AAS 09-075

Lunar Landing Trajectory Design for Onboard Hazard Detection & Avoidance

Steve Paschall, Tye Brady, Tom Fill, Ron Sostaric

The Charles Stark Draper Laboratory / NASA Johnson Space Center

32nd ANNUAL AAS GUIDANCE AND CONTROL CONFERENCE

January 31 - February 4, 2009
Breckenridge, Colorado

Sponsored by
Rocky Mountain Section

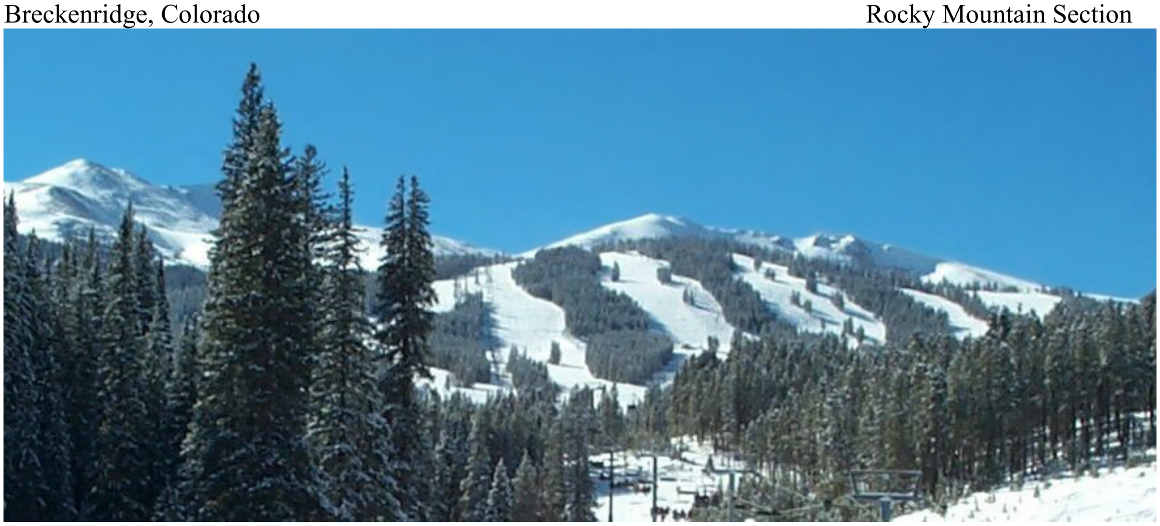

AAS Publications Office, P.O. Box 28130 - San Diego, California 92198

LUNAR LANDING TRAJECTORY DESIGN FOR ONBOARD HAZARD DETECTION & AVOIDANCE

Steve Paschall,[*] Tye Brady[†], Tom Fill,[‡] and Ron Sostaric[§]

The Autonomous Landing and Hazard Avoidance Technology (ALHAT) Project is developing the software and hardware technology needed to support a safe and precise landing for the next generation of lunar missions. ALHAT provides this capability through terrain-relative navigation measurements to enhance global-scale precision, an onboard hazard detection system to select safe landing locations, and an Autonomous Guidance, Navigation, and Control (AGNC) capability to process these measurements and safely direct the vehicle to a landing location. This paper focuses on the key trajectory design issues relevant to providing an onboard Hazard Detection and Avoidance (HDA) capability for the lander. Hazard detection can be accomplished by the crew visually scanning the terrain through a window, a sensor system imaging the terrain, or some combination of both. For ALHAT, this hazard detection activity is provided by a sensor system, which either augments the crew's perception or entirely replaces the crew in the case of a robotic landing. Detecting hazards influences the trajectory design by requiring the proper perspective, range to the landing site, and sufficient time to view the terrain. Following this, the trajectory design must provide additional time to process this information and make a decision about where to safely land. During the final part of the HDA process, the trajectory design must provide sufficient margin to enable a hazard avoidance maneuver. In order to demonstrate the effects of these constraints on the landing trajectory, a tradespace of trajectory designs was created for the initial ALHAT Design Analysis Cycle (ALDAC-1) and each case evaluated with these HDA constraints active. The ALHAT analysis process, described in this paper, narrows down this tradespace and subsequently better defines the trajectory design needed to support onboard HDA. Future ALDACs will enhance this trajectory design by balancing these issues and others in an overall system design process.

[*] Senior Member Technical Staff, GN&C Systems Division, The Charles Stark Draper Laboratory, Inc., 555 Technology Square, Cambridge, MA 02139. spaschall@draper.com
[†] Principal Member Technical Staff, Systems Engineering Division, The Charles Stark Draper Laboratory, Inc., 555 Technology Square, Cambridge, MA 02139. tye@draper.com
[‡] Principal Member Technical Staff, GN&C Systems Division, The Charles Stark Draper Laboratory, Inc., 555 Technology Square, Cambridge, MA 02139. tfill@draper.com
[§] Aerospace Engineer (AST), Flight Mechanics and Trajectory Design, NASA Johnson Space Center, 2101 Nasa Parkway, Houston, TX 77058. ronald.r.sostaric@nasa.gov

INTRODUCTION

The return of humans to the Moon will require increased capability beyond that of the previous Apollo missions. Longer stay times and greater flexibility with regard to landing locations are among the many improvements planned. A descent and landing system that can land the vehicle more accurately than Apollo with a greater ability to detect and avoid hazards is important to the development of a lunar outpost, and also for increasing the number of potentially accessible lunar sortie locations. A flexible descent and landing system will allow landings in more challenging terrain and provide more opportunity with regard to mission timing and lighting considerations, while maintaining safety as the top priority. The ALHAT (Autonomous Landing and Hazard Avoidance Technology) project is developing a self-contained lunar landing system, which will address this need. The ALHAT system will use Terrain Relative Navigation (TRN) techniques to enable high-precision navigation and will detect hazards at the landing site by utilizing an onboard Hazard Detection and Avoidance (HDA) sensor system. These TRN and HDA capabilities operating in conjunction with the Autonomous GNC system will enable the lander to land safely and precisely, independent from ground control, without lunar navigation infrastructure, and without *a priori* knowledge of a safe landing location. Further, the system will not be restricted by local lighting conditions, as was the case for Apollo[1]. The system is designed to allow for crew reach-in to the AGNC system so that the ALHAT system can support both robotic and crewed missions.

The overall trajectory profile is a fundamental design parameter to the lander system design. The profile must be designed such that it provides a low ΔV solution to minimize consumed propellant mass during landing, while also accommodating the various other needs of the lander. Some of these other needs are: terrain clearance during the deorbit and powered braking phases, dispersion correction margin for maneuver execution and navigation errors, and increased time and improved perspective of the landing site during the final approach. This paper focuses on the tradeoffs and constraints surrounding the design of the final approach and landing trajectory to support an onboard HDA capability.

WHY ONBOARD HDA?

Safety is considered the most critical issue for a crewed spacecraft system. For the purposes of this paper, safety is discussed only in the context of achieving a safe landing. Safe landing is defined in terms of the lander's orientation and velocity relative to the surface and the terrain character at that location. Terrain is judged hazard-free if it is sufficiently level and free from any rocks or holes that could damage or tip the lander[*]. Several different mission design strategies can be used to address landing safety.

One such strategy would involve a direct, physically certified safe landing location. The terrain at this selected location would be judged prior to the mission to be hazard-free and safe for landing. In this scenario the lander's primary task would be to fly to this spot precisely and execute a safe touchdown. To support this landing approach, high-quality terrain information would need to be collected via either high-resolution orbital reconnaissance or via a robotic pre-

[*] The definition of a terrain hazard is tied to the lander physical design. The lander is said to have a specific "hazard tolerance" based on its ability to cope with various boulders, holes, slopes, etc.

cursor mission that collects this data from the lunar surface. Another possibility would be to create a safe landing point (e.g. construct a landing pad) on an earlier mission.

A second strategy would involve an indirect, statistically certified safe landing area. This mission design concept implies that the landing area of interest has been observed with enough detail to statistically characterize the terrain and extrapolate to some level of probability that safe landing locations are present within that area. With this strategy, no further information on the landing area is available in advance, so the mission design can follow one of two design paths. The first is that the lander reaches this landing area and makes a blind landing. This means that the actual landing location will be some random point within the landing area with only statistical insight into the terrain hazards. For this approach to provide a high safe-landing probability, the local terrain character would need to be extremely benign or the lander would need a very high hazard tolerance. If more risk were acceptable for the landing (i.e. lower safe-landing probability for a robotic lander), this blind landing approach would be feasible without imposing those terrain or vehicle design restrictions. The other design path for this statistical approach would require that the lander identify a safe landing location, fly to this spot precisely, and then safely touch down. Such a strategy would require the lander to locate and maneuver to a safe landing location by employing an onboard HDA capability. An advantage of this approach is that landings are not as reliant on precursor missions or local infrastructure. Landings can occur at many locations on the Moon by using statistical hazard frequency information gathered from lower resolution *a priori* terrain data. No directly identified or prepared safe site is necessary. Another advantage is that a high safe-landing probability may be achievable even in very hazardous terrain where a blind, random landing might have little chance of success.

An onboard HDA capability is desirable because it enables a safe, successful landing in a much wider range of locations, without local infrastructure or significant precursor support missions. In the extreme, an HDA capability also provides information about the lack of any safe landing site early enough in the final approach to abort the landing and save the vehicle. The other mission design concepts mentioned here can also work for lunar landing, but lack the flexibility that an onboard HDA capability provides. Over time, infrastructure will likely be put in place to either identify or prepare safe landing sites. Once this occurs, an onboard HDA capability can continue to add value by providing additional robustness to the lander in the event of off-nominal landings.

LANDING TRAJECTORY OVERVIEW

Figure 1 shows the major phases of a typical lunar descent trajectory from a parking orbit. For ALHAT this initial orbit is assumed circular with a 100 km altitude. The deorbit burn begins the sequence of maneuvers necessary to land on the lunar surface. The deorbit burn targets a 100 x 15 km transfer orbit. The periapse of 15 km was chosen to minimize propellant usage without exceeding safety margin needed for terrain clearance, as well as a passive abort in case of a failed Powered Descent Initiation (PDI). The deorbit burn is followed by a coast to PDI of about one hour duration.

The powered descent phase, depicted in Figure 2, begins at PDI and continues until touchdown on the lunar surface. The engine remains on throughout this phase. Powered descent consists of 4 sub-phases: Braking, Pitch-up, Approach, and Terminal Descent. The objective of the braking maneuver is to remove the majority of the orbital velocity as efficiently as possible while targeting a certain altitude and range from the landing site. The timing of PDI (i.e. the start of the braking maneuver) is chosen so that the trajectory meets these target conditions efficiently given the vehicle's thrust-to-weight ratio. During braking, the engine throttle remains at a high and relatively constant setting, and the vehicle is in a relatively horizontal orientation. During this maneuver, any trajectory dispersions due to burn execution errors or improved navigation state knowledge can be corrected. Errors in the along-track direction are mitigated by modulating the engine throttle, while cross-track errors are mitigated by directing the thrust out-of-plane to the trajectory.

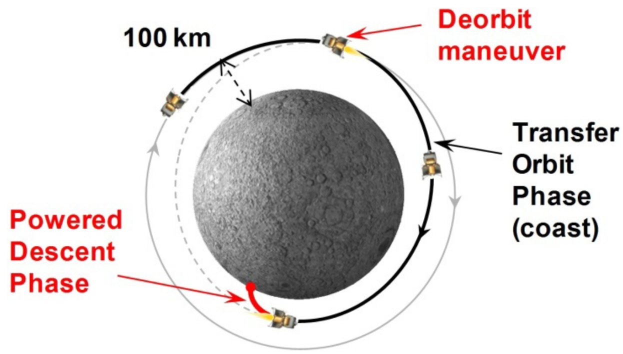

Figure 1 – Major Phases for Lunar Descent Trajectory

The approach phase is purposely designed to have a more vertical attitude and lower acceleration level than the braking phase. The vertical attitude provides better visibility of the landing area, while the lower acceleration level provides for both slower speeds and more observation time while approaching the target. The intermediate pitch-up maneuver phase provides a smooth transition in acceleration level and vehicle attitude from the high thrust, near horizontal braking conditions to these desired approach conditions. By this point in the trajectory, the closed-loop Guidance & Control systems have corrected trajectory dispersions to the level of the Navigation system accuracy.

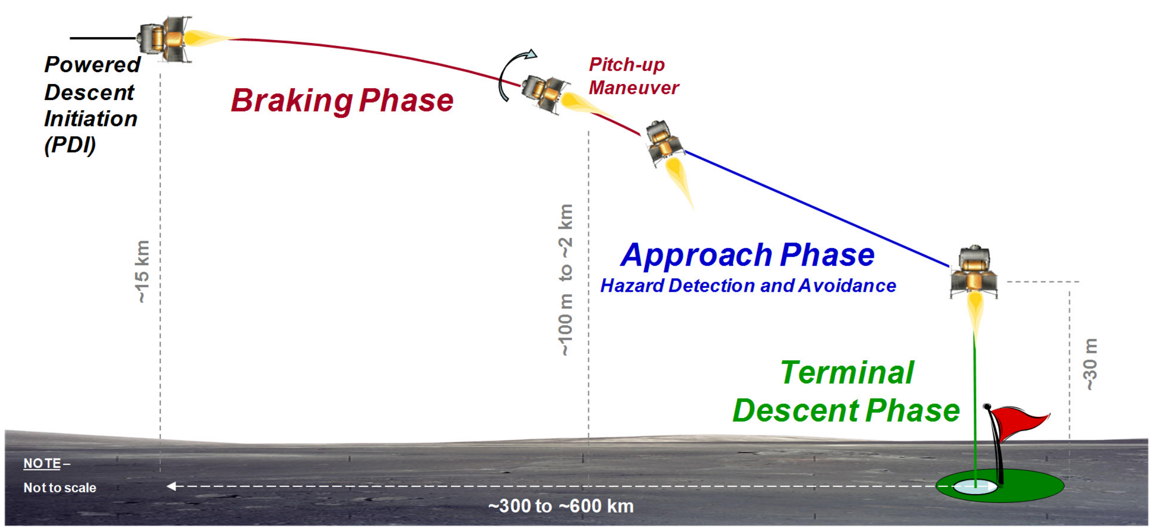

Figure 2 – Powered Descent Sub-Phases and Events

The objective of the final terminal descent maneuver is to descend slowly to the landing site in a near vertical orientation, staying directly over the landing target and nulling out any remaining horizontal velocity. At this point the landing target will likely no longer be visible to the crew and/or sensors onboard the vehicle. This is because the vehicle is descending directly from above the target, and lunar dust scattered from the engine exhaust will likely obscure the terrain below. Based on Apollo experience the terminal descent maneuver for ALHAT begins at 30m altitude.

The ALHAT profile parallels the Apollo descent profile in many respects[2]. The Apollo descent profile included a braking phase, an approach phase, and terminal descent. The beginning of the braking phase (also referred to as PDI) typically began from a 15.24 km perilune altitude elliptical parking orbit. From there, the crew began PDI by entering Program 63 (P63) into the guidance computer. The lunar module would continue in P63 until the guidance target for the beginning of the approach phase (referred to "High Gate"), at 2.2 km altitude and 7.8 km slant range, was reached and then automatically switch to the approach phase guidance, Program 64 (P64). The sudden switch to P64 at High Gate would produce a pitch-up of the vehicle to aid viewing of the landing site. P64 guided the spacecraft to "low gate", which occurred at 30 m altitude and 32 m slant range. If the crew did not intervene with the guidance programs, the computer would automatically proceed to Program 65 (P65) after passing through low gate. Velocity nulling guidance (P65) was the guidance computer's fully automatic landing mode. In this mode, the horizontal velocity rates were zeroed as the vehicle descended vertically at a constant rate to the surface.

If fuel efficiency (low landing ΔV) were the dominant trajectory constraint during landing, the trajectory could be designed without an approach maneuver so that braking would extend from PDI until the terminal descent maneuver. Adding the approach maneuver between the pitch-up and the final terminal descent compromises this optimal fuel efficiency to provide more flight time near the landing target and a better view of the target from the lander. This increased time and improved perspective are key aspects in providing an onboard HDA capability to the lander. The system level trades, which ultimately determine the design of the HDA capability and consequently the design of this approach maneuver, are the focus of the remainder of this paper.

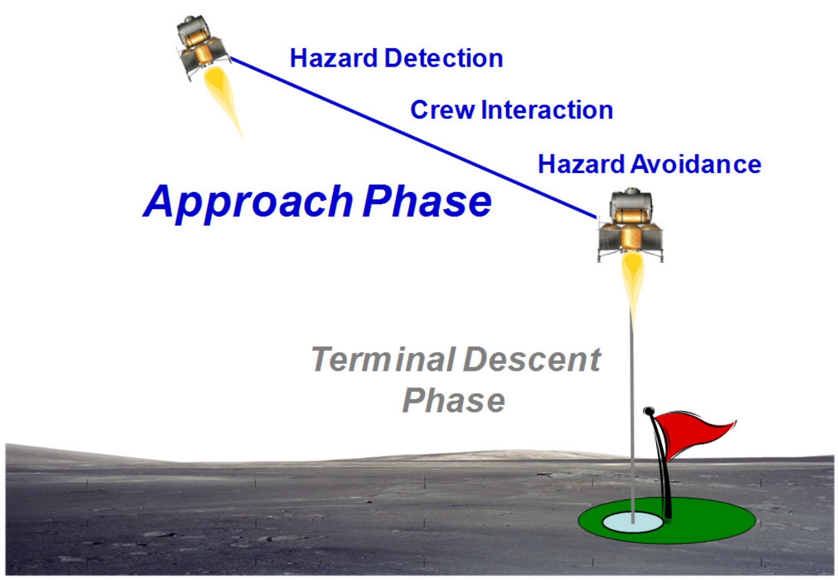

Figure 3 – HDA Activities

SYSTEM TRADEOFFS AND CONSTRAINTS FOR ONBOARD HDA

To provide an onboard HDA capability, the landing system must possess the ability to detect terrain hazards at the targeted landing location, be capable of avoiding those hazards, and still make a safe, soft landing. HDA begins after the lander has pitched up and is flying the approach phase, which provides an improved view of the landing area. The HDA function can be divided into three activities, each of which carries its own set of tradeoffs and constraints. These are Hazard Detection (HD), Crew Interaction (CI), and Hazard Avoidance (HA). Figure 3 highlights these three HDA activities along the trajectory profile.

Hazard Detection

Hazard Detection (HD) is the process of observing the terrain near the landing site and identifying the presence of hazards. An area that is hazard-free can generally be thought of as terrain that is level within some tolerance and free of small craters and boulders, which could tip the lander or impinge its underside[3]. As mentioned previously, Apollo accomplished this process through the onboard crew visually inspecting the landing area through the LM window after the Pitch-up maneuver. Several computer vision strategies exist for detecting hazards, which could employ a camera, stereo camera, radar, or LIDAR. ALHAT has chosen to invest in flash LIDAR development. While not the only viable approach for performing HD, flash LIDAR offers significant advantages over other technologies. It provides a direct, high-resolution measure of the terrain shape for evaluation and unlike a passive sensor is not dependent on ambient light (i.e. Sunshine or Earthshine) to operate. The terrain resolution of LIDAR is higher than is possible with radar, allowing it to image small but hazardous terrain features. Compared to scanning LIDAR technology, the flash variety acquires data much more quickly, enabling more terrain to be imaged during the landing.

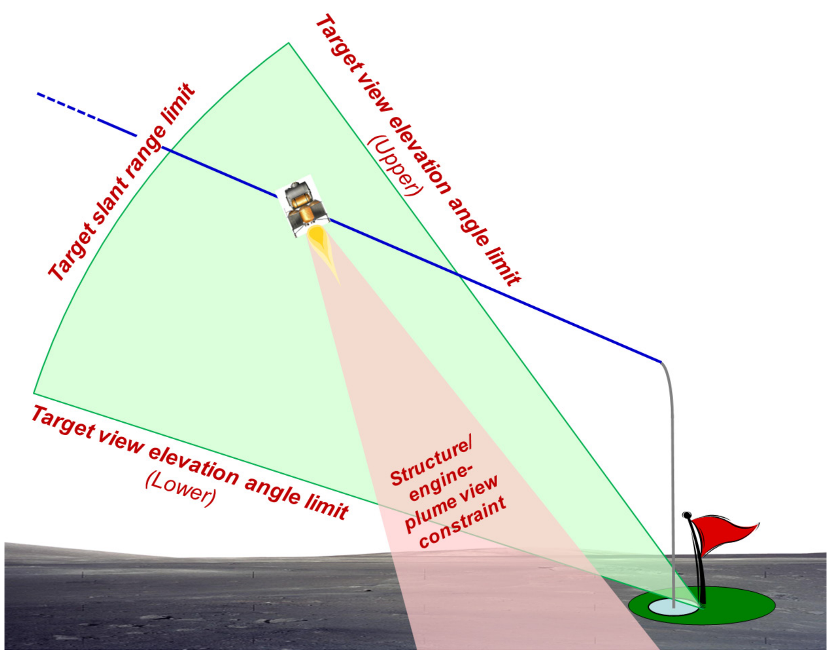

Figure 4 – Notional Hazard Detection Perspective Constraints

Several constraints need to be accommodated by the final system design for onboard HD to be successful. If the HD approach relies on ambient light (i.e. passive sensors like human eyes or cameras), then the timing and location of the landing must be chosen properly so that this light shines on the landing target with the proper elevation and azimuth. The trajectory design must provide proper perspective relative to the landing area and sufficient time to accomplish this detection process. The perspective constraints break down into three parts: a maximum slant range from the target to view hazards, a proper target line-of-sight elevation angle, and a proper unobstructed view of the landing target around the lander structure and/or engine plume (Figure 4). Time to detect hazards is needed for both the sensor hardware acquiring the data and the algorithms, which must interpret the raw data.

The numeric values placed against these constraints ultimately depends on the sensing technique used by the lander. For example, active hazard sensors take direct measurements of the terrain, so their performance is best when viewing the terrain from above (i.e. trajectory descents more vertically toward target). Passive sensors, on the other hand, use contrast variations and terrain shadows to infer the hazards and therefore are constrained to trajectory path angles slightly larger than the elevation angle of the ambient light, much like the constraints imposed on the Apollo crew[4]. Another example is the available view corridor around the structure/engine-plume obstructions. If this corridor is meant to support a window view, then the bounds will be limited by feasible window installation locations on the lander. If the crew capsule is sitting on top of the fuel tanks and engine, it is unlikely that a window view could be provided down through these things. If this corridor is to support a sensor, then perhaps it could be mounted on the underside of the vehicle to be less constrained by the lander structure. The amount of time necessary to perform the HD activity can vary significantly depending on the size of the landing area to be assessed, the method of acquiring the data (eyes, camera, LIDAR, etc), and the method used to process the terrain data.

Crew interaction

The Crew Interaction (CI) portion of HDA represents the time needed by the onboard crew to interpret the HD information (and any other landing-relevant information) and choose a safe point to land within the surveyed landing area. For Apollo this activity was seamless with the HD activity since Crew both visually inspected and evaluated the terrain for safe landing locations. If the lander is operating robotically then this aspect of the HDA process becomes the responsibility of an onboard algorithm to select a safe landing location from landing site hazard information gathered by onboard sensors. For ALHAT, this is the responsibility of the Autonomous Flight Manager (AFM), which uses this hazard information in conjunction with other metrics (e.g. fuel cost contours, proximity to targets of interest, etc.) to create a prioritized list of preferred safe landing locations. This information is provided to the crew (if present) to consider in choosing the final landing location. In robotic mode, the AFM automatically selects the top choice from this prioritized list. ALHAT also enables the crew to use their view out the window of the landing area (if properly lit) to aid in the selection of this safe landing target.

The constraints on the CI activity from a trajectory perspective is essentially time. The crew needs time to process the hazard information and make their decision. The time required depends on many aspects of the crew/system interface design[5]. For example, if the hazard information is in a very raw format (i.e. view out a window or camera view), then cognition time would be longer than if the crew is presented with a small list of discrete, vetted options. Once the crew has made a decision, this information needs to be input to the lander flight computer. The design of this interface would also play a factor in this CI time.

Hazard avoidance

Once the safe landing target is selected, the lander must accurately divert to this point in preparation for the terminal descent and touchdown. The GNC system in conjunction with the lander engine and RCS systems enable this Hazard Avoidance (HA) divert maneuver. This portion of the HDA process impacts the trajectory design because the trajectory profile must provide for reasonable vehicle states (i.e. position, velocity, attitude) at the time when the HA command is received so that the maneuver can be completed safely and successfully. This means that the vehicle state at the beginning of the HA maneuver must not preclude the vehicle from reaching any safe point within the landing area. This could happen, for example, if the vehicle were too close to touchdown or moving too fast toward the ground given its control authority (i.e. maximum engine acceleration and attitude rate). The constraints on the HA activity are therefore the maximum required divert distance, the lander control authority, and the allowed divert propellant budget.

SYSTEM DESIGN PROCESS FOR ONBOARD HDA

The constraints of these activities must be balanced against one another in order to properly design a trajectory that accommodates onboard HDA. This is a particularly challenging aspect of the lander system design since it involves the interaction of many vehicle systems, the crew, and the local lunar environment in a complex, coupled fashion. These interactions occur while the vehicle is moving along a high velocity trajectory with a rapidly narrowing time horizon and diminishing control authority needed to execute this maneuver to a safe landing location. Since the lander is mass and therefore propellant constrained, loiter time near the landing site is limited. This time constraint means that relaxing the constraints on one of the HDA activities effectively tightens the constraints on the others.

For a particular approach phase design, the trajectory provides a particular timeline, landing site perspective, etc. Adopting a trajectory design would affect the HDA activities according to their constraints and would yield a unique HDA system performance. To evaluate this HDA system performance, a broad tradespace of feasible approach phase trajectory designs were created. These trajectories were parameterized in terms of trajectory path angle at hazard detection, slant range at hazard detection, and the engine acceleration profile during approach. Figure 5 depicts this tradespace and shows the values that have been selected for these parameters. The intersection of each unique value for the three parameters produces a unique trajectory profile (currently 252 profiles exist within the tradespace). Figure 6 defines the initial slant range and path angle.

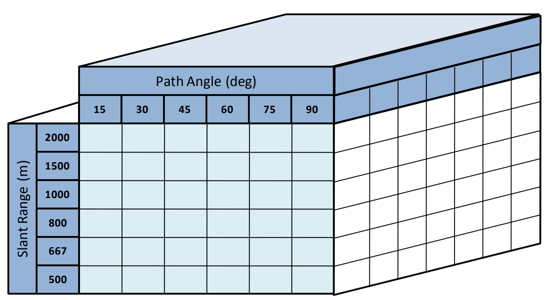

Figure 5 – ALHAT Trajectory Trade Space

The initial path angle parameter is chosen to vary through all possible values in 15-degree increments. The initial slant range varies from 500m to 2km to represent values considered reasonable for the start of the HD activity, which for ALHAT is performed by an active sensor system. The maximum value for the initial slant range could be extended further from the target if necessary (e.g. the Apollo approach began at approximately 8km slant range[2]), but for ALDAC-1 is limited to 2km since this nears reasonable range limits of a LIDAR sensor system. The acceleration profile varies from a constant 1.05 lunar-g to a 2.0 lunar-g design. On the low end, this results in a very slow, near-hovering trajectory while on the other end results in a very fast, dynamic landing. The acceleration profile is not required to be a constant, but is designed this way initially for simplicity*.

The HDA process was simulated along each of these trajectories. Any trajectories that do not provide adequate HDA performance can be eliminated from the tradespace, therefore better defining the approach phase design. By following this tradespace analysis strategy, the realm of HDA system designs can be evaluated and traded against one another to ultimately define the best system-optimized HDA capability and trajectory profile design.

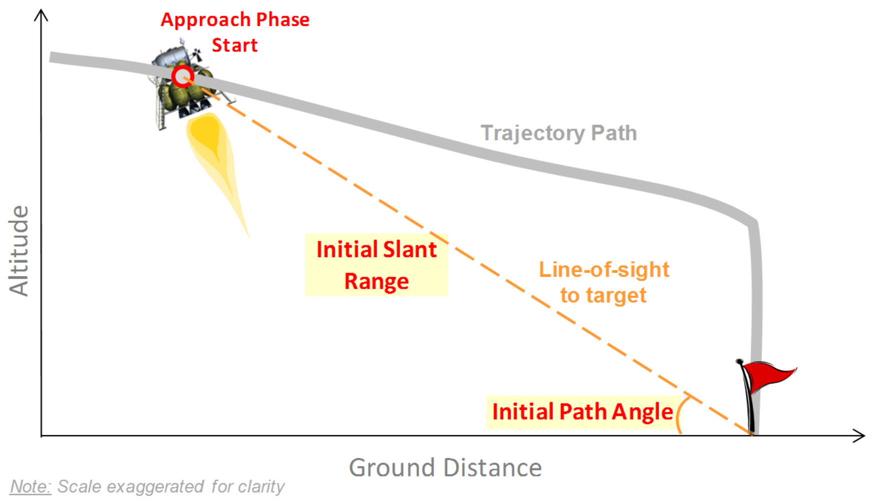

Figure 6 – Initial Slant Range and Path Angle Defined

TRAJECTORY ANALYSIS RESULTS FOR ONBOARD HDA

As of ALDAC-1, the details of the hazard detection and crew interaction activities are still under investigation. As a result, no final decisions have been made with regard to the trajectory design for ALHAT. The ALDAC-1 analysis evaluated across the trajectory tradespace, though, has provided valuable insight into how the trajectory should be designed to support onboard HDA. The following sections present examples of the ALDAC-1 trajectory analysis† and the resulting insights into the system design.

Hazard Detection Analysis

During the HD process, proper perspective relative to the landing area and sufficient time are key parameters that the approach phase trajectory design must accommodate. In terms of the tra-

---

* Additional and potentially more complex trajectory design schemes will be considered in future ALDAC's. The initial strategy is to start simple and evaluate the capability of the integrated system. Other trajectory schemes could be created which would tailor segments of the approach phase to better accommodate any constraints of the HDA activities.

† The ALHAT Guidance and Control algorithms were used closed-loop in a lunar landing simulation tool to create the data presented in this paper.

jectory design, this can be described by the slant range, trajectory path angle, and look angle time trends relative to the target.

This section focuses on the trends of the look angle and its impact on the trajectory design. The look angle is defined as the angle between the lander body-frame “down” direction (i.e. vehicle longitudinal axis) and the line-of-sight to something outside of the lander. The primary look angle of interest is the target look angle, defined as angle “a” in Figure 7. This angle represents the direction the crew and/or hazard detection sensor must look to see the landing target relative to the lander orientation. Having the proper target look angle is important since this determines whether the landing target is obstructed by the lander structure and/or engine plume.

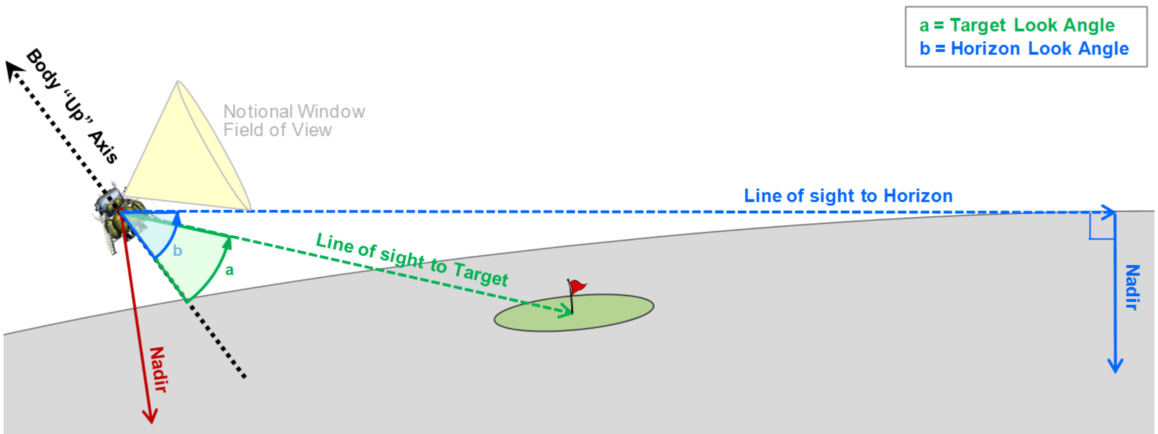

Figure 7 – Look Angles Defined

The time history of the look angle is a function of the lander position and orientation relative to the target of interest. The time history of the lander position and orientation are in turn a function of the approach phase trajectory design, which means that the design is important in determining both crew and sensor viewing of the landing site. To achieve a particular approach phase trajectory design, the thrust acceleration magnitude and pointing direction must follow a specific time history profile. Since the lander longitudinal axis defines the thrust acceleration pointing direction, this creates the relationship between the look angle and trajectory design.

The most fuel-efficient approach trajectory requires a maximum thrust, nearly constant retrograde attitude (i.e., a continuation of the braking phase), yet produces very unfavorable look angle profiles. At the expense of fuel, more favorable look angle profiles can be achieved by deviating from this timeline late in the braking phase. The most straightforward deviation, which was employed by Apollo and this initial ALHAT analysis, is a planned reduction in acceleration level from that of the braking phase. The reduction produces a pitch-up to a more vertical thrust attitude to counter gravitational losses at the lower acceleration level. Any acceleration that is applied laterally in a retrograde direction continues to reduce the lateral lander velocity in preparation for touchdown. This strategy (as was done in Apollo and in this ALDAC) produces a fairly constant approach attitude but variable look angle profile as will be demonstrated shortly. If greater control of the look angle is required, variable vertical and lateral acceleration profiles can be employed to either improve the initial look angle or provide a more constant look angle profile. This trajectory design strategy, though, would require more aggressive and variable approach attitude profiles and incur greater fuel penalties. The tradeoffs among these different strategies will be studied in detail in future ALHAT design analysis cycles.

Figure 8 shows the look angle trend versus time-to-go (until touchdown) for a sample approach phase design taken from the ALHAT trajectory tradespace. A look angle of zero degrees represents looking down through the floor of the lander, in the direction of the main engine exhaust. A look angle of +90 degrees represents looking straight ahead relative to the lander from the perspective of crew standing inside. The heavy red curve shows the target look angle trend with the red dot at 150 sec representing the end of the pitch-up maneuver and start of the approach phase where hazard detection begins. The green shaded region represents the look angle range necessary to see the entire landing site, assuming that the site is 180 meters in diameter*. The blue shaded region represents the additional look angle range needed to see the lunar horizon simultaneously with the landing site.

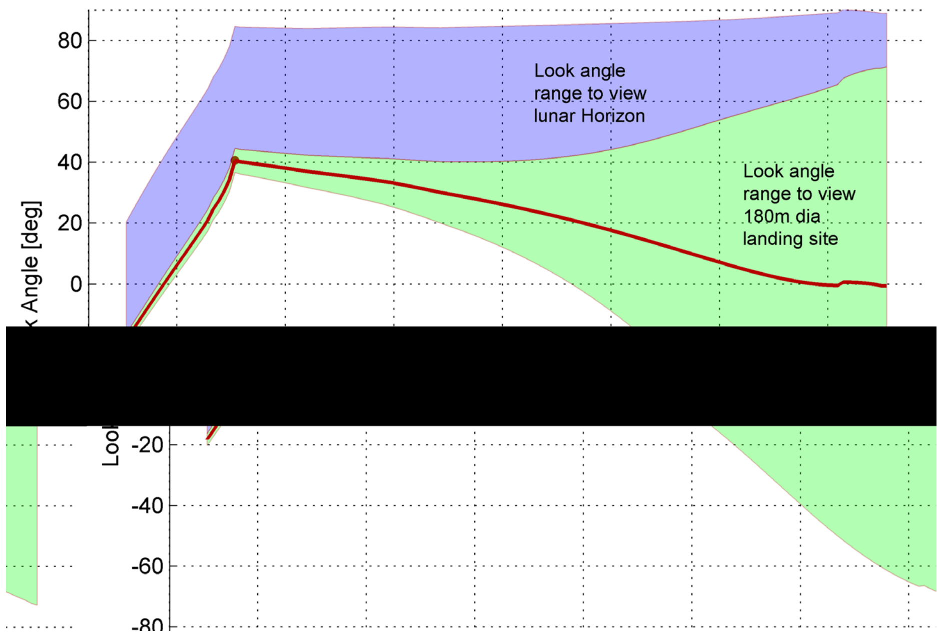

Figure 8 – Target Look Angle during Example Approach Phase Trajectory (1000m initial slant range, 45deg initial path angle, 1.05 lunar-g acceleration profile)

Crew Considerations

If the crew is performing the HD activity visually through a window for a specified duration, the minimum and maximum target look angle over that time specifies the required vertical field-of-view (FOV) and body location for the lander window. For example, if the crew required 30 seconds to perform the HD operation and need to see the entire landing site and horizon simultaneously then the lander window requirements can be extracted from Figure 8 for this sample trajectory. In this case the window would require at least a 60 degree full-cone FOV (min look angle of +25 degrees to max look angle of +85 degrees) and be centered at a 55 degrees look angle.

For the Apollo LM, the lower window edge was a +25 degree look angle[6], which meant that the landing target was only visible to the crew if it had a look angle greater than that value. If the lander window configuration is already defined, then the target look angle trend specifies how

* This landing site size was chosen because the ALHAT precision landing requirement is currently specified as 90 meters 3-sigma in radius.

much time the crew has available to observe the landing target. For this sample trajectory in Figure 8, a +25 degree target look angle limit would allow the crew to see the entire landing site for 30 seconds and the exact landing target (center of the landing site) for 50 seconds.

Sensor Considerations

If a sensor is performing the hazard detection activity then this same target look angle data is useful. If the sensor is rigidly attached to the lander then it is subject to the same FOV and body location/orientation issues as the crew window. For example, a rigidly mounted camera that is required to observe the landing site and horizon simultaneously for 30 seconds would similarly need a 60-degree full-cone FOV and mounted with its boresight centered at a 55-degrees look angle.

If the sensor were gimbal mounted to track the landing site during the approach phase, then Figure 8 shows that for this same 30 second observation period the gimbal would require an 8-degree traverse to keep its boresight centered on the site (+41 degree to +33 degree look angle) with a maximum FOV of 16 degrees (full-cone) in order to see the entire landing site.

Look Angle Sensitivities

Figure 9 shows the sensitivity of the target look angle time history to the approach phase trajectory design. For this example, 27 trajectories from the ALHAT trajectory trade space were selected in order to show how the look angle trend changes when adjusting the initial approach phase slant range, path angle, and acceleration profile.

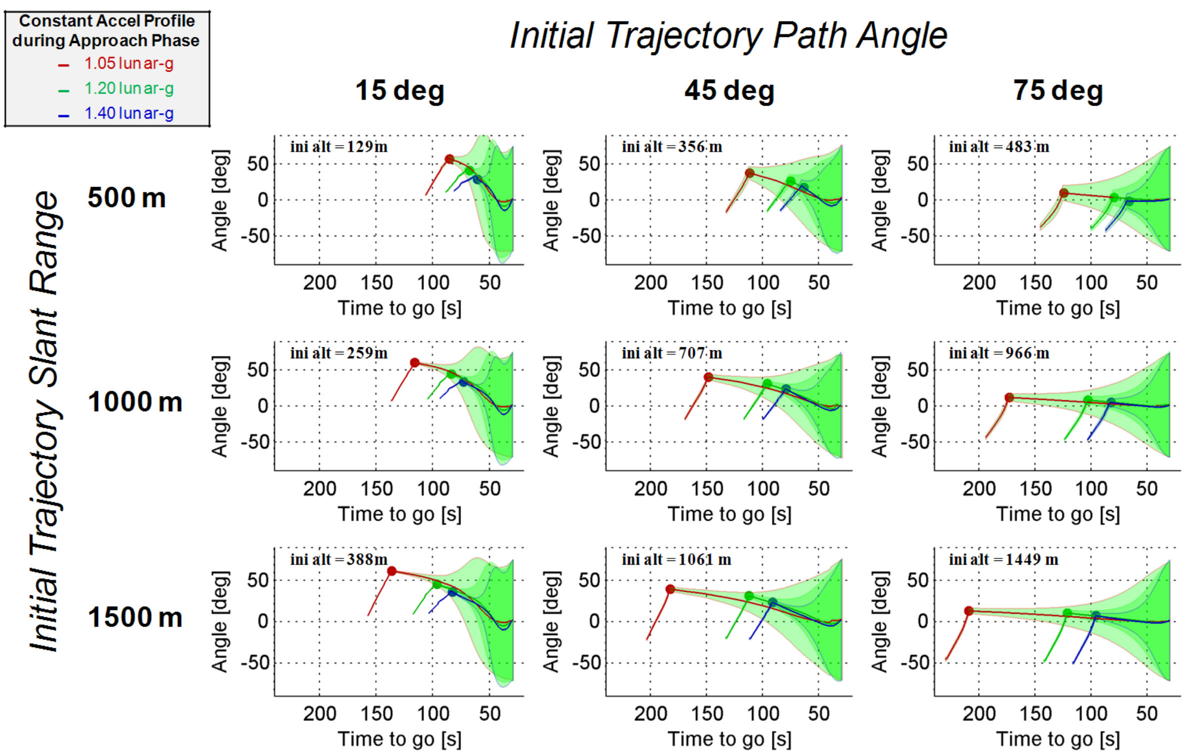

Figure 9 – Sensitivity of Target Look Angle Time History to Varying Approach Phase Trajectory Design

One clearly visible trend is that using a lower acceleration profile or lengthening the initial slant range extends the approach phase timeline and consequently provides more time to view the landing site, all other parameters being equal. Another trend is that designing the approach phase trajectory with a shallower initial path angle yields a larger initial look angle (generally better for viewing the site around lander structure and/or engine plume obstructions), but also undergoes a greater total look angle change during approach (possibly requiring a larger window/sensor FOV or gimbal traverse to view the landing site). This same trend is observed by lowering the acceleration profile in order to fly this phase more slowly (i.e. blue to green to red curves). The 75-degree initial path angle trajectories all have target look angles of less than +15 degrees, which indicates that providing the crew with a window view of the target is unlikely assuming a lander window configuration similar to the Apollo LM. A hazard detection sensor (or camera, in lieu of a window), on the other hand, could likely be mounted on the lander to allow it to see closer to a 0 degree look angle, possibly allowing these steep trajectories to be feasible for a sensor-based hazard detection process. An obstructed view due to the engine plume might still be a problem for these steep approach phase scenarios.

Initial Target Look Angle [deg]

	Slant Range (m)	Trajectory Path Angle (deg)					
		15	30	45	60	75	90
A=1.05 (lunar-g's)	2000	64	55	40	27	14	-1
	1500	62	53	39	27	13	1
	1000	60	52	41	26	12	1
	800	61	51	40	25	11	2
	667	59	51	40	25	10	0
	500	59	51	38	23	10	1
A=1.1	2000	58	51	37	24	13	0
	1500	56	49	36	25	12	0
	1000	55	47	37	26	10	0
	800	53	48	36	24	9	1
	667	54	46	36	23	8	0
	500	52	45	33	20	7	1
A=1.2	2000	51	44	32	21	11	0
	1500	46	42	31	20	10	0
	1000	46	39	32	22	8	1
	800	45	40	31	20	6	-1
	667	46	39	30	18	5	1
	500	45	37	27	16	4	1
A=1.3	2000	44	39	28	19	10	0
	1500	41	36	27	18	9	0
	1000	41	34	27	19	6	2
	800	40	34	26	17	5	1
	667	39	33	26	15	4	2
	500	38	33	23	11	2	0
A=1.4	2000	41	36	25	17	9	0
	1500	38	33	23	15	7	0
	1000	36	30	23	16	5	1
	800	37	30	24	15	2	1
	667	35	30	23	12	2	1
	500	34	29	20	11	1	2
A=1.5	2000	36	31	21	14	8	0
	1500	35	30	21	13	7	0
	1000	32	27	21	14	4	0
	800	32	27	22	13	1	1
	667	32	27	20	10	0	2
	500	31	26	17	9	1	2
A=2.0	2000	25	23	15	8	5	0
	1500	23	20	15	7	4	2
	1000	23	18	11	10	3	0
	800	22	18	13	7	3	0
	667	22	17	13	5	4	1
	500	21	16	11	4	3	1

Time Target Look Angle > 25 deg [sec]

	Slant Range (m)	Trajectory Path Angle (deg)					
		15	30	45	60	75	90
A=1.05 (lunar-g's)	2000	88	98	82	15		
	1500	73	80	66	19		
	1000	54	59	53	9		
	800	45	50	44	2		
	667	39	44	38			
	500	31	35	28			
A=1.1	2000	67	67	49			
	1500	55	54	40	1		
	1000	41	40	30	3		
	800	35	34	26			
	667	31	29	22			
	500	23	22	15			
A=1.2	2000	46	40	21			
	1500	38	31	16			
	1000	27	20	12			
	800	23	18	9			
	667	20	15	7			
	500	15	12	3			
A=1.3	2000	35	23	6			
	1500	25	17	4			
	1000	19	11	3			
	800	15	9	2			
	667	14	8	1			
	500	11	5				
A=1.4	2000	24	14				
	1500	17	9				
	1000	13	6				
	800	10	5				
	667	9	4				
	500	6	2				
A=1.5	2000	16	7				
	1500	11	5				
	1000	7	3				
	800	6	2				
	667	4	2				
	500	1	1				
A=2.0	2000						
	1500						
	1000						
	800						
	667						
	500						

Time Target Look Angle > 5 deg [sec]

	Slant Range (m)	Trajectory Path Angle (deg)					
		15	30	45	60	75	90
A=1.05 (lunar-g's)	2000	105	127	144	149	128	
	1500	88	106	118	126	102	
	1000	67	82	94	95	72	
	800	58	71	81	80	55	
	667	51	63	71	70	43	
	500	40	52	59	56	35	
A=1.1	2000	81	93	102	103	87	
	1500	68	79	86	87	65	
	1000	53	60	66	69	46	
	800	46	52	58	57	35	
	667	41	45	51	50	23	
	500	32	36	41	39	13	
A=1.2	2000	61	67	69	65	52	
	1500	51	55	57	53	39	
	1000	39	41	44	42	23	
	800	34	36	37	36	9	
	667	30	31	33	30		
	500	24	25	25	23		
A=1.3	2000	52	52	52	47	31	
	1500	42	43	42	39	22	
	1000	32	32	32	29	8	
	800	27	27	27	24		
	667	24	24	23	19		
	500	20	18	18	13		
A=1.4	2000	46	43	39	33	19	
	1500	36	34	31	25	11	
	1000	28	26	23	20	1	
	800	23	22	20	17		
	667	21	19	17	12		
	500	18	15	13	6		
A=1.5	2000	40	36	29	22	12	
	1500	31	29	23	17	7	
	1000	25	21	18	14		
	800	22	18	16	11		
	667	19	15	13	8		
	500	15	12	9	2		
A=2.0	2000	25	16	9	4	1	
	1500	19	12	8	3		
	1000	14	8	5	4		
	800	12	7	5	2		
	667	11	6	4	1		
	500	7	4	2			

Figure 10 – Example Target Look Angle Trends across ALHAT Trajectory Tradespace

Figure 10 shows three example target look angle trends as a function of the entire ALHAT trajectory tradespace. The first example is the initial target look angle at the start of the approach phase. This corresponds to the look angle value at the dots in the previous figure. The color shading changes from blue to fuchsia for large to small look angles, respectively. This shows that having a slow and shallow approach phase design (upper left region) provides a larger initial target look angle, which is consistent with the trend visible in the previous figure. The converse situation (fast and steep approach design) provides a small initial target look angle. Figure 11 below shows two example trajectories from the tradespace to more clearly illustrate this effect of the trajectory acceleration profile on the look angle behavior. In this figure, both trajectories have the same initial approach phase slant range and path angle. The only difference is that the one on the left flies a constant 1.05 lunar-g acceleration profile while the other flies a 2.0 lunar-g profile. What this shows is that for all other parameters being equal, a faster approach phase trajectory designs require the vehicle to be oriented more horizontally in order to reduce the high lateral velocity. This more horizontal orientation results in a smaller target look angle throughout the approach phase (but is more fuel-efficient).

The second and third examples in Figure 10 show the time during the approach phase when the target look angle is great than 25 degrees and 5 degrees respectively. The color shading varies from green (long times) to red (short times). Any grayed-out cells in these two examples represent trajectories where the look angle is never greater than the selected value. The 25-degree limit was selected to illustrate how a crew window-based HD strategy might constrain the approach phase trajectory design. With this strategy, only the slower and shallower trajectories from this tradespace are viable options. The 5-degree limit was selected to illustrate a sensor based hazard detection strategy with the assumption that a more favorable positioning of the sensor as compared to a window will result in a more lenient look angle constraint. With both of these examples, further narrowing of the trajectory tradespace will occur as the time required for viewing the target increases. For example, if both scenarios required at least 30 seconds to perform HD, then only the trajectory cases noted in yellow-green to green shades would be viable. To put these times in context with the early Apollo trajectory design, the Apollo look angle was greater than 25 degrees for approximately 35 seconds beginning at 1km slant range and approximately 60 seconds beginning at 2km[6]. This result compares favorably with the similar ALHAT approach trajectory (e.g. 41 seconds for the 1km slant range, 15-degree path angle, 1.1 lunar-g acceleration profile. 67 seconds for the 2km slant range, 15-degree path angle, 1.1 lunar-g acceleration profile.).

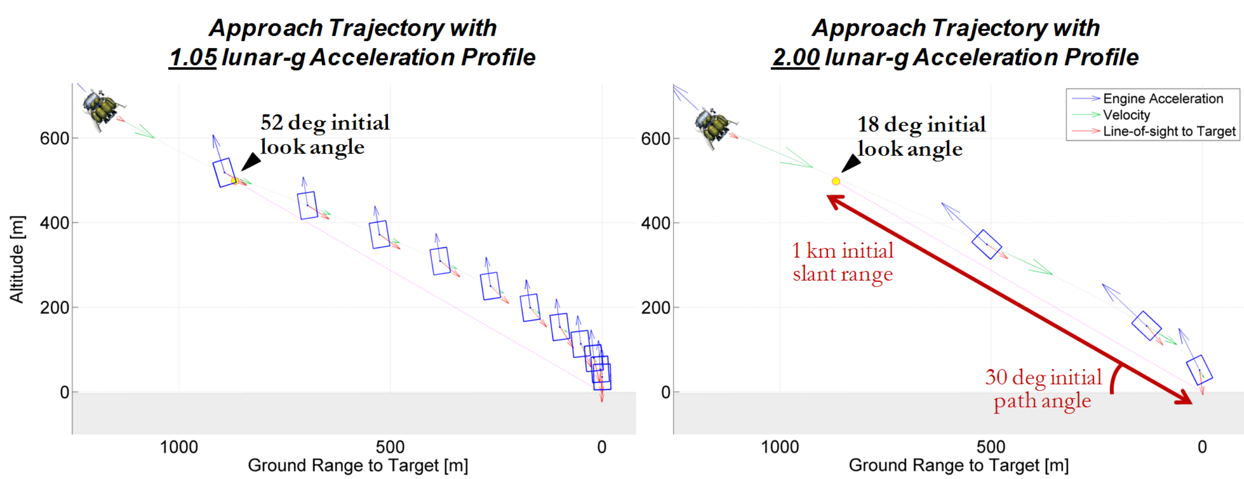

Figure 11 – Comparison of Lander Orientation and Target Look Angle Behavior as a Function of Acceleration Profile

This section discussed how the target look angle and landing site FOV trends represent important perspective and timing issues relevant for the hazard detection activity. Further, it showed how specific look angle requirements, based on a particular HD technique, would constrain the trajectory design. There are other perspective and timing issues relevant to HD that would impose additional constraints on the trajectory design. One example is the slant range time history, which will affect the sensor power to image the target and crew/sensor resolution of the terrain at the target. Another is the trajectory path angle time history, which determines the perspective relative to the local horizontal and effects how an active sensor signal bounces back from the ground to the lander and how a passive sensor/crew see the terrain in terms of local lighting and shadows. Similar to the look angle constraints, both the initial value and amount of change over time for the slant range and path angle affect the HD activity. Consideration of these perspective and timing issues relevant to the selected HD technique will impose constraints to the approach phase design and lead to an appropriate trajectory to support onboard hazard detection.

Hazard Avoidance Analysis

Once the hazards have been detected and the crew or AFM have selected a safe landing point, the lander must execute a hazard avoidance divert maneuver to reach this point. The constraints on this HA activity are the maximum required divert distance, the lander control authority, and the allowed divert propellant budget. The maximum divert distance will be a requirement imposed on the lander and will be limited by the terrain area scanned for hazards. For ALHAT this means the lander will need to be capable of diverting up to the edge of the area scanned by the hazard detection LIDAR system, which currently is assumed to range from 45m to 360m across. The other two HA constraints are driven by the lander design. For ALDAC-1 the lander control authority and divert propellant budget were significant, but the guidance mechanization for this analysis was limited to a fixed approach phase time based on the nominal trajectory design. The impact of this is that the lander must receive the divert command sufficiently early for a given trajectory design and divert distance so that it is complete before the terminal descent altitude is reached. Future ALDACs will allow greater flexibility in guidance so that the approach phase time can be extended as needed to execute a divert. This means that the HA divert maneuver can be any distance and last any amount of time as long as there is sufficient propellant to execute the maneuver. Since the propellant budget for this maneuver will be limited, guidance will still require that the divert command be received sufficiently early to limit propellant consumption.

This section focuses on the propellant (ΔV) use and timing related to the HA divert maneuver for ALDAC-1. Figure 12 shows the total lunar landing ΔV and approach phase duration for each of the trajectories within the ALHAT trajectory tradespace. These data correspond to nominal trajectories with no HA maneuver. The nominal ΔV is measured from the deorbit burn to touchdown. The approach phase duration is measured as the time between the end of pitch-up and terminal descent. Figure 12 shows that there is a significant variation of the nominal required ΔV, which translates into required propellant. Based on this information alone, the trajectory design could be further constrained if the landing ΔV budget is known. The figure also shows how significant the approach phase duration varies across the tradespace. There is an order-of-magnitude difference in the nominal duration from the minimum to the maximum values.

For the HA maneuver, a key metric of interest is the maximum time the maneuver can be delayed while still being able to achieve some maximum divert distance. This maximum delay time is important because it is during this period that the hazard detection and crew interaction functions must occur for onboard HDA to be successful. If the combined HD and CI time exceed this maximum delay time, then the lander risks not being able to reach the selected safe landing point. As mentioned previously, future ALDACs will allow guidance to extend the approach phase timeline as needed to accomplish a HA divert. For that situation, the maximum delay time equals the nominal approach phase duration shown in Figure 12. For ALDAC-1, though, the maximum delay time was a function of the divert distance.

Nominal Landing ΔV [m/s]

A (lunar-g's)	Slant Range (m)	Trajectory Path Angle (deg) 15	30	45	60	75	90
A=1.05	2000	2006	2064	2117	2153	2179	2192
	1500	1983	2030	2075	2107	2130	2142
	1000	1953	1989	2022	2056	2073	2082
	800	1938	1970	2005	2030	2046	2054
	667	1929	1958	1988	2009	2026	2032
	500	1910	1940	1965	1980	1995	2003
A=1.1	2000	1958	1997	2034	2061	2083	2096
	1500	1940	1975	2006	2031	2050	2061
	1000	1920	1948	1972	1992	2007	2017
	800	1911	1934	1957	1975	1988	1996
	667	1903	1925	1945	1962	1975	1982
	500	1891	1910	1928	1944	1955	1962
A=1.2	2000	1918	1945	1972	1996	2016	2031
	1500	1906	1929	1953	1974	1991	2005
	1000	1892	1911	1931	1948	1961	1974
	800	1886	1902	1920	1935	1948	1958
	667	1881	1896	1912	1927	1938	1948
	500	1874	1887	1901	1913	1922	1932
A=1.3	2000	1900	1921	1943	1967	1987	2005
	1500	1889	1909	1928	1948	1966	1982
	1000	1879	1894	1911	1928	1943	1955
	800	1873	1888	1903	1917	1931	1942
	667	1869	1883	1897	1910	1923	1934
	500	1865	1875	1887	1899	1911	1920
A=1.4	2000	1888	1905	1927	1949	1971	1990
	1500	1879	1895	1915	1934	1953	1970
	1000	1870	1883	1899	1915	1931	1945
	800	1866	1877	1892	1906	1920	1934
	667	1863	1873	1887	1900	1913	1925
	500	1859	1868	1879	1891	1902	1913
A=1.5	2000	1879	1895	1916	1938	1960	1982
	1500	1870	1886	1905	1924	1943	1963
	1000	1864	1876	1891	1907	1923	1940
	800	1861	1871	1885	1899	1914	1929
	667	1859	1867	1879	1893	1907	1921
	500	1856	1862	1873	1885	1897	1910
A=2.0	2000	1856	1867	1886	1911	1938	1969
	1500	1852	1863	1879	1900	1924	1950
	1000	1848	1856	1870	1887	1908	1930
	800	1846	1853	1866	1881	1900	1919
	667	1844	1851	1862	1877	1894	1913
	500	1843	1848	1858	1871	1886	1903

Approach Phase Duration [sec]

A (lunar-g's)	Slant Range (m)	Trajectory Path Angle (deg) 15	30	45	60	75	90
A=1.05	2000	120	151	176	194	205	208
	1500	102	129	151	166	175	178
	1000	81	103	120	133	140	143
	800	71	90	106	117	124	126
	667	64	81	95	105	112	114
	500	53	68	80	89	94	96
A=1.1	2000	95	113	129	140	147	150
	1500	81	97	111	121	127	129
	1000	65	78	89	97	102	103
	800	58	69	78	86	90	91
	667	52	62	71	77	81	83
	500	44	52	60	66	69	70
A=1.2	2000	75	86	95	101	106	107
	1500	65	74	81	87	91	92
	1000	52	59	66	70	73	74
	800	46	53	58	62	65	66
	667	42	48	53	56	59	60
	500	36	40	45	48	50	51
A=1.3	2000	66	73	79	84	87	88
	1500	57	63	68	72	75	76
	1000	46	51	55	58	60	61
	800	41	45	49	52	54	54
	667	37	41	44	47	49	49
	500	32	35	38	40	41	42
A=1.4	2000	60	65	70	73	76	76
	1500	51	56	60	63	65	66
	1000	42	45	49	51	53	53
	800	37	40	43	45	47	47
	667	34	37	39	41	42	43
	500	29	31	33	35	36	37
A=1.5	2000	55	60	63	66	68	68
	1500	48	51	55	57	58	59
	1000	39	42	44	46	47	48
	800	34	37	39	41	42	42
	667	31	34	36	37	38	38
	500	27	29	30	32	33	33
A=2.0	2000	43	45	47	48	48	49
	1500	37	39	40	41	42	42
	1000	30	32	33	33	34	34
	800	27	28	29	30	30	30
	667	25	26	26	27	27	27
	500	21	22	23	23	23	23

Figure 12 – Nominal Landing ΔV and Approach Phase Duration across ALHAT Trajectory Tradespace

Table 1 provides a summary of the results for an ALDAC-1 study, which determined the maximum delay time and resulting additional ΔV cost to execute maximum HA diverts for each trajectory within the trajectory tradespace[7]. The first column of the table is the size of the region scanned for hazards by the HD sensor system. The second column is the divert distance which corresponds to half of the HD area's diameter minus the lander's footprint radius. For this example, the lander radius is assumed to be 10 meters, which also includes a buffer around the vehicle to account for touchdown accuracy limitations. The next column provides time metrics on the reduction in the approach phase duration needed to enable this HA divert. The last column provides metrics on the additional ΔV cost incurred by this max delayed max divert. This summary table is provided instead of the detailed trajectory tradespace results because the values only weakly correlate with the variations in the approach phase trajectory design. The dominant factor affecting these divert-related times and ΔV costs for ALDAC-1, then, is primarily the required max divert distance. The total HA divert performance results are found by combining the divert data from Table 1 with the nominal data found in Figure 12. For example, the maximum delay time before executing the HA divert (to allow for the HD and CI activities) is found by subtracting the Table 1 time values (column 3) from the nominal approach phase times from Figure 12. By comparing the magnitude of the data between the nominal and divert data, it is clear that the nominal performance trends dominate the overall HA divert performance results.

Similar to the look angle durations presented in Figure 10, this max delay max divert data can be used to constrain the trajectory design once the time requirements for the HD and CI activities are defined. For example, if detecting terrain hazards and selecting a safe site takes at least 40 seconds, then none of the trajectories shaded orange in Figure 12 are viable. Another issue is that larger areas scanned for hazards both require more time to complete the HD/CI activities and have less delay time available due to the larger required divert.

Table 1 – Summary of Hazard Avoidance Performance Results for Max Delayed Max Divert

Hazard Detection Area Diameter [m]	Divert Distance [m]	Approach Phase Time Reduction Due to HA Divert [s] { **mean** *(min–max)* }	Additional ΔV Cost Due to HA Divert [m/s] { **mean** *(min–max)* }
45	12.5	**6.1** *(1.9–7.2)*	**1.5** *(0.8–2.0)*
90	35	**8.5** *(5.1–9.2)*	**5.8** *(4.8–7.0)*
180	80	**11.6** *(8.7–13.1)*	**12.7** *(10.7–15.3)*
360	170	**15.2** *(12.5–17.8)*	**25.0** *(18.8–30.1)*

CONCLUSION

This paper focused on the tradeoffs and constraints surrounding the design of the final approach and landing trajectory to support an onboard HDA capability for a future lunar lander. The overall trajectory profile is a fundamental design parameter to the lander system and must be designed such that it provides a low ΔV solution to minimize consumed propellant mass during landing. The trajectory profile must also accommodate the various other needs of the lander, in this case onboard HDA. Hazard detection can be accomplished by the crew visually scanning the terrain through a window, a sensor system imaging the terrain, or some combination of both. The trajectory design will be greatly influenced by the hazard detection activity since it will supply the proper perspective, proper range to the landing site, and sufficient time to view the terrain. Once the terrain has been viewed, the trajectory design must provide additional time to process this information and allow a decision to be made about where to safely land. Finally, the trajectory design must provide sufficient margin to enable a hazard avoidance maneuver in order to reach this safe landing location.

A tradespace of feasible approach phase trajectory designs were investigated to best understand the impacts and trades associated with a lunar lander HDA capability. These trajectories were parameterized in terms of trajectory path angle at hazard detection, slant range at hazard detection, and the engine acceleration profile during approach. Each of the values were varied to represent the broadest range of feasible trajectory designs to be considered by a future lunar lander in balance with projected capabilities of a future hazard detection system. On one end of the spectrum, a very slow, near-hovering trajectory while on the other end, a very fast, dynamic landing. Regardless where in the spectrum a particular trajectory design is chosen, adopting it affects the hazard detection and avoidance activities according to their constraints, which yields a unique HDA system performance.

The time history of the lander position and orientation are a function of the approach phase trajectory design, which are important in determining both crew and sensor viewing of the landing site. To achieve a particular approach phase design, the thrust acceleration magnitude and pointing direction must follow a specific time history profile. Since the lander longitudinal axis defines the thrust acceleration pointing direction, this creates the relationship between the look angle and trajectory design. Having the proper target look angle is important since this determines whether the landing target is obstructed by the lander structure and/or engine plume. The most fuel-efficient approach trajectory requires a maximum thrust, nearly constant retrograde attitude (i.e., a continuation of the braking phase), yet produces very unfavorable look angle profiles. If greater control of the look angle is required, variable vertical and lateral acceleration profiles can be employed to either improve the initial look angle or provide a more constant look angle profile. This trajectory design strategy, though, would require more aggressive and variable approach attitude profiles and incur greater fuel penalties.

If the crew is performing the hazard detection activity visually through a window for a specified duration, the minimum and maximum target look angle over that time specifies the required vertical field-of-view (FOV) and body location for the lander window. Using a lower acceleration profile or lengthening the initial slant range extends the approach phase timeline and consequently provides more time to view the landing site, all other parameters being equal. Designing the approach phase trajectory with a shallower initial path angle yields a larger initial look angle (generally better for viewing the site around lander structure and/or engine plume obstructions), but also undergoes a greater total look angle change during approach (possibly requiring a larger window/sensor FOV or gimbal traverse to view the landing site). A hazard detection sensor (or camera, in lieu of a window), on the other hand, could likely be mounted on the lander to allow it to see closer to a 0-degree look angle, possibly allowing steep trajectories to be feasible for a sensor-based hazard detection process. An obstructed view due to the engine plume might still be a problem for these steep approach phase scenarios.

Once the hazards have been detected and the crew or AFM have selected a safe landing point, the lander must execute a hazard avoidance divert maneuver to reach this point. The constraints on this hazard avoidance activity are the maximum required divert distance, the lander control authority, and the allowed divert propellant budget. The maximum divert distance will be a requirement imposed on the lander and will be limited by the terrain area scanned for hazards. Analysis within this paper has shown how significant the approach phase duration varies across the tradespace indicating an order-of-magnitude difference in the nominal duration from the minimum to the maximum values. Additionally, if the combined hazard detection and crew interaction time exceed a maximum hazard avoidance delay time, then the lander risks not being able to reach the selected safe landing point.

This paper discussed how the target look angle and landing site FOV trends represent important perspective and timing issues key to the hazard detection activity. Further, it showed how specific look angle requirements, based on a particular hazard detection technique, would constrain the trajectory design and that for all other parameters being equal, a faster approach phase trajectory requires the vehicle to be oriented more horizontally in order to reduce the high lateral velocity. This more horizontal orientation results in a smaller target look angle throughout the approach phase but has the attractive quality of being more fuel-efficient.

Hazard detection and avoidance, albeit as demonstrated by Apollo or on our future return to the Moon, will continue to be a critical activity for all safe lunar landings. Accordingly, lunar landing trajectory designs that support onboard HDA must be examined in an overall system design process to develop a robust, safe landing technique for the next generation of lunar landing vehicles.

ACRONYM LIST

AFM: Autonomous Flight Manager
AGNC: Autonomous Guidance, Navigation, and Control
ALDAC: ALHAT Design Analysis Cycle
ALHAT: Autonomous Landing and Hazard Avoidance Technology
CI: Crew Interaction
FOV: Field of View
GNC: Guidance, Navigation, and Control
HA: Hazard Avoidance
HD: Hazard Detection
HDA: Hazard Detection and Avoidance
LIDAR: Light Detection and Ranging
LM: (Apollo) Lunar Module
PDI: Powered Descent Initiation
TRN: Terrain Relative Navigation

REFERENCES

[1] Paschall. S.C., Brady, T., Cohanim, B.E., Sostaric, R. “A Self Contained Method for Safe & Precise Lunar Landing”. IEEE Aerospace Conference. March 2008.

[2] Klumpp, A.R., “Apollo Lunar-Descent Guidance” MIT Instrumentation Lab Report R-695. 1971.

[3] Barton, G., “Status Report on Trade Study – FLO 92:38 Safe Lunar Landing Definition”, C.S. Draper Laboratory, ESB-93-122A, 22 March 1993.

[4] Cappellari, Jr., J. O., “Where on the Moon? An Apollo Systems Engineering Problem," The Bell System Technical Journal 51 (5) (1972): 955.

[5] Chua, Z.K. and Major, L.M., "Task Modeling for Lunar Landing Point Redesignation", accepted for AIAA Info-Tech@Aerospace, Seattle, WA, April 2009. AIAA-2009-2056.

[6] Klumpp, A.R., “A Manually Retargeted Automatic Landing System for the Lunar Module (LM)”, Journal of Spacecraft and Rockets, Vol 5 No 2, Feb 1968.

[7] Cohanim, B.E., Fill, T.J., Paschall, S.C., Major, L.M., Brady, T., “Approach Phase ΔV Considerations for Lunar Landing”, IEEE Aerospace Conference. March 2009.