Flight testing a Real-Time Hazard Detection System for Safe Lunar Landing on the Rocket-Powered Morpheus Vehicle

Nikolas Trawny*, Andres Huertas†, Michael E. Luna‡, Carlos Y. Villalpando§, Keith E. Martin¶, John M. Carson‖, and Andrew E. Johnson**

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109

Carolina Restrepo††

NASA Johnson Space Center, Houston, TX 77058

Vincent E. Roback‡‡

NASA Langley Research Center, Hampton, VA 23681

The Hazard Detection System (HDS) is a component of the ALHAT (Autonomous Landing and Hazard Avoidance Technology) sensor suite, which together provide a lander Guidance, Navigation and Control (GN&C) system with the relevant measurements necessary to enable safe precision landing under any lighting conditions. The HDS consists of a stand-alone compute element (CE), an Inertial Measurement Unit (IMU), and a gimbaled flash LIDAR sensor that are used, in real-time, to generate a Digital Elevation Map (DEM) of the landing terrain, detect candidate safe landing sites for the vehicle through Hazard Detection (HD), and generate hazard-relative navigation (HRN) measurements used for safe precision landing. Following an extensive ground and helicopter test campaign, ALHAT was integrated onto the Morpheus rocket-powered terrestrial test vehicle in March 2014. Morpheus and ALHAT then performed five successful free flights at the simulated lunar hazard field constructed at the Shuttle Landing Facility (SLF) at Kennedy Space Center, for the first time testing the full system on a lunar-like approach geometry in a relevant dynamic environment. During these flights, the HDS successfully generated DEMs, correctly identified safe landing sites and provided HRN measurements to the vehicle, marking the first autonomous landing of a NASA rocket-powered vehicle in hazardous terrain. This paper provides a brief overview of the HDS architecture and describes its in-flight performance.

I. Introduction

Since 2006, the cross-NASA center ALHAT project has been conducting focused technology development for sensors systems and associated algorithms that allow autonomous and safe precision-landing on the Moon under any lighting conditions.[1] The ALHAT sensor suite consists of a flash LIDAR-based Hazard

*GN&C Hardware and Testbed Development Group, AIAA Member.
†Maritime & Aerial Perception Systems Group.
‡Guidance & Control Operations Group, AIAA Senior Member.
§Advanced Computer Systems & Technologies Group.
¶Simulation & Support Equipment Group.
‖GN&C Hardware and Testbed Development Group, currently on IPA at JSC, AIAA Associate Fellow.
**Guidance & Control Section Staff, AIAA Senior Member.
††Integrated GN&C Analysis Branch, AIAA Member.
‡‡Remote Sensing Flight Systems Branch, AIAA Senior Member.

Detection System (HDS),[2] a Doppler LIDAR velocimeter,[3] and a long-range Laser Altimeter.[4] The HDS was developed as a stand-alone, bolt-on system to detect available safe landing sites within a 100 m × 100 m region of the Lunar surface from 1 km slant range, 500 m altitude, in real-time and under any surface lighting conditions, and to subsequently provide hazard-relative position measurements to the Host Vehicle (HV)'s GN&C subsystem to improve landing accuracy.[5,6]

The current generation HDS system was fully integrated with the Morpheus avionics and tested during a helicopter flight test campaign (Field Test #5) at the simulated lunar hazard field at KSC in December 2012.[7] These flight test demonstrated the full range of HDS operations – DEM generation, safe site detection, and hazard relative navigation – and provided rich datasets for performance evaluation and tuning. In March 2014, the ALHAT was integrated onto the Morpheus vehicle and performed a total of five free flights, FF10–FF14, at the KSC SLF. This paper describes the HDS performance during this flight campaign.

II. Hazard Detection System Architecture

The HDS consists of a stand-alone compute element (CE), a power distribution unit (PDU) and battery box (BB), a flash LIDAR Sensor Head (FLSH), a flash LIDAR electronics box (LEB), a two-axis gimbal, and a dedicated inertial measurement unit (IMU) (see Fig. 1).

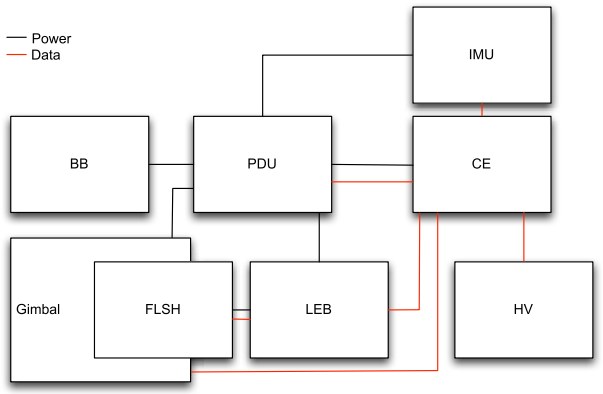

(a)

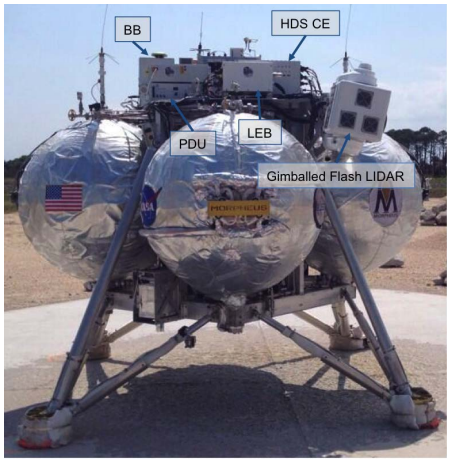

(b)

Figure 1. HDS blockdiagram, and installed on the Morpheus vehicle.

The HDS CE combines a Xilinx FPGA (surrogate for the path-to-flight Virtex-5 FPGA) and a multicore Tilera Tile-64 multicore processor (chosen for its common architecture with the Maestro processor developed under the Opera program[8]), linked through a high-speed, 10 Gb/s XAUI connection.[9] The FPGA provides a precise hardware-based time stamp for all sensor data, and a time synchronization mechanism between the HDS and HV via an analog PPS pulse. The CE contains an optional data logger and console system (DLC). This system is based on a COTS PC/104+ x86 single board computer connected to a solid-state drive, and communicates with the CE via ethernet.

The 3-D imaging flash LIDAR built by LaRC is a second generation, compact, real-time, air-cooled, 20 Hz time-of-flight sensor system based on 3-D imaging camera technology developed by Advanced Scientific Concepts (ASC).[10] The FLSH integrates a 1.064 μm class IV laser, 1° field of view (FOV) f/7.3 receiver optics, 1° divergence transmitter optics, and a 128 pixel × 128 pixel focal plane array (FPA) which resides in the ASC camera. The flash LIDAR receiver optics are bore-sighted with a small optical witness camera mounted to the FLSH belly which provides a larger (8° along the diagonal) FOV context for the LIDAR image. The separate LEB houses the laser electronics unit, the PC-104-based sensor controller, the power conditioning and distribution subsystem, as well as signal conditioning and thermal control circuitry. To provide for terrestrial operations in humid environments, the flash LIDAR has provisions for dry-air purge during operations. The flash LIDAR achieves a range precision of 8 cm (1-σ), an absolute range accuracy of better than 35 cm over the full operational range of 50 m to over 750 m, and provides the necessary 10 cm GSD at a 750 m range to allow reliable detection of 40 cm hazards. Several real-time signal processing steps

are applied to the raw images, including application of the precision calibration, application of the range accuracy calibration, masking of dead or intermittent pixels, application of a five-point median filter, and insertion of a header which contains various house-keeping data parameters.

III. ALHAT/Morpheus Free Flight Campaign

A. Concept of Operations

The FF10-FF14 flight profile was a 96 s–98 s trajectory from launch pad 5a to concrete pad 1 in the lower right corner of the hazard field (see Fig. 5), about 406 m downtrack. As shown in Fig. 2, after a straight ascent to 245 m, the vehicle started to descend on a roughly 30° downward slope towards the landing site. Shortly after apogee, the pitch rate had sufficiently dampened to start a 6 s mosaic scan of a 60 m × 60 m area centered on the intended landing point (ILP), at a slant range of about 460 m. The resulting LIDAR images were then assembled into a 3-D DEM and analyzed for hazards. Within ∼14 s of the start of the mosaic, the HDS computed and transmitted a list of five safe sites ranked by safety probability to the onboard GN&C system, which then diverted the vehicle landing location to the safest recommended landing location. To be accepted by the onboard vehicle safety checks, the safe site location had to be within a conservative 2 m circle of the surveyed landing pad center, otherwise the vehicle would reject the HDS solution and land on the pad center. The HDS then pointed the gimbal to track a high-relief feature in the mapped area. A correlation-based matching process determined the location of the feature in the LIDAR FOV and, if successful, provided this information as an HRN measurement to the vehicle GN&C system at 1 Hz. At 150 m slant range, the HDS was commanded to shut down, which stowed the gimbal and flash LIDAR in a safe configuration for landing and powered down the system.

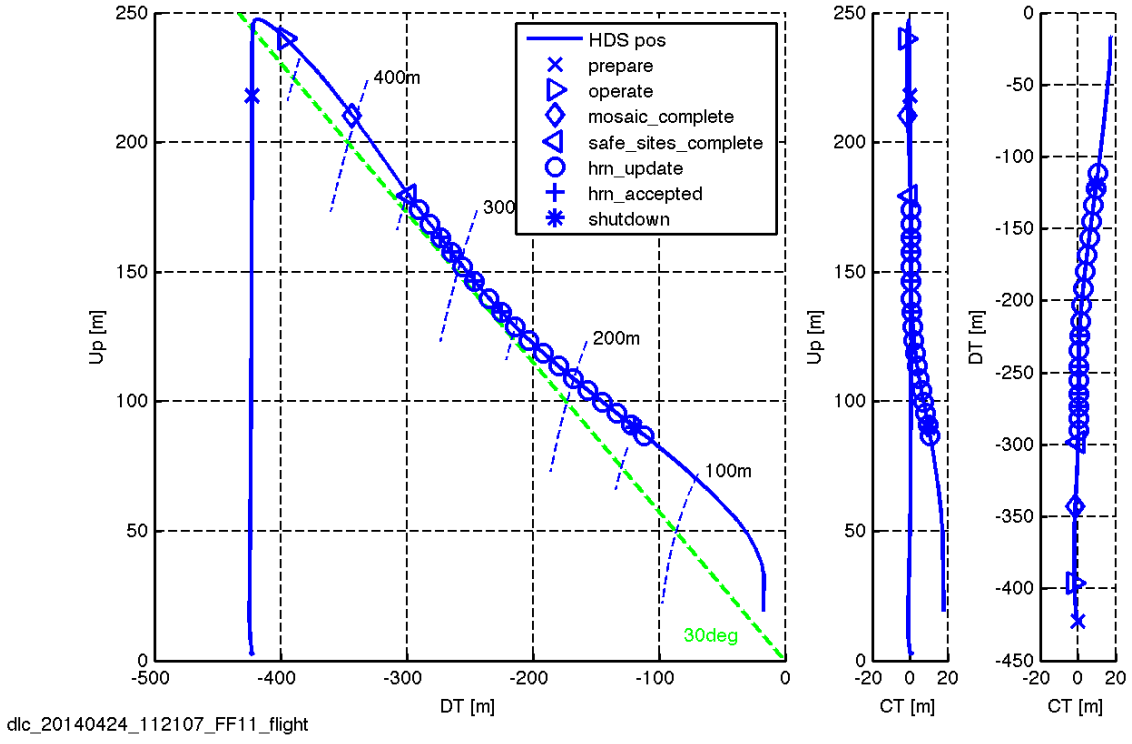

Figure 2. Trajectory of FF11 with HDS-related events.

Note that the original concept of operations was to scan a 90 m × 90 m DEM from 750 m slant range. However, the Morpheus engine performance and fuel margins for this flight campaign only allowed a maximum slant range of 460 m. The DEM size was adapted to the achievable trajectory geometry in a way that the number of LIDAR frames needed to cover the map area stayed constant. The number of frames

to assemble during DEM generation is the dominant factor driving the algorithm processing requirements. The HDS time performance demonstrated in this flight campaign implies the HDS will be able to cover the full 90 m × 90 m DEM from 750 m slant range within the timing requirements once the more capable HD5 engine (350 lbs more thrust, ∼3 % higher I_{sp}) has been fully flight certified.

The Morpheus GN&C system was running a dual string navigation filter, one computing a GPS/INS solution, the other fusing IMU data with the ALHAT sensor measurements. During open-loop flights, the vehicle was navigated using the GPS/INS solution and landed at the pre-determined pad center, with the ALHAT filter and HDS running in parallel for data collection. A so-called 'advanced' open-loop flight, FF12, used the GPS/INS filter for navigation but listened to the HDS safe site recommendation. Closed-loop flights used the ALHAT filter as the primary navigation solution. During the closed-loop flights, the discrepancy between both filter solutions was continuously monitored for consistency. If a safety threshold was exceeded, the filter would switch to the INS/GPS system in flight to safely land the vehicle. This switch happened in both closed-loop flights.

A summary of the overall flight performance can be found in Table 1.

Table 1. Flight summary

	FF10	FF11	FF12	FF13	FF14
Date	04/02/2014	04/24/2014	04/30/2014	05/22/2014	05/28/2014
Type	Open-loop Day	Open-loop Day	Adv. open-loop Day	Closed-loop Day	Closed-loop Night
DEM quality	good	good	moderate	good	moderate
HD Performance	very good	very good	very good	very good	very good
HD Timing [s]	12	12	13	11	11
Top Safe Site Safety Probability [%]	97	98	96	97	95
Distance to pad center [m]	1.1	0.8	1.4	0.4	2.6
HRN Performance	limited	limited	limited	limited	limited
HRN attempts/feasible/valid	16/9/0	18/15/4	17/13/1	18/14/2	17/13/2
Comment	Spurious pixels in flash LIDAR images, results are for post-processed data	best overall DEM quality	dust, pre-triggering during HRN		

IV. Hazard Detection System Performance

The HDS performed mostly as expected in flight, thanks in large part to the extensive preceeding laboratory, ground, and helicopter test campaigns.[11,12] In addition, the team performed a large number of simulations in close cooperation with the Morpheus GN&C team to analyze and predict DEM coverage, safe site selection, number of HRN measurements, and timing performance for the planned Morpheus trajectory and flight dynamics.[13] In what follows we describe the performance of the HDS subsystems in more detail.

A. Flash LIDAR

The maximum slant range at which the flash LIDAR was operated onboard Morpheus at KSC was approximately 470 m.[14] Fig. 3(a) shows a range contour plot for one frame of FF11 data with several rock hazards visible at a slant range of approximately 470 m just before commencement of mosaic operations. The mission elapsed time (MET) is effectively the time since launch. The color gradient from top to bottom in the image is due to apparent slope caused by the viewing angle. Past flight testing at KSC onboard a NASA UH-1H helicopter has shown the flash LIDARs maximum operational range to be approximately 1350 m.[15] The minimum operational range of the flash LIDAR is dominated by saturation and defocus. The flash LIDAR incorporates an Automatic Gain Correction (AGC) algorithm which can hold the image intensity within the sensor's dynamic range down to ranges of 150 m. The flash LIDAR receiver lens is set for maximum depth of field in order to provide in-focus images at all ranges above 250 m. As an example of the hazard detection capability of the LIDAR, Fig. 3(b) shows a DEM which was constructed from a LIDAR range image showing a pair of 30 cm high rock hazards that were visually detected from a slant range of 425 m. For reference,

the pixel size at the range of Fig. 3(b) is 6 cm.

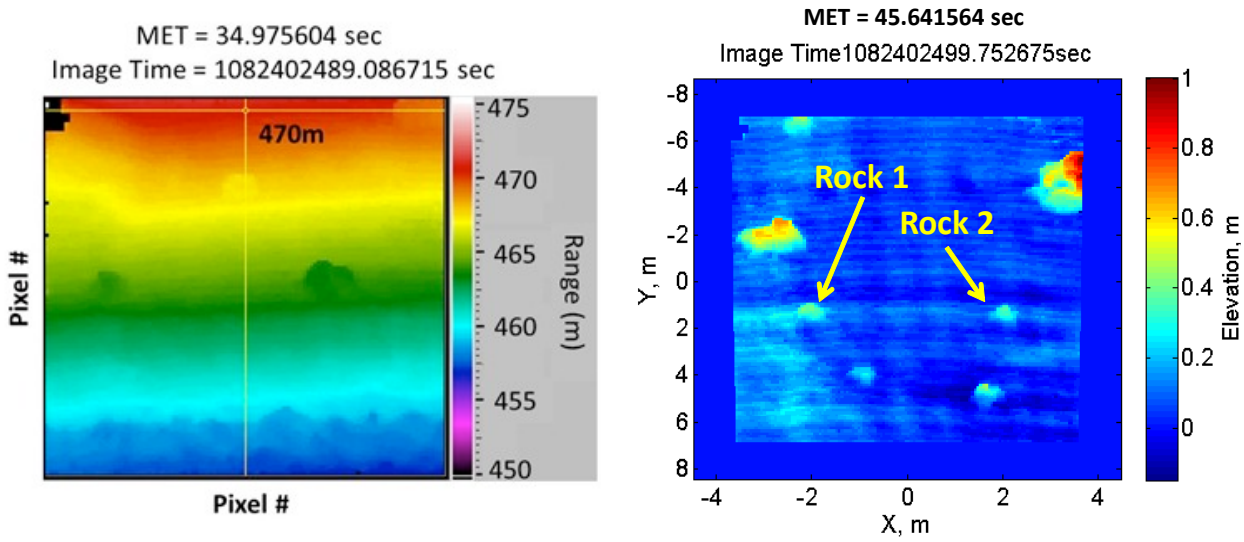

(a) (b)

Figure 3. FF11 range contour image (a) of the KSC hazard field on approach from approximately 470 m. The color gradient from top to bottom in the range image is due to apparent slope caused by the viewing angle. (b) FF11 Digital Elevation Map (DEM) taken from a 425 m slant range of a portion of the hazard field showing a pair of 30 cm hazards that were imaged. Note that the pixel size is 6 cm.

B. Gimbal

The gimbal performance requirements differ during HD and HRN. During the HD mosaic the requirements are driven by the 50% crosstrack and 20% downtrack flash overlap fractions selected to ensure DEM coverage. Given the flash LIDAR's 1° FOV, this translates into at least 0.5° gimbal control error in the azimuth and 0.2° gimbal control error in the elevation axis.[6] For all flights, these requirements were met, with the exception of the azimuth control error briefly exceeding the requirement when switching from row-to-row. This is expected behavior and does not impact the map. During HRN, the gimbal needs to track the feature well enough to ensure it stays within the image search region, resulting in a required maximum control error of 0.25°. This was met easily during FF10–FF13. During the closed loop flights, the control error exceeded the requirements during discrete disturbances, in particular the in-flight filter switch in FF13 and the filter corrections after processing the HRN updates in FF14, but showed similar good steady-state performance as in the open-loop flights.

C. Navigation

The HDS navigation system uses the HDS IMU to provide HDS position, velocity, and attitude estimates of the HDS body frame with respect to an internally computed map frame at 400 Hz (see Figs. 2-4). The HDS navigation system is bootstrapped to the 50 Hz navigation state provided by the host vehicle, except during gimbal mosaics where bootstrapping is suspended to ensure a smooth pose estimate without discontinuities due to filter updates. During the mosaic, the HDS navigation system is propagating the navigation state open loop. Note that no filtering of the HV navigation solution or HDS IMU noise nor estimation of IMU biases is attempted, so that accurate alignment between the two systems, both in space and in time, is of critical importance. In flight, this period of open-loop (internal) propagation lasted ∼6 s. The drift between HV and HDS navigation states is of great interest, because errors in the HDS solution during this time will manifest as distortions in the HazDet DEM. This drift is caused by several parameters, such as errors in the initial state at the start of the mosaic, e.g., due to actual navigation errors or time lag, IMU noise and biases, and misalignments. The maximum position drift was less than ∼0.9 m per axis (mostly downrange). The drift rate for the last two (closed-loop) flights was distinctly smaller than for the first three flights. We

estimate that the drift rate during the closed-loop flights approaches the limit of what is achievable without explicit in-flight bias or misalignment estimation.

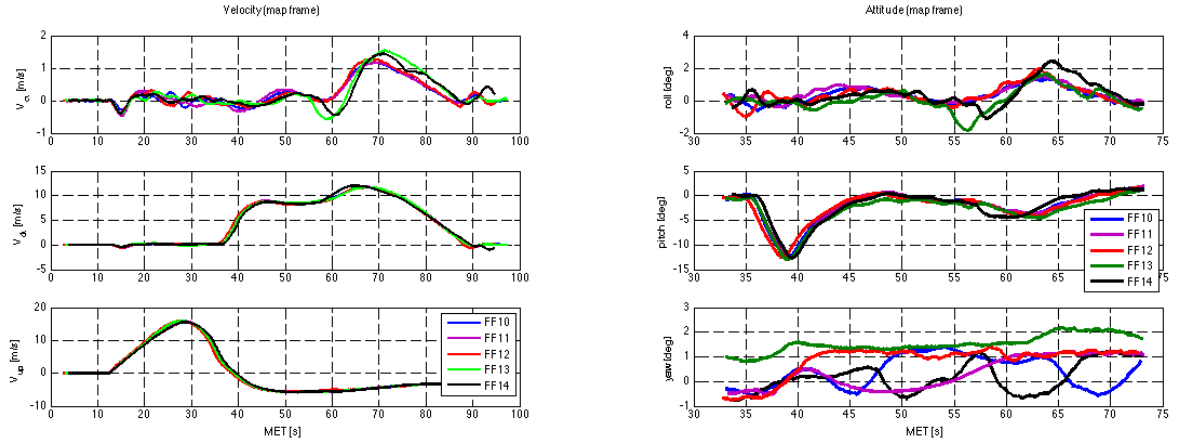

(a) (b)

Figure 4. Velocity and attitude profiles for FF10-FF14.

D. DEM Generation

The ability to reconstruct the scanned terrain surface into a high fidelity DEM is critical to the subsequent processes that characterize that surface in terms of features hazardous to the lander. The features in the terrain become hazardous, in turn, as a function of the lander mechanical characteristics and the dynamics of the touchdown event. DEM quality thus primarily refers to the fidelity with which the reconstructed terrain models the actual terrain. Several basic aspects need to be considered and evaluated. These include imaging the intended terrain patch, placing the DEM correctly in 3-D space, understanding and dealing with the uncertainties in range measurements, mosaicking with minimal misleading artifacts, prevention of gaps, other sources of noise, navigation dynamics, and so on. In order to evaluate the quality of the DEMs, a survey of the hazard field was conducted prior to the helicopter test and data collection campaign early in 2012.[7] The resulting reference DEM is shown in Fig. 5. The hazard field has been adjusted and modified during the Morpheus/ALHAT free flight test campaign to ensure pad 1 was the safest site (dictated by vehicle safety and post-flight accessibility concerns) and thus, the flight LIDAR DEMs are not directly comparable in their entirety to the surveyed data.

Table 2 summarizes basic metrics that can be computed when comparing the flight DEMs to the reference DEM. One aspect of LIDAR DEM quality is the elevation map density. To maximize that, the mosaic plan implements a 50% overlap crossrange and 20% overlap downrange. At 0.1 m resolution, the LIDAR DEMs achieve greater than 99% density. The small gaps in the LIDAR DEMs are expected from rock and crater shadows due to the nominal 30° trajectory path angle at the time of the mosaicking process. The choice of 50-20 overlap fractions is due to real-time considerations, where a 50-50 overlap fractions would result in more evenly distributed noise over the LIDAR DEM. Another aspect of DEM quality refers to the uncertainty in the LIDAR DEM elevations. This uncertainty is propagated from the LIDAR sensor range precision capability. From LIDAR sensor calibration efforts, the expected range precision is 0.08 m for a front-to-parallel surface. The noise component for the near flat and slightly inclined underlying surface of the hazard field terrain observed with a 30° path angle is half of that, or 0.04 m. Table 2 indicates that this expectation was met during FF10, FF11, and FF13. To measure these, we considered the area covered by the concrete pad installed in safe zone 1. The true pad surface is known to be flat and level and consequently, the standard deviation of elevations there is a reasonable metric for the elevation uncertainty associated with the flight DEMs. Fig. 6 illustrates details of the pad surface, a close-up useful to appreciate the quality of the LIDAR DEMs. The mean of these surface pixels is compared, also in Table 2, to the ground truth elevation of 1.3 m. The last aspect of quality is reflected by the measured flatness (planarity) and levelness (orientation) of the pad surface in the LIDAR DEMs. A flat level plane has zero planarity and zero levelness. Planarity refers to deviation from a plane, and levelness refers to angular deviation from horizontal. Given the gridded elevations

(a) (b)

Figure 5. 100 m × 100 m hazard field (a) and reference DEM (b), with the 60 m × 60 m submap imaged in flight shown in the bottom right. The labeled five 10 m diameter safe zones are hazard-free. Heat-resistant concrete landing pads were built on zone 1 and zone 2. Some modifications to the hazard field were carried out prior and after FF10 to ensure pad 1 was the safest landing site.

of the pad surface, the cross-covariance matrix represents planarity (ratio of smallest to largest eigenvalues) and levelness (orientation of the smallest eigenvector.) These metrics, shown in Table 2, illustrate better flatness and levelness for the pad surfaces in FF10, FF11 and FF13, a result that is consistent with the elevation uncertainty metric described above. Taking all these basic factors into consideration, there were no instances of poorly reconstructed terrain during the ALHAT test flight campaign. Subsequent surface characterization and hazard assessment tasks based on these LIDAR DEMs allows the system to successfully select truly safe landing sites. The quality of the reconstruction however, does propagate to the assessment of landing risk.

An example of a complete DEM is shown in Fig. 7, with the five safest sites identified in flight marked by dotted circles. Overall, the general topographic characteristics of the hazard field terrain are well represented in the LIDAR DEMs. Note the LIDAR DEMs are gently sloped per the actual slope of the hazard field, with the concrete pad region in the lower left being higher in elevation than the rest of the DEMs underlying surface. Horizontal distortion maps between the ground reference and the inflight DEMs show a map tie error (mid-field) of between 0.9 m and 3.2 m, mostly in downtrack direction. There is a left-to-right increase in downtrack error, which is related to an interplay between the single flash DEM alignment algorithm and an underestimation of the ground plane slope, which may be caused by an underlying attitude estimation or timing error and is still being investigated.

Table 2. DEM Quality Metrics

	FF10	FF11	FF12	FF13	FF14
DEM density [%]	99.7	99.6	99.6	99.6	99.7
DEM Distortion MID-field					
Downtrack [m]	1.4	0.9	2.1	3.2	1.4
Crosstrack [m]	0.1	0.0	0.1	-0.1	0.2
DEM fidelity mean (landing pad vs. ground truth) [m]	1.1	-0.1	1.0	-0.7	0.1
DEM fidelity std dev (landing pad) [cm]	4	4	5	4	6
DEM Quality at landing pad					
flatness	1.1	1.8	2.3	1.5	2.8
levelness	0.04	0.04	0.08	0.03	0.05

E. Hazard Detection

Terrain features become hazards as a function of the lander mechanical design. The ALHAT HDS system therefore implements two major approaches that maximize the accuracy of the hazard assessment process:[5]

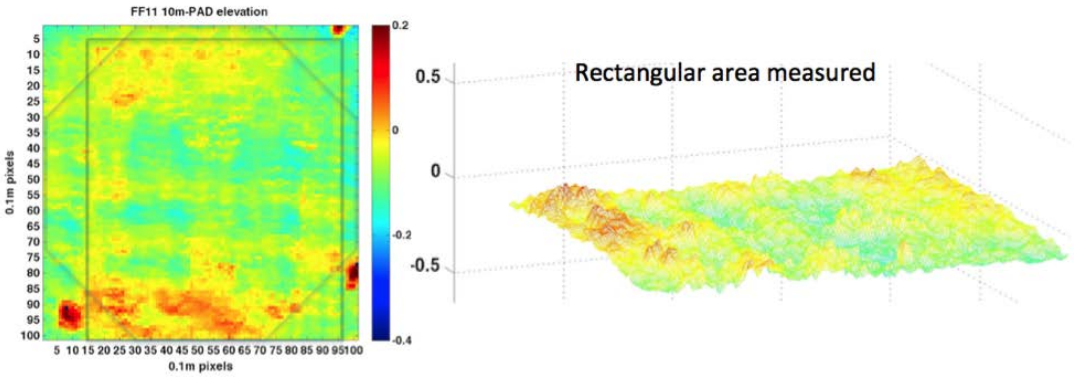

Figure 6. LIDAR DEM portions of the concrete landing pad for FF11, which gives an indication of the noise and roughness in the reconstruction of the otherwise flat and level pad surface.

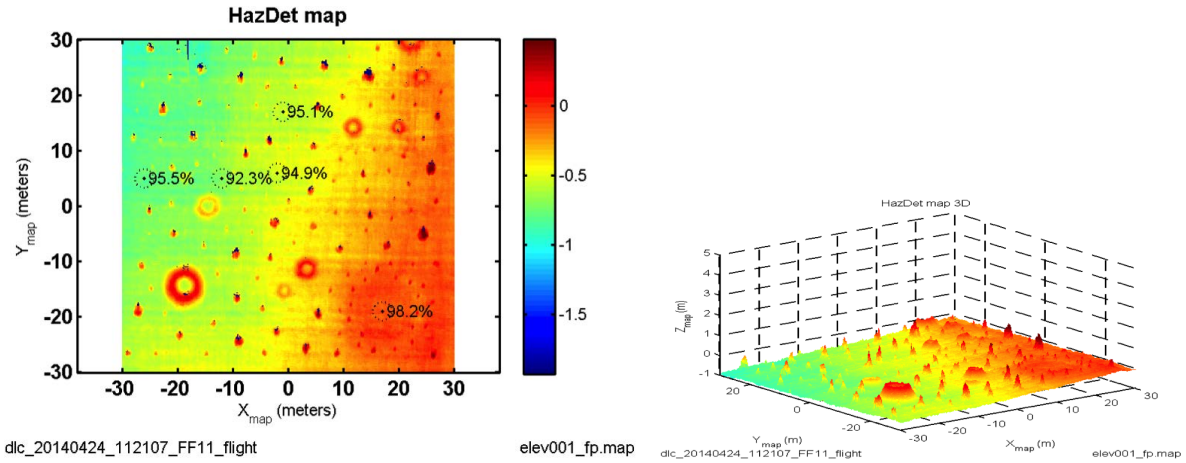

(a) (b)

Figure 7. DEM with identified safe sites for FF11, and associated horizontal distortion map compared to the reference DEM.

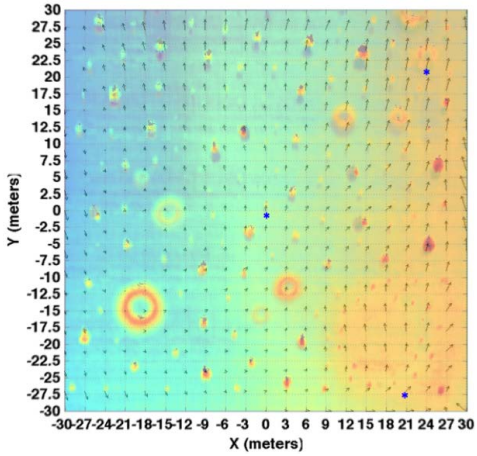

Figure 8. DEM with identified safe sites for FF11, and associated horizontal distortion map compared to the reference DEM.

1) LIDAR DEMs are analyzed at the scale of the lander, i.e., with the highest spatial fidelity. 2) Hazard assessment is based on the mechanical tolerances of the lander, specifically, effective slope and effective roughness. Subsequent text refers to slope meaning effective slope and the same for roughness. Slope is assessed deterministically while roughness is assessed probabilistically. Slope and roughness hazards are combined into a probabilistic hazard map that describes the probability of hazardness, or inversely, the probability of safety.

1. Leg-Pad Maps

The purpose of the leg-pad map (see example in Fig. 9(a)) is to pre-compute the position in space of the lander's footpads, assuming that the touchdown dynamics does not result in crushed rocks or pads sinking into soft soil. It also assumes that the pads are rigid. Pre-computing the location of the pads allows the computation (rather than estimation) of effective slope. Morpheus pads are small (0.25 m diameter). The plane-fit method would give up to 2° slope uncertainty for a lander that size (3.4 m diameter). Because the HDS also computes 3-pixel and 5-pixel median-filtered DEMs (to reduce flash LIDAR artifacts in HRN processing) it was determined that the 3x3 median-filtered DEM was an appropriate candidate for leg-pad maps.

2. Effective Slope Maps

Effective slope is calculated using the leg-pad maps. For all possible placements of the lander pads (the fourth is either on the surface or elevated) the worst slope is taken for each pixel every 1 m. For expediency only the placements deviating 10° from the lander heading were considered rather than the full 45° span. For Morpheus the slope tolerance for hazard assessment is 10°. For context, only rocks taller than 0.6 m would induce slopes greater than 10° and that represents less than 5% risk for the entire DEM.

The 1 m resolution slope map for FF11 is shown in Fig. 9(b). Red pixels mark areas that exceed Morpheus's slope tolerance. As in previous plots, the selected safe sites are also illustrated to indicate that none is located on a slope hazard. Recall that these measurements are made at the scale of the lander. Therefore each pixel (1 m^2) encodes safety for the entire lander (3.4 m diameter.)

3. Effective Full Roughness Maps

Roughness estimation is determined probabilistically by assessing the chance that a roughness measure represents a hazardous condition. By examining at each pixel and all the surrounding pixels under the lander (full roughness) the largest deviation from the lander bottom to all possible slope planes for that pixel, the chance that there is a roughness hazard under the lander is assessed in conjunction with the associated elevation uncertainties of those pixels. As with slope hazard assessment, the estimates are computed 1 m apart but based on the full resolution DEM data. The roughness maps presented below, therefore, have a 1 m resolution. For convenience, the roughness maps are represented by inverse roughness (denoting lack of roughness, or safety), and incorporate in the assessment the landers mechanical tolerance to roughness, which is 0.4 m for Morpheus. The roughness hazard map for FF11 is illustrated in Fig. 9(c). The color scale indicates an arbitrary safety probability of 0.7 or better. As before, the safe landing sites are illustrated to indicate that they are located in safe areas.

4. Safety Probability Maps

Roughness probability maps and slope maps are combined into a hazard probability map. To take into account navigation uncertainty (1 m, 1-σ) the hazard probability map is convolved with a Gaussian filter with space constant 1 m. The results are attenuated safety probability maps where the site selection algorithm looks for safe landing sites. See Fig. 9(d) for the example safety map for FF11. As before, the top five safe sites are illustrated to indicate the safety of the sites in green areas and away from hazards.

5. Safe Site Selection, Evaluation and Verification

The safe site selection algorithm applies a morphological dome operator to detect all the regional maxima in the surface represented by the safety map. The peaks in the regional maxima are true unambiguous peaks, i.e., all the pixel neighbors have values smaller than the peak value. To prevent multiple peaks from

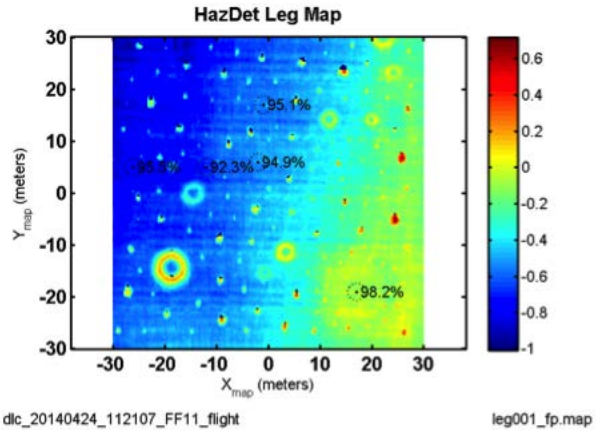

(a)

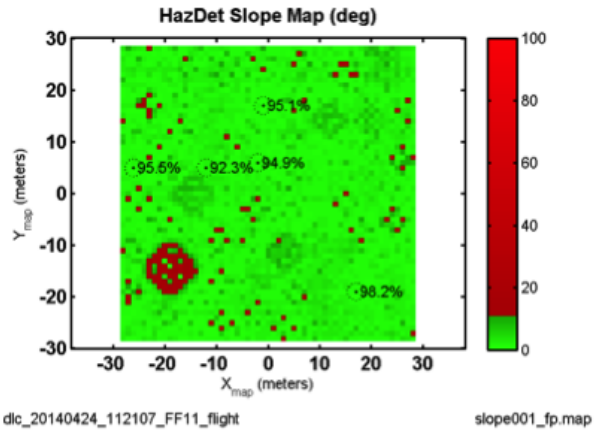

(b)

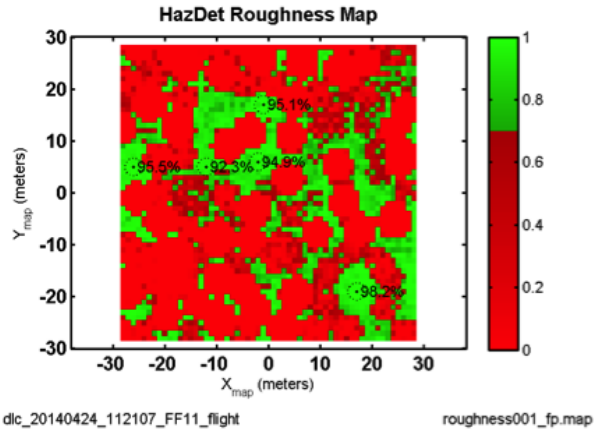

(c)

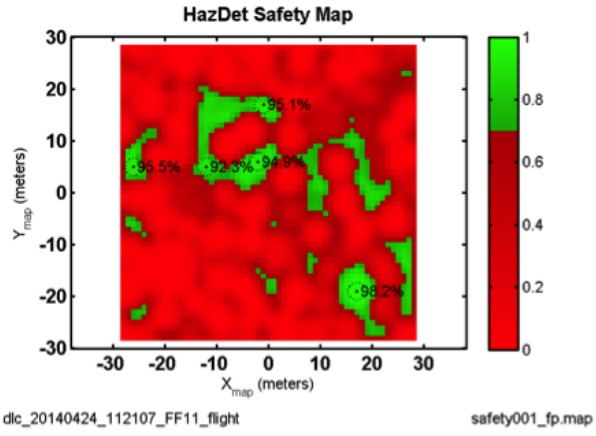

(d)

Figure 9. Hazard detection maps for FF11.

clustering in nearby locations, the algorithm enforces a parameterized minimum separation among the peak locations. The process assembles a probability-ranked list of peaks and reports the top five peak locations to the lander. Note that the hazard detection flight code did not change from flight to flight, so differences in safe site location and safety probabilities are due to variances in the quality of the LIDAR DEMs. FF10 (with the spurious LIDAR pixels removed in post processing) and FF11 have the highest quality DEMs, and the safety probabilities for the top sites are 0.975 and 0.984, respectively. The HDS has no prior knowledge of the topography of the terrain, as indicated earlier, but the flat and level concrete pad would topographically represent the safest zone in the LIDAR DEMs. As a consequence, the safest site would be expected to be found on the surface of the concrete pad, which happened consistently in all flights.

In order to validate the safety of all selected landing sites for FF10-FF14, we inspect the ground truth safety map derived from the noiseless reference DEM. Fig. 10 shows the top five sites for all flights overlaid on the ground truth DEM and 0.1 m safety map. Safety zones 1 (the landing pad), 3 and 4 are shown by 10 m diameter outlines. Note that hazardous rocks have been added in zone 3 that are not reflected in the ground truth maps. Also note that the LIDAR DEMs for FF12 and FF13 extend closer than the edge of the ground truth safety map. The ground truth safety map is produced by the actual HDS algorithms by processing simulated LIDAR flashes rendered from the noiseless DEM. The full-roughness algorithm thus extends the extent of the hazards by at least one-lander radius. Finally the safety map is convolved with the expected 1 m (1-σ) navigation uncertainty. The green areas therefore have probability 1.0. Note that the safe sites are on the green area, with the exception of the two sites in the near-field. These are on or beyond the edge of the ground truth safety map and are therefore assigned a truth zero safety probability, even though in reality they lie on flat, grassy terrain which led (correctly) to a very high safety probability in the in-flight computations. Site 5 in FF12, on the right of zone 1, has safety probability 0.96 in the ground truth and 0.92 in the LIDAR safety map.

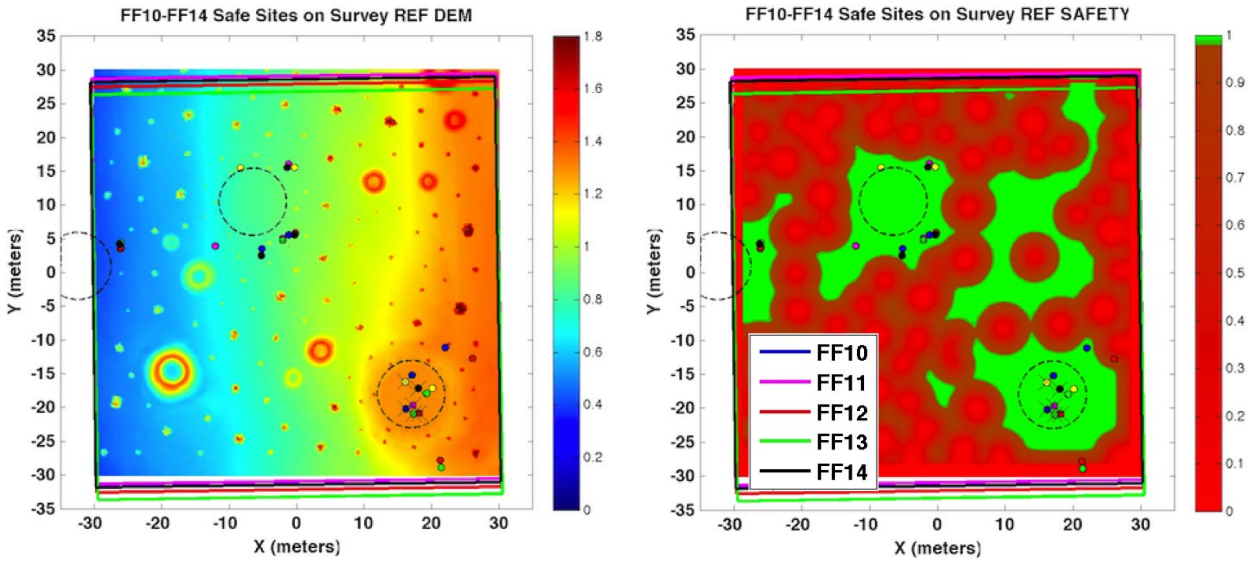

(a) (b)

Figure 10. Safe sites for all flights, overlaid on ground truth terrain and safety maps. Yellow markers indicate the FF10 flight results, with the spurious LIDAR range data present.

F. Hazard-Relative Navigation

HRN selects and tracks a relevant fixed feature (e.g., a large rock in the LIDAR DEM) nearby, so that during descent, the apparent motion of the fixed feature can be measured. Such motion can be interpreted as a change in the navigation state error since the time the safe site for landing was identified. HRN assumes that the attitude errors are very small and that the apparent motion of the tracked feature is largely due to errors in the position of the lander. In flight, each incoming LIDAR image is converted to a single flash DEM. A template centered on the dominant terrain feature in this DEM is first flattened and then correlated with the HD DEM in a search region around the predicted location. The location of the best

match is computed with subpixel accuracy using biquadratic interpolation of the correlation surface. Each match is then evaluated against quality metrics on the correlation scores, the measurement covariance, and the correlation peak quality (sharpness) before being passed to the host vehicle. The HRN measurements the host vehicle receives include the map location of the surface feature being tracked, the location of the surface feature relative to the HDS, in a sensor-fixed reference frame, and the uncertainty of the HDS-relative feature position. Note that in flight the decreasing slant range to the landing site results in an increasing mismatch between map and image resolution, or, equivalently, a smaller and smaller LIDAR footprint in the map. This, together with the limited depth of field of the LIDAR eventually precludes successful HRN measurements as the vehicle approaches the landing site. Due to the short trajectories, no re-mosaicing steps to create finer resolution maps in flight was executed.

Table 1 summarizes the HRN performance for FF10-FF14. The number of measurement attempts denotes the number of 1 Hz flash LIDAR images considered for HRN. The number of feasible attempts denotes range frames that have sufficient spatial extent for constructing a template for matching. Finally, the number of valid measurements count those attempts where the template was successfully correlated to the HD DEM, the sub-pixel location of the correlation peak could be computed, and the match met all quality criteria.

The ALHAT/Morpheus free flight campaign provided the first opportunity to test the various improvements of the HRN algorithm implemented in the wake of the 2012 helicopter tests under dynamic conditions.[7] Changes included bug fixes to the feature selection algorithm, increased robustness of the correlation algorithm to missing data, improvements to the DEM quality, and overall greatly increased pointing accuracy and stability. As expected, the free flights revealed a number of new issues that were addressed after post-flight analyses and corrected between flights. These included tuning the HRN feature search region, improved numerical conditioning of range image pre-processing, degraded range image quality due to aerosols or erroneous pixels, and the interaction of HRN measurements with the ALHAT navigation filter.

Since over the course of the campaign several important changes had been implemented in the HRN algorithm, all flights were post-processed with the HRN software flown in FF14, which more than doubled the number of valid measurements (see Fig. 11(a)). Visual inspection of all reprojected feature vectors in the LIDAR images demonstrated that the feature was correctly matched (see Fig. 11(b)). Unfortunately, even the latest updates flown in FF14 only yielded two valid HRN measurements. This low number was due to the fact that the feature drifted outside of the search region after HRN measurements were processed by the Morpheus navigation filter. This finding motivated a revision of the HRN measurement simulator and measurement model, and an exhaustive simulation campaign using flight and simulated data. For the upcoming flight campaign, several modifications to LIDAR operations (to mitigate spurious LIDAR pixels), to the flight rules (to avoid wind-blown plume and dust in LIDAR imagery[a]), and to the vehicle navigation filter (to reduce feature drift) have been implemented.

V. Conclusion

The 2014 free flight campaign of the ALHAT sensor suite on Morpheus marked NASA's first flight demonstration of a real-time hazard detection system and demonstrated full operability of the ALHAT sensor suite on a relevant, lunar-like trajectory with realistic flight dynamics and landing terrain. The HDS performed DEM generation, hazard detection, and hazard relative navigation within the timing requirements, and produced five truly safe sites in every flight, the safest of which was always on the landing pad as desired for vehicle safety. Several updates were made to the HRN algorithm during the campaign to address and mitigate anomalies experienced in flight. These updates have been tested in flight and in post-processing. In all, the flight campaign demonstrated for the first time autonomous precision-landing and hazard avoidance abilities on a NASA rocket-powered vehicle, advancing ALHAT technology to TRL 6 and underscoring its readiness for infusion into future flight missions.

Acknowledgments

We would like to acknowledge the hard work and dedication of the ALHAT and Morpheus teams at JPL, LaRC, JSC, and KSC, who have together taken the HDS from a paper design to multiple in-flight demonstrations on a free-flying rocket vehicle. The work described herein was performed at the Jet Propul-

[a]The dust of concern here was generated at launch, not at touchdown, and would not exist in a lunar landing scenario.

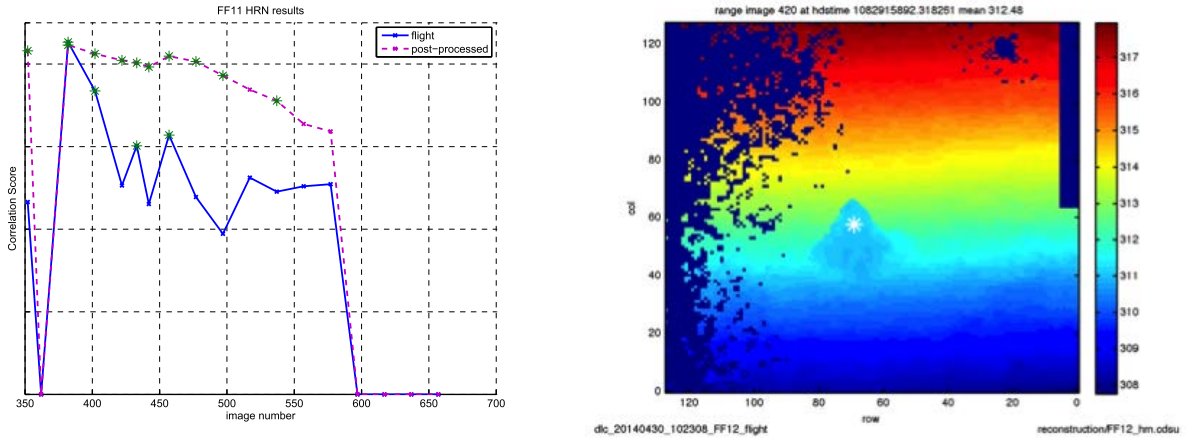

(a) (b)

Figure 11. (a) FF11 correlation scores in flight (solid line) and after post-processing with FF14 HRN algorithm and settings (dashed line). Valid HRN measurements are denoted by a green asterisk. (b) FF12 example of HRN feature (rock) and reprojected match (white asterisk). Notice the effect of pixel masks (top right) and pre-triggering due to dust (left), which in this case the HRN algorithm was able to tolerate.

sion Laboratory and at the Langley Research Center within the Autonomous precision Landing and Hazard detection and Avoidance Technology (ALHAT) project. The work at the Jet Propulsion Laboratory, California Institute of Technology, was performed under contract with the National Aeronautics and Space Administration. Government sponsorship acknowledged

References

[1]Epp, C., Robertson, E., and Brady, T., "Autonomous Landing and Hazard Avoidance Technology (ALHAT)," *Proc. IEEE Aerospace Conference (AEROCONF 2008)*, March 2008.

[2]Keim, J., Mobasser, S., Kuang, D., Cheng, Y., Ivanov, T., Johnson, A., Goldberg, H., Khanoyan, G., and Natzic, D., "Field test implementation to evaluate a flash LIDAR as a primary sensor for safe lunar landing," *IEEE Aerospace Conference (AEROCONF 2010)*, March 2010.

[3]Amzajerdian, F., Pierrottet, D., Petway, L., Hines, G., and Barnes, B., "Doppler lidar sensor for precision navigation in GPS-deprived environment," *Proc. International Society for Optics and Photonics (SPIE)*, June 2013.

[4]Amzajerdian, F., Pierrottet, D. F., Petway, L. B., Hines, G. D., Roback, V. E., and Reisse, R. A., "Lidar Sensors for Autonomous Landing and Hazard Avoidance," *Proc. AIAA SPACE 2013 Conference & Exposition*, San Diego, CA, 10–12 Sept. 2013.

[5]Ivanov, T., Huertas, A., and Carson, J., "Probabilistic Hazard Detection for Autonomous Safe Landing," *Proc. AIAA Guidance, Navigation, and Control Conference*, August 2013.

[6]Carson, J. M., Bailey, E. S., Trawny, N., Johnson, A. E., Roback, V. E., Amzajerdian, F., and Werner, R. A., "Operations Concept, Hardware Implementation and Ground-Test Verification of a Hazard Detection System for Autonomous and Safe Precision Lunar Landing," *Proc. AAS/AIAA Astrodynamics Specialist Conference*, Hilton Head Island, SC, Aug. 11–15 2013.

[7]Trawny, N., Carson, J. M., Huertas, A., Luna, M. E., Roback, V. E., Johnson, A. E., Martin, K. E., and Villalpando, C. Y., "Helicopter Flight Testing of a Real-Time Hazard Detection System for Safe Lunar Landing," *Proc. AIAA SPACE 2013 Conference & Exposition*, San Diego, CA, 10–12 Sept. 2013.

[8]Villalpando, C., Johnson, A., Oberlin, J., and Goldbert, S., "Investigation of the Tilera Processor for Real-time Hazard Detection and Avoidance on the Altair Lunar Lander," *Proc. IEEE Aerospace Conference (AEROCONF 2010)*, March 2010.

[9]Villalpando, C., Werner, R., III, J. C., Khanoyan, G., Stern, R., and Trawny, N., "A Hybrid FPGA/Tilera Compute Element for Autonomous Hazard Detection and Navigation," *Proc. IEEE Aerospace Conference (AEROCONF 2013)*, March 2013.

[10]Stettner, R., Bailey, H., and Richmond, R., "Eye Safe Laser Radar Focal Plane Array for Three-Dimensional Imaging," *Proc. International Society for Optics and Photonics (SPIE)*, 2001.

[11]Carson, J. M., Robertson, E. A., Pierottet, D. F., Roback, V. E., Trawny, N., Devolites, J. L., Hart, J. J., Estes, J. N., and Gaddis, G. S., "Preparation and Integration of ALHAT Precision Landing Technology for Morpheus Flight Testing," *Proc. AIAA Space 2014 Conference*, Aug. 2014.

[12]Carson III, J. M., Hirsh, R. L., Roback, V. E., Villalpando, C. Y., Busa, J. L., Pierrottet, D. F., Trawny, N., Martin, K. E., and Hines, G. D., "Interfacing and Verifying ALHAT Safe Precision Landing Systems with the Morpheus Vehicle," *Proc. AIAA SciTech 2015 Conference & Exposition*, Kissimmee, FL, Jan. 2015.

[13]Luna, M. E., Huertas, A., Trawny, N., Villalpando, C. Y., Martin, K. E., Wilson, W., and Restrepo, C. I., "Simulations of

the Hazard Detection System for Approach Trajectories of the Morpheus Lunar Lander," *Proc. AIAA SciTech 2015 Conference & Exposition*, Kissimmee, FL, Jan. 2015.

[14]Roback, V. E., Pierrottet, D. F., Amzajerdian, F., Barnes, B. W., Bulyshev, A. E., Hines, G. D., Petway, L. B., Brewster, P. F., and Kempton, K. S., "Lidar Sensor Performance in Closed-Loop Flight Testing of the Morpheus Rocket-Propelled Lander to a Lunar-Like Hazard Field," *Proc. AIAA SciTech 2015 Conference & Exposition*, Kissimmee, FL, 5 - 9 Jan. 2015.

[15]Roback, V. E., Bulyshev, A. E., Amzajerdian, F., Brewster, P. F., Barnes, B. W., Kempton, K. S., and Reisse, R. A., "Helicopter Flight Test of a Compact, Real-Time 3-D Flash Lidar for Imaging Hazardous Terrain during Planetary Landing," *Proc. AIAA SPACE 2013 Conference & Exposition*, No. 2013-5383, San Diego, CA, 10–12 Sept. 2013.