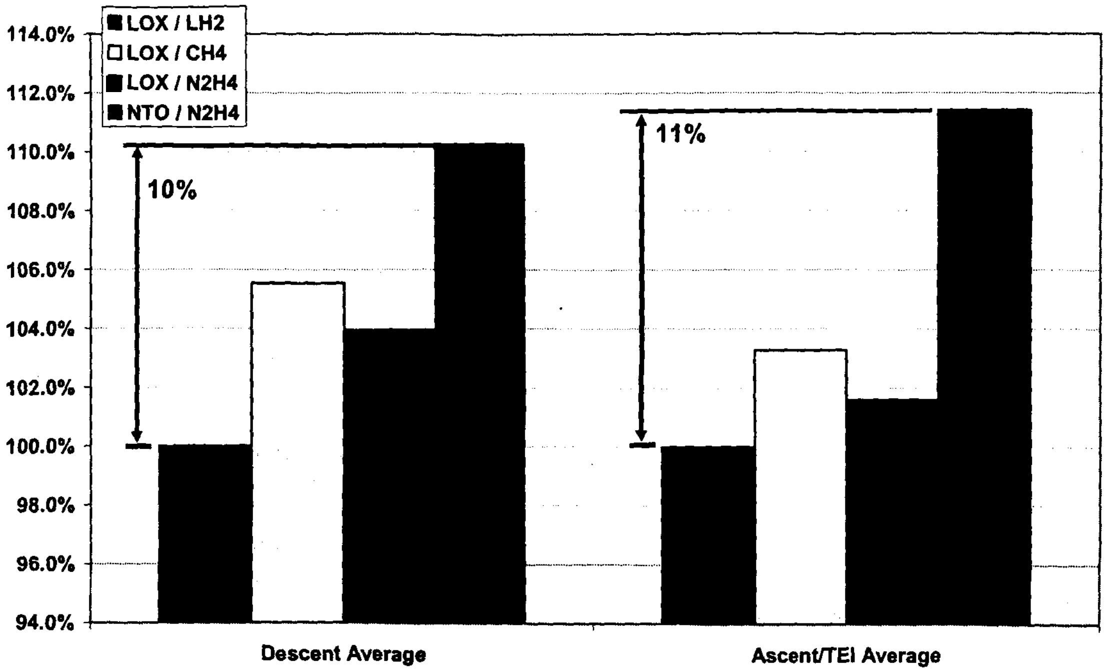

Figure 4: Lunar Direct Return Architecture Sensitivity to Propellant Choice

The same propellant choice trade was assessed for the LOR architecture. These results can be viewed in Figure 5. In the case of the LOR architecture, the descent stage is by far the dominant stage in the architecture with a 13% gap in IMLEO requirements between the best and worst performing propellant choice. The TEI stage is the least significant stage and the percent differenced indicated for the ascent and the TEI stages is statistically insignificant given the data set created by the trade analysis. This would indicate that choices made regarding payload requirements and propellants for the descent stages will have the greatest effect on the overall LOR architecture performance while choices made with respect to the TEI stage are insignificant.

A correlation can be drawn between a stage's "energy requirement" and its relative contribution to the overall architecture performance. The energy requirement of a stage is a qualitative assessment of what is being physically required of a stage. Stages with high payload masses and high delta-v requirements have high energy requirements while stages with low payload masses and low delta-v requirements have low energy requirements. Using this distinction it stands to reason that the descent stage would dominate LOR architecture performance. Of the three stages considered, it is required to carry significantly more payload and has delta-v requirements higher than any other stage. On the opposite end of the spectrum, the TEI stage has a significant delta-v requirement but very small payload requirements.

The results of the propellant stage trade also provide insight into the results of the overall comparison of LDR and LOR architectures. By splitting the habitats and minimizing the size of the ascent stage, the payload requirement on the descent stage of the LOR architecture is also reduced. The entry vehicle and TEI stage also never travel to the surface, further reducing the payload requirement on the descent stage. Since the descent stage dominates the LOR architecture performance, that particular architecture will benefit greatly from this reduced requirement. Looking at the LDR architecture, there is a much different picture. Payload masses for the ascent stage are significantly higher than in the LOR architecture and the delta-v requirement for the ascent stage is also higher. This not only makes the ascent stage a significant contributor to architecture performance but increases the payload requirement on the descent stage. Since both stages are significant contributors to the LDR architecture performance, it is easy to see why the LOR architecture performs 22% better on average than the LDR architectures.

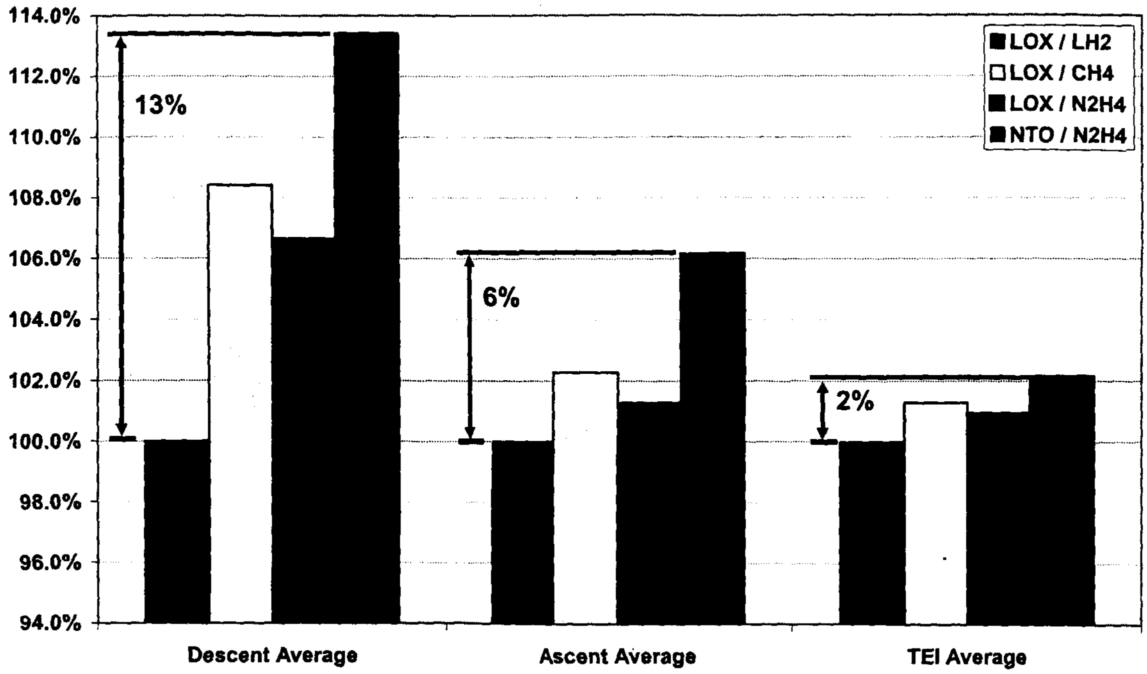

Figure 5: Lunar Orbit Rendezvous Architecture Sensitivity to Propellant Choice

C. Propellant Specific Impulse Trade

For assessing the architecture sensitivities to variations in propellant specific impulse, the same overall averages discussed in section A were used. Variations of plus and minus five seconds of specific impulse were used on each stage in succession generating variations in IMLEO requirements that reflect the increase and decrease in propellant performance. The results of this trade can be seen in Figure 6. For both architectures, the results are the same. As expected, as propellant performance increases, indicated by a five second increase in specific impulse, the IMLEO requirements are reduced. The inverse is also true. The only notable result presented here is the difference in the variance between the Lox/LH2 variations and the other propellant choices. This result is purely an artifact of the ground rules of the study. While descent and ascent stages were considered high energy, and therefore dominant stages, neither is nearly as dominant as the TLI stages. With significantly higher requirements in both payload and delta-v, changes made to the TLI stages have a truly significant effect on the performance of the architectures. It was for this reason that these stages were limited to the highest performing propellant in the study, Lox/LH2.

However, since the specific impulse variations only affected the TLI stages during the Lox/LH2 specific impulse variations, the Lox/LH2 appears to be more significantly affected by this variation. Excluding the TLI stages from the specific impulse analysis brings the IMLEO variation for Lox/LH2 to the same level as the other three propellant choices. The trends among propellant choice are the same between the two architectures, however the sensitivity of the architectures varies slightly when compared to each other. The LOR architecture does show slightly greater variation from the baseline for each propellant choice, but none of these variations are statistically significant. In order to find significant changes in IMLEO requirements, specific impulse must be greatly varied. This is better represented by the data in the previous section which trades propellant types, an extreme form of specific impulse variation.

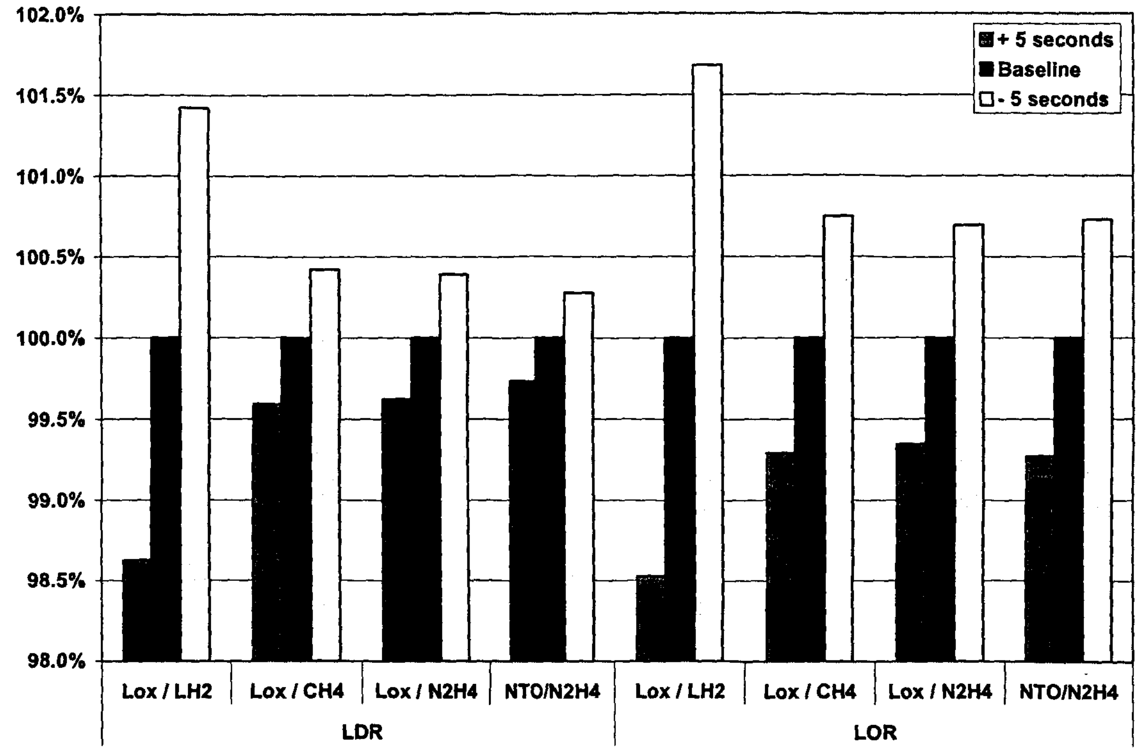

Figure 6: Architecture Sensitivity to Specific Impulse Variation

D. Propellant Storage Technology Trade

The analysis of the propellant storage technology impact on the architecture performance compares the use of a zero boil-off (ZBO) system versus the use of passive systems that allow cryogenic propellant to boil off. With a ZBO system, the mass of a cryo-cooler assembly and the power systems that drive that assembly must be added to the dry mass of the stage being assessed. However, no additional propellant must be added to the stage to account for propellant that is boiled and released from the stage during the time the stage is exposed to the space environment. With a passive storage system, this additional propellant must be added to assure that a sufficient amount of propellant will still be in the stage in order to perform the burns required. The real question in assessing the difference between the two methods is when does the extra propellant and tank mass inherent in the passive system become heavier than the ZBO system, thus making the ZBO system a more efficient method for propellant storage?

For this trade, the ZBO systems were approximated by establishing curves that fit previous design data and are dependent on the propellant mass of the stage[3,4]. The averages of the trade runs using ZBO systems were compared to the averages of the trade runs using passive systems and the percent variation of IMLEO was used as the comparative characteristic as in previous trades. The exposure time of each stage was assumed to be one month and the boil-off rates were assumed to be 4% mass per month for hydrogen and 1% mass per month for oxygen and methane. Hydrazine and nitrogen-tetroxide are both considered storable propellants and were not assessed in this trade.

The results of the trades for the two candidate architectures can be seen in Figure 7. The trends are the same for both architectures. These results show that for one month exposures in space, the ZBO system increases the IMLEO requirement when applied to all propellant combinations. While the 0.8% to 0.1% differentials are not considered statistically significant, this does reveal that for these stages and this exposure time architectures are more efficient when allowing boil-off to occur rather than adding the extra mass of a ZBO system.

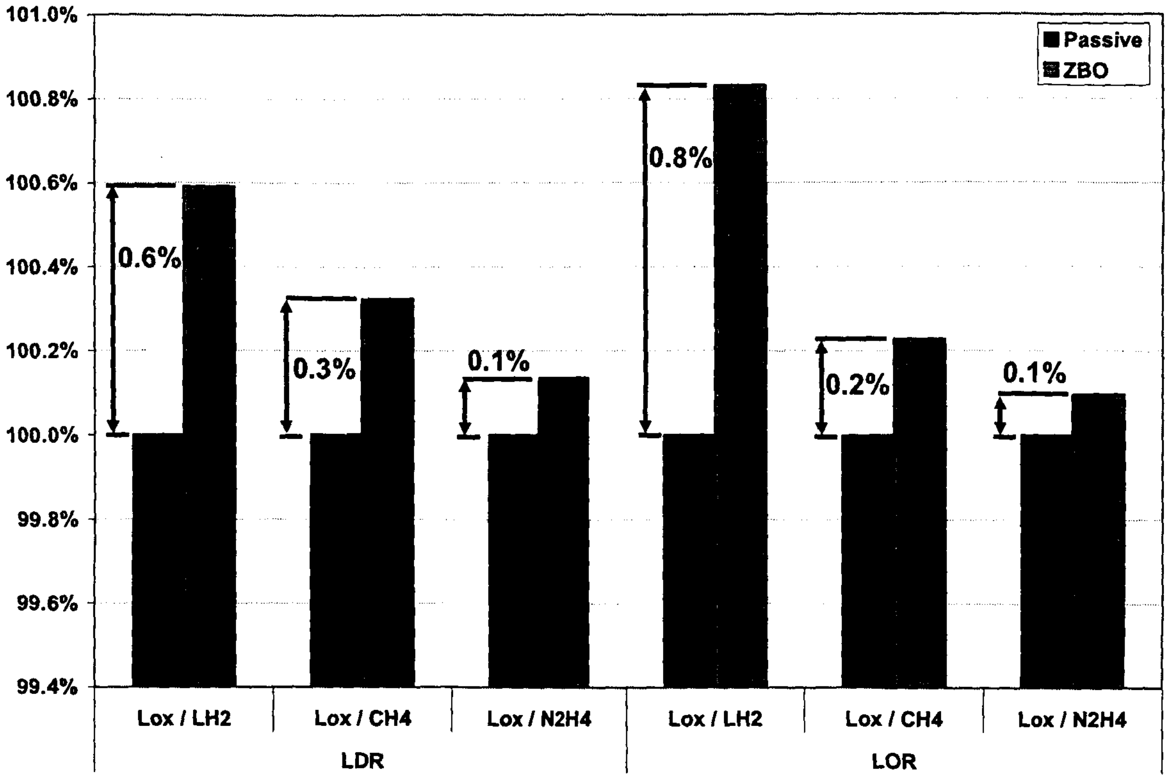

Figure 7: Architecture Trade: Propellant Storage Technology

Early results of the study indicated that the high energy stages, especially the TLI stage due to their exclusive use of hydrogen, would begin to benefit from ZBO systems as the exposure time of the stages increased. This result is not unexpected for two reasons. The first is the fact that boiled propellant mass increases with exposure time and therefore requires more initial propellant mass to ensure enough propellant remains for the burn. The second reason for these results is that ZBO systems are dependent on propellant mass and not dependent on exposure time. If one establishes the propellant mass required for a burn, the zero boil-off system will be the same mass independent of how long the stage is exposed to the space environment. With passive systems, the increase in propellant mass with expected exposure time will eventually drive the mass of the stage over the mass of a stage with a ZBO system.

E. Launch Variation Sensitivities

For the original propellant storage technology trade, a standard one month exposure time was assumed for all stages. While this begins to reveal the trends inherent in architectures employing these two propellant storage technologies, the exposure time does not reflect the real exposure times that would be encountered in a true mission scenario. The mission timeline for the LOR and LDR architectures is very similar, but what will establish the exposure time of each stage is the launch operations characteristics. The assumption was made early in the study that all architectures would consist of four launches and account for some assembly time in low Earth orbit. To reflect a more realistic set of exposure times, variations in launch operations were investigated for their effects on the architecture performance.

The two main factors that will vary with respect to the launch operations are the order in which the stages are launched into low Earth orbit, and the time gap that is present between launches. The stage that is launched first will, in general, be exposed to the space environment longer than the stage that is launched last. The magnitude of this exposure time will be dependent on how much time passes between launches. Several assumptions were made to simplify this portion of the analysis. A general assumption regarding mission planning is that the crew launch must always include the TEI stage in order to provide low Earth orbit abort capabilities in the event that rendezvous and docking fails and the crew must be returned to Earth before the mission begins. A further assumption made for this analysis is that the crew will always be the last element of the mission launched into space. Given that each architecture has four launches and one is the crew launch, there are now three launches available for launch sequence trades. This leads to six possible launch sequence combinations to be traded.

With respect to the propellant choices, a specific propellant was set as the baseline for each stage and no propellant trades were performed during the launch variation analyses. As per the original ground rules of the study, the TLI stages were assumed to use Lox/LH2 propellant only. The results reviewed previously indicated that the descent stage is a major contributor to the performance of both architectures. For this reason, the descent stage was assumed to use Lox/LH2 propellant only. The ascent and TEI stages were assumed to be NTO/N2H4 propellant to reflect the choice most likely to be made. This choice reflects considerations for crew safety given the reliability and simplicity of pressure-fed hypergolic propellant systems.

First investigated was the time between launches. For this trade, four combinations of passive and ZBO propellant storage technologies on the TLI and descent stages were investigated. The baseline time between launches was set at two weeks with trades made at four, six and eight weeks. The results can be seen in Figures 8 and 9. Several conclusions can be drawn from these results. The first is that as the gap between launches increases, so does the required IMLEO for each architecture. The LDR architecture variations are slightly greater than the LOR architectures due to the fact that LDR architectures are, on average, more massive than LDR architectures. What is most interesting about these results is the fact that even with launch gaps as large as two months (indicating a total time to complete one manned lunar mission of 188 days) the largest increase in IMLEO is statistically insignificant at 1.6%.

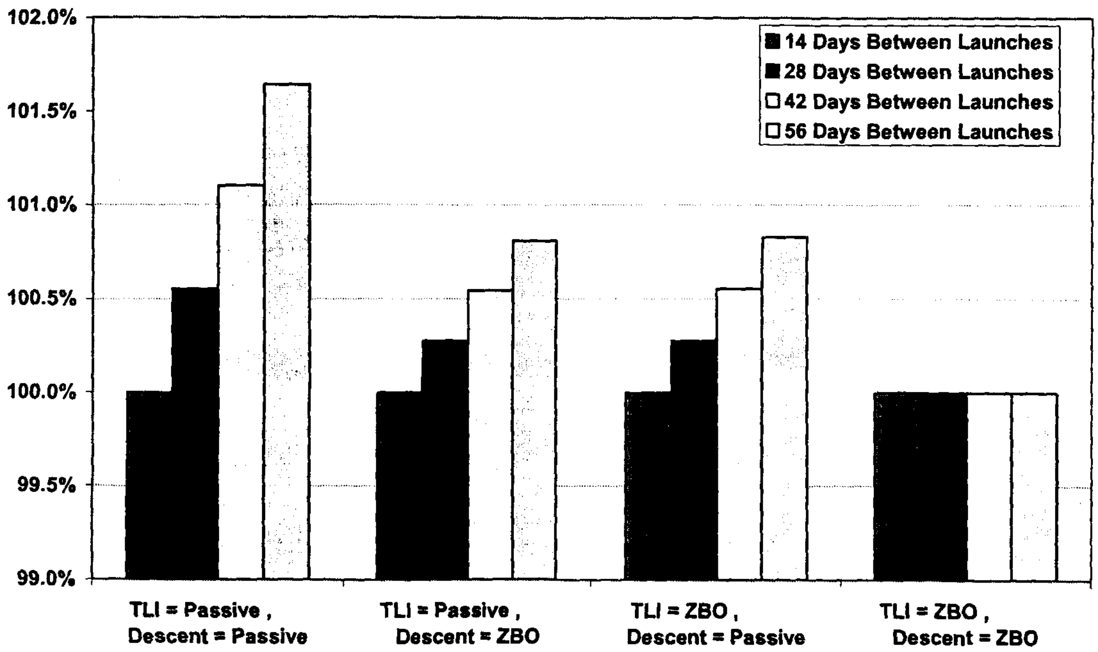

Figure 8: Lunar Direct Return Architecture Sensitivity to Time Between Launches

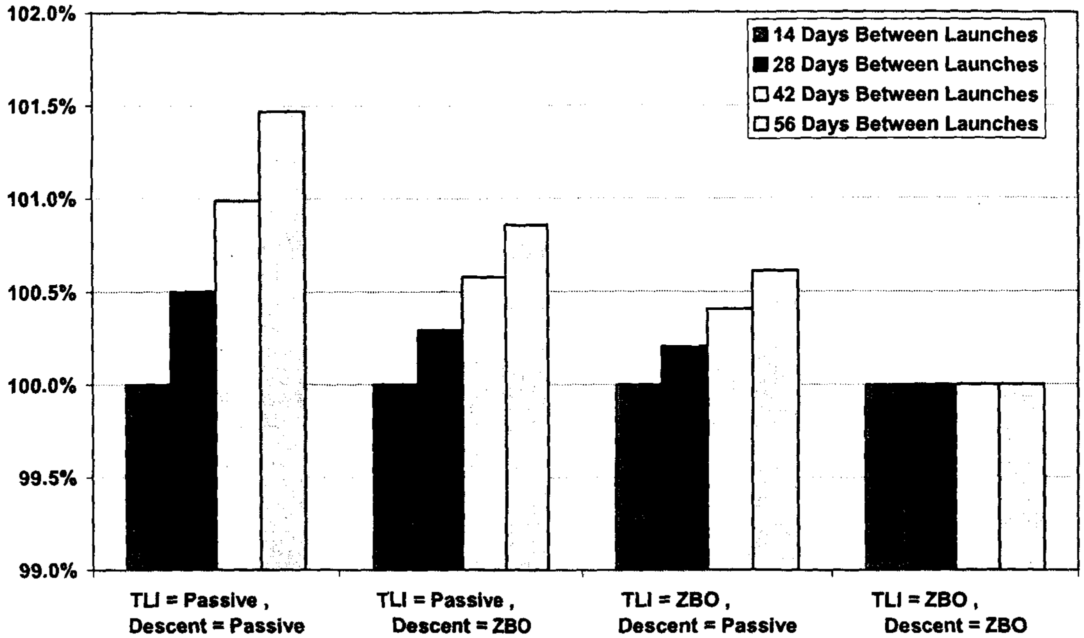

Figure 9: Lunar Orbit Rendezvous Architecture Sensitivity to Time Between Launches

Trading the propellant storage technology employed by the stages offers further insight into the possible benefit of ZBO systems. For the LDR architectures, the individual TLI stages and the descent stage have very similar energy requirements. Therefore, applying a ZBO system to either stage produces approximately the same reduction in IMLEO variation when trading time between launches. When observing this same trade for the LOR architectures, the results indicate that the ZBO system benefits the architecture performance more when applied to the TLI stages than to the descent stage. Again, this follows the energy requirements of the stages given that the descent stage for the LOR architectures has a lower energy requirement than the TLI stages. For all architectures, by applying ZBO systems to both stages and using a storable propellant for all other stages, the IMLEO variation is non-existent.

After investigating the sensitivity to time between launches, the focus turned to the sequence of those launches. The six sequential options were assessed with the same ground rules as indicated for the trade of time between launches. These sequences were assessed over the four launch gap values previously investigate. The results of this trade can be seen in Figures 10 and 11. Table 1 shows the breakdown of stages for each sequence. The results of the launch sequence trade are similar for each architecture. For each of the two architectures, these results show that the launch sequence effect is the same regardless of the time between launches. The overall trend also shows that the magnitude of IMLEO variations with respect to launch sequence increase as the time between launches increases. This result is consistent with what was shown in the previous trade and the results further indicate that choosing a poor launch sequence can magnify the effect of longer gaps between launches. However, even with a poor sequence chosen with the longest time between launches, the IMLEO variation is still not statistically significant.

Table 1: Launch Sequence Breakdown

Lunar Direct Return Architecuture				
	Launch 1	Launch 2	Launch 3	Launch 4
Sequence 1	TLI 1	TLI 2	Descent	Ascent/TEI
Sequence 2	TLI 1	Descent	TLI 2	Ascent/TEI
Sequence 3	TLI 2	TLI 1	Descent	Ascent/TEI
Sequence 4	Descent	TLI 1	TLI 2	Ascent/TEI
Sequence 5	TLI 2	Descent	TLI 1	Ascent/TEI
Sequence 6	Descent	TLI 2	TLI 1	Ascent/TEI

Lunar Orbit Rendezvous Architecture				
	Launch 1	Launch 2	Launch 3	Launch 4
Sequence 1	TLI 1	TLI 2	Desc/Asc	TEI
Sequence 2	TLI 1	Desc/Asc	TLI 2	TEI
Sequence 3	TLI 2	TLI 1	Desc/Asc	TEI
Sequence 4	Desc/Asc	TLI 1	TLI 2	TEI
Sequence 5	TLI 2	Desc/Asc	TLI 1	TEI
Sequence 6	Desc/Asc	TLI 2	TLI 1	TEI

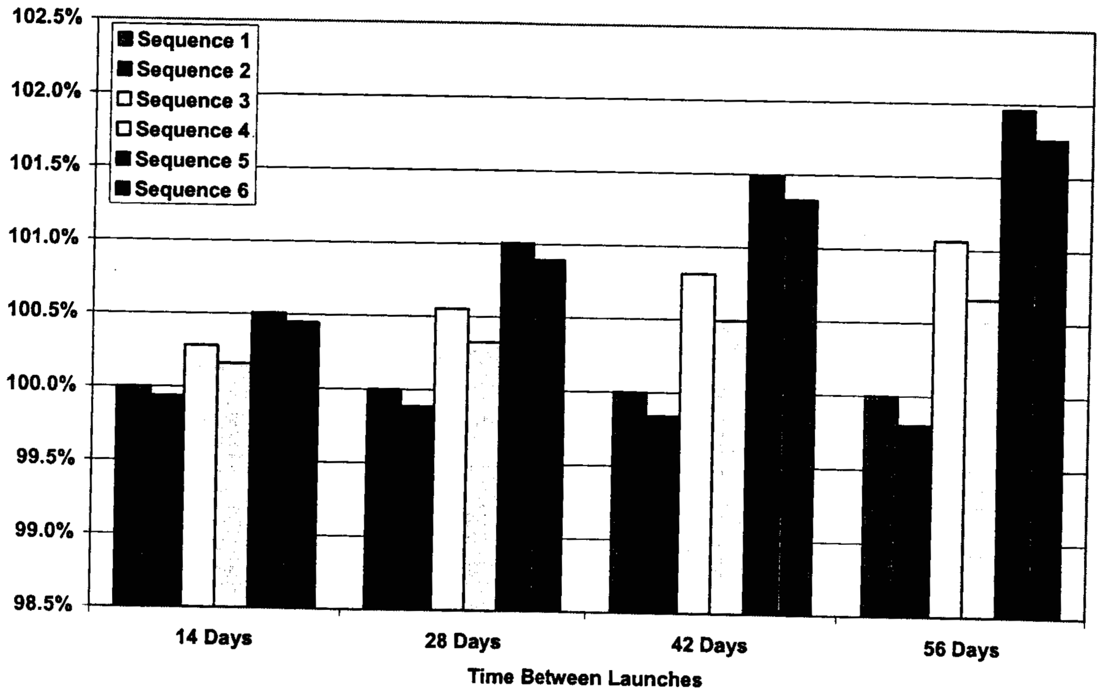

Figure 10: Lunar Direct Return Architecture Sensitivity to Launch Sequence

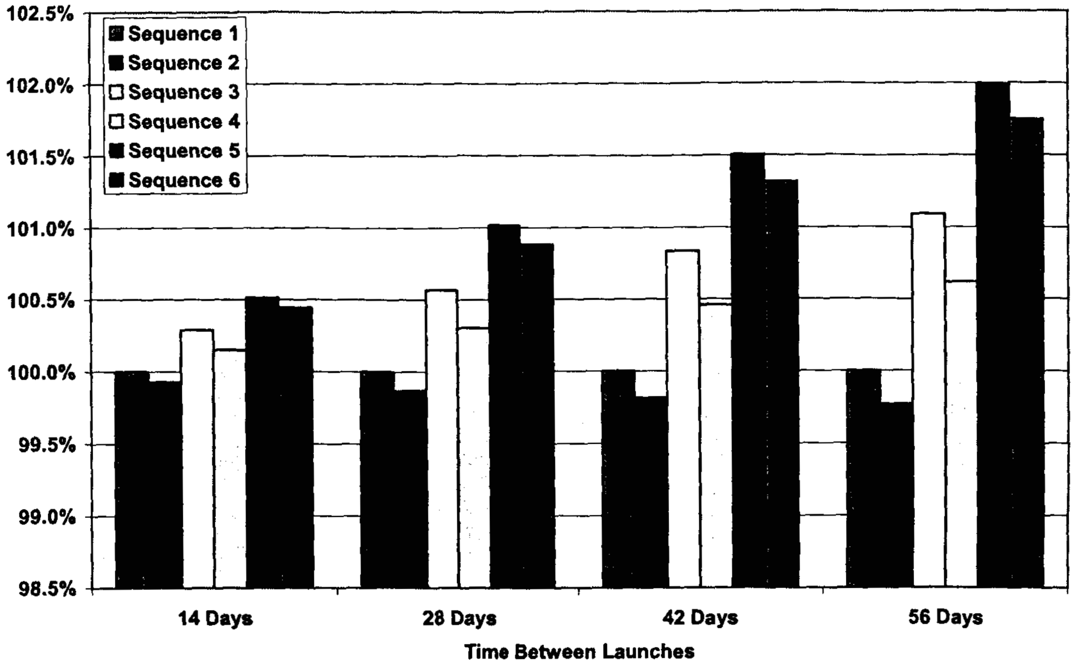

Figure 11: Lunar Orbit Rendezvous Architecture Sensitivity to Launch Sequence

V. Conclusions

When performing high-level analyses of lunar architectures, it is important to know which questions to ask. Through this investigation, the true drivers of architecture performance were isolated and other characteristics believed to be possible drivers were not. For missions to the moon involving chemical propulsion, the majority of the driving factors will be propulsion related. Those investigated in this study, propellant choice, specific impulse variations, propellant storage technologies and launch considerations, are some of the major driving factors. Several conclusions can be drawn from this study.

The first is that determining the qualitative energy requirements of the individual stages required by an architecture will determine which stages will have the greatest impact on the architecture performance. When trading propellant types with stages, results indicated that both architectures were greatly affected by changes in the descent stage propellant. Assessing the delta-v and payload requirements quickly shows that the descent stage has the highest energy requirement of any stage other than the TLI stages for these architectures. This high energy stage is also a significant architecture performance driver, as indicated by the trades performed. The energy of a stage will indicate that stage's level of significance in the performance of the architecture.

Another conclusion to draw is that small variations in specific impulse for the propellant choices traded have little effect on the architecture performance. Also indicated from the trades is the fact that, even with a two month gap between element launches, the IMLEO variation with respect to launch gap and launch sequence is low. This shows that for the shorter missions investigated the use of a ZBO system for propellant storage is not required. For longer stays at the moon or for chemical Mars mission architectures this will most likely not be the case, but early lunar missions reap little or no benefit from employing ZBO propellant storage systems.

This study investigated a small portion of the overall lunar mission architecture trade space. Other architectures are being considered for lunar missions, all with their own set of advantages and disadvantages. This study was also limited to architectures with seven day surface stays, and was not constrained by the selection of a specific launch vehicle. Payloads, such as crew habitats and equipment, were restricted to a specific mass and not traded. All of these characteristics could potentially affect the performance of the architecture, especially payload masses, which can affect the energy requirement of a stage. For a truly complete assessment of the lunar mission trade space, all of these issues should be considered. What this study has accomplished is identifying the key trades that should be run to get a good understanding of each architecture at a high level as well as lending some insight into the performance of the two specific architectures investigated. Applying these types of trade analyses to all lunar architectures will begin to reveal which architectures have the greatest potential for helping NASA realize its goal of returning humans to the surface of the moon.

Acknowledgements

The work described in this paper was funded in whole or in part by the NASA Space Architect, NASA Headquarters, Washington, D.C.

References

[1]ESMD-RQ-0006,. "Lunar Architecture Broad Trade Study Final Report", Version Baseline, 1 September, 2004.

[2]Stevens, K. et. al., "Lunar Engine Comparison Study" NASA TM. In Press.

[3]ter Brake, H. J. M. and Wiegerinck, G. G. M., Low-power Cryocooler Survey, Cryogenics, 42, 705-718, 2002.

[4]Glaister, D. S, Donabedian, M., Curran, D. G. T., and Davis, T., An Assessment of the State of Cryocooler Technology for Space Applications, The Aerospace Corporation, Report No. TOR-98 (1057)-3, 1998.