Evaluation of the Simple Safe Site Selection (S4) Hazard Detection Algorithm using Helicopter Field Test Data

Michael E. Luna,[*] Eduardo Almeida,[†] Gary Spiers,[‡] Carlos Y. Villalpando,[§] Andrew E. Johnson,[¶] and Nikolas Trawny[‖]

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

Small scale terrain hazards, such as rocks, slopes, and craters, can pose significant risk to landing spacecraft and rover or payload deployment. Onboard Hazard Detection and Avoidance (HDA) systems scan and analyze the landing area for these hazards in real time during descent, and divert the spacecraft to the safest touchdown site. The computationally efficient Simple Safe Site Selection (S4) algorithm combined with a flash LIDAR is an HDA system geared towards small robotic spacecraft. Rather than creating and analyzing a digital elevation map (DEM) from potentially many overlapping range images, S4 operates directly on a single flash LIDAR image. Extending prior work that has analyzed S4 performance for Mars landing using extensive simulations, this paper evaluates S4 performance using actual flash LIDAR images of an artificial hazard field acquired during a 2014 helicopter field test in Death Valley, CA. In particular, we describe LIDAR characterization and calibration, creation of ground truth elevation and safety maps, creation of ground truth sensor poses, actual S4 algorithm processing, and performance analysis. The results show that the safety cost images produced by S4 are remarkably close to the ground truth safety map (computed offline by an HDA algorithm developed under the Autonomous Landing and Hazard Avoidance (ALHAT) project) at significantly reduced computational cost, confirming S4 as a viable candidate algorithm for onboard spacecraft HDA.

I. Introduction

Scientifically interesting sites for planetary surface exploration often contain terrain hazards such as rocks, slopes, and craters that can pose risk to safely landing a spacecraft. These risks can be mitigated through onboard autonomy, e.g., using technologies like terrain-relative navigation (TRN) and hazard detection and avoidance (HDA). TRN matches sensed terrain during descent to an onboard a-priori map created from orbital reconnaissance in order to precisely localize the spacecraft, which allows avoiding hazards detectable from orbit. In contrast, HDA senses the terrain during descent and can therefore operate without precise localization and can sense smaller hazards below the resolution of orbital imagery.

Multiple sensing modalities for active hazard detection have been proposed,[1,2] including passive optical monocular camera-,[3–8] stereo camera-,[9,10] radar-,[11] and LIDAR-based[12–16] approaches. The European Space Agency has investigated onboard HD for a variety of future lander missions[17–20] . NASA's ALHAT project[21–27] demonstrated LIDAR-based HD onboard a helicopter and a terrestrial rocket vehicle. The chinese Chang'e-3 mission was the first to demonstrate onboard HD in an actual mission, by landing safely on the moon in December 2013.[28]

The Simple Safe Site Selection (S4) algorithm[29] was developed for flash LIDAR-based HD on small robotic landers. It detects safe sites within a single flash LIDAR range image, thus avoiding the need to

[*]Guidance & Control Operations Group, AIAA Senior Member.
[†]Maritime & Aerial Perception Systems Group.
[‡]Optics Section Staff.
[§]Advanced Computer Systems & Technologies Group.
[¶]Guidance & Control Section Staff, AIAA Senior Member.
[‖]GN&C Hardware and Testbed Development Group, AIAA Senior Member.

generate a digital elevation map (DEM) and resulting in a very computationally efficient algorithm. Prior work demonstrated the effectiveness of S4 on simulated Mars terrain and approach geometries. This paper evaluates S4 performance on data acquired during a helicopter field test against a ground truth safety map generated using the ALHAT HDA algorithm.[30] The paper is organized as follows: the field test is described in Section II, followed by a discussion of ground truth elevation and safety map generation in Section III, ground truth sensor pose computation in Section IV, and LIDAR characterization in Section V. Section VI summarizes the S4 algorithm. Finally, Section VII presents results of S4 processing of some field test LIDAR images.

II. Field Test Description

In February 2014, the Jet Propulsion Laboratory (JPL) conducted a helicopter flight test of safe and precise landing technology in Death Valley, CA.[31] A sensor suite, comprising among others a commercial ASC TigerEye Flash LIDAR sensor, with a 128 pixel × 128 pixel detector array and 3 deg FOV optics, operating at 10 Hz, was mounted to a two-axis gimbal. This sensor platform, combined with an Applanix GPS/INS and high resolution gimbal encoders for navigation ground truth was mounted to the nose of a Eurocopter AS350 AStar helicopter (see Fig. 1). Field test avionics and support equipment in the helicopter cabin provided power, a graphical user interface for the operator, as well as real time logging of ground truth and sensor data. Prior to the flight tests, a metrology campaign determined geometric precision alignment between all sensors and the gimbal. Unfortunately, at the time of the field tests the flash LIDAR – being only a secondary test objective – was not yet fully integrated with the remaining sensor suite, resulting in a lack of time-synchronization to the rest of the sensor payload, in turn creating a significant challenge to establish ground truth LIDAR poses.

Painted plastic hemispheres of various known sizes were laid out near the Furnace Creek Airport to form an artificial hazard field (see Fig. 1). Over the course of the test campaign, the team acquired several LIDAR data sets from the helicopter descending over the tarmac (later used for fixed pattern noise calibration), and from it hovering over the hazard field. The small LIDAR FOV, coupled with no real-time visual feedback from a context camera, made systematic imaging of the hazard field surprisingly difficult, and resulted in some of the data sets containing only few images of actual hazards. To increase the chance of imaging hazards, for some flights the gimbal was set to a sinusoidal sweeping pattern.

(a) (b)

Figure 1. (a) Gimbaled sensor head mounted to the nose of a Eurocopter AS350 AStar helicopter. (b) Helicopter flying over the artificial hazard field in Death Valley, CA.

III. Ground Truth Digital Elevation and Safety Map Generation

The synthetic hazard field was comprised of 52 plastic hemispherical hazards, varying in diameter from 0.10 m to 1.22 m, painted light grey, dark grey, black, and white, and randomly distributed over an approximately 20 m × 20 m area. The Northwest, Southwest, and Southeast corners of the hazard field were surveyed using long term dwell GPS measurements. In addition, the top centers of all hazards, as well as the three corners, were surveyed using a total station. The rigid body transform between the GPS frame

(WGS84) and the total station frame was estimated using the three corner point correspondences. This allowed all hazards to be mapped in the same frame as the ground truth helicopter trajectory. A synthetic DEM was then created by adding hemispheres of the right diameter at the measured East/North locations to an approximate, level ground plane around the hazard field center (see Fig. 2). This DEM can be sampled at any resolution and was used to generate the truth safety map.

The DEM was then processed using a high fidelity probabilistic hazard detection algorithm[30] to generate the ground truth safety map. In intermediate steps, the algorithm creates slope and roughness hazard maps depending on the hazard tolerance of the vehicle. For the assumed vehicle roughness tolerance of 35 cm, only the taller hemispheres represent actual landing hazards. These two maps are combined into a hazard probability map. To take into account navigation uncertainty (1m, 1σ) the hazard probability map is convolved with a Gaussian filter with space constant 1m, resulting in the final safety map (see Fig. 3).

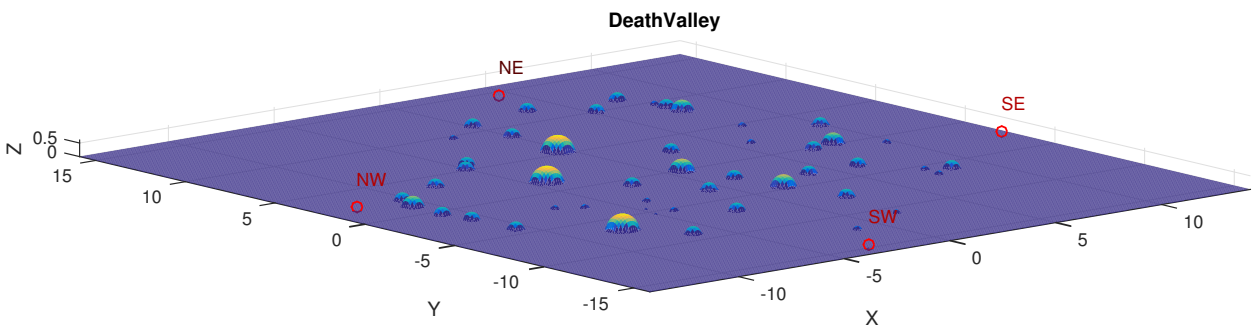

Figure 2. Synthetic ground truth DEM generated from the surveyed hazards.

IV. Ground Truth Sensor Pose Determination

To evaluate the S4 HDA performance on the LIDAR images, it is necessary to determine the true safety of the recommended safe site in the ground truth safety map. This in turn requires precise knowledge of the LIDAR sensor pose at the time the image was acquired. Due to the lack of time-synchronization with GPS ground truth at flight time, the LIDAR data had to be precision aligned in post processing. To this end, we determined a coarse time alignment per flight using cross correlation of the measured and expected boresight range to the ground plane (see Fig. 4). A residual meter-level position error per image, possibly due to a calibration error, was removed using an optimal translation shift estimated from feature matching of a single hazard between range image and map (see Fig. 5). Since this process has not yet been automated, only a select set of five images has been processed for this paper.

V. LIDAR characterization

The flash LIDAR images suffered from invalid pixels and fixed pattern noise which required calibration prior to S4 processing.

A. LIDAR pixel mask creation

A pixel mask was created by identifying stuck, biased, or excessively noisy pixels over several thousand range images. Also included was the Automatic Range Correction (ARC) pixel in the top right corner of the image used to improve the LIDAR zero time reference and to calibrate the readout integrated circuit timing clock, resulting in improved LIDAR range accuracy. To identify pixels to be masked, we used a dataset containing more than 4000 frames, acquired at near constant altitude. Hot or excessively noisy pixels could be identified from the per pixel standard deviations, essentially high-pass filtering the images. Conversely, a biased pixel could be identified by considering the per pixel mean, corresponding to a low-pass filtered image sequence. In order to reliably reject the ARC pixel, we chose to add a manual rectangular mask in the top right corner. The final mask obtained by combining all three criteria contained 98 pixels (including 27 singletons) or 0.6% of the range image.

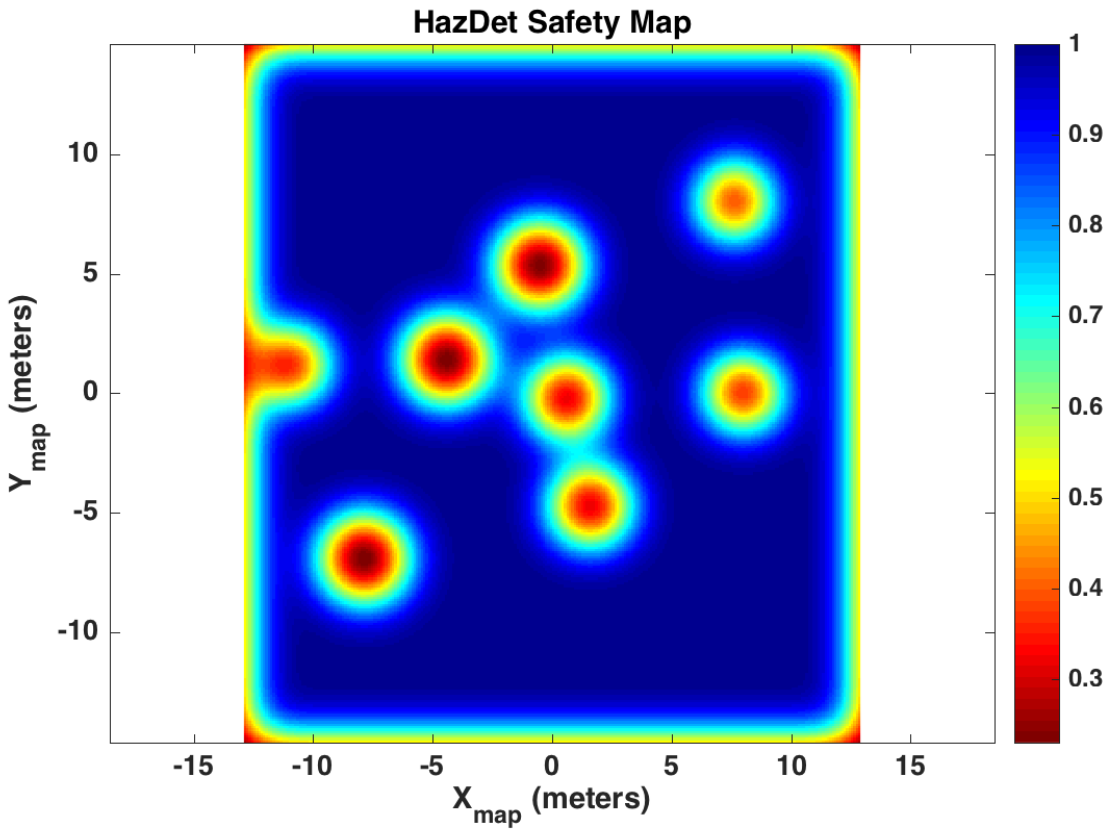

Figure 3. Ground truth safety map for the Death Valley hazard field.

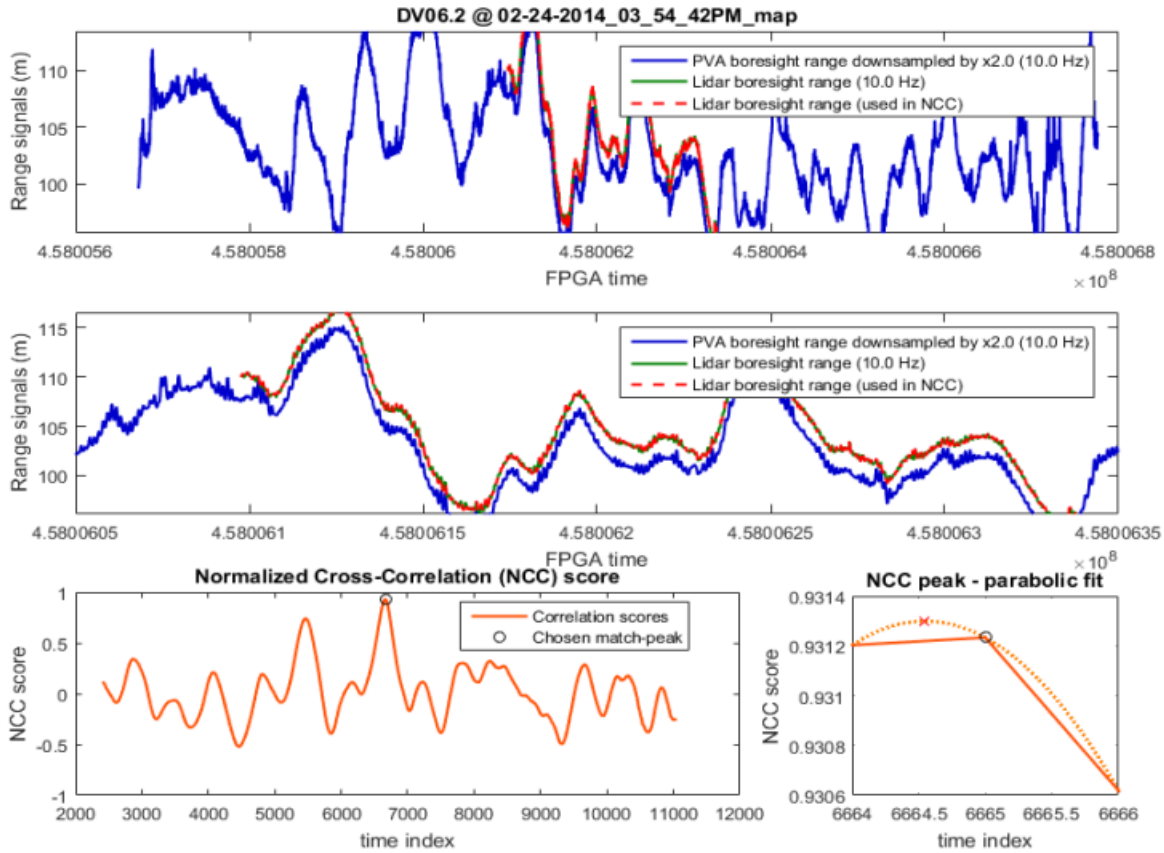

Figure 4. Example of coarse time alignment using normalized cross correlation of the boresight range. A post-fit overlay of measured vs expected boresight range (top) shows good agreement, but a zoom-in (middle) reveals a small residual bias. Quadratic interpolation of the correlation score (bottom) yields the optimal start time.

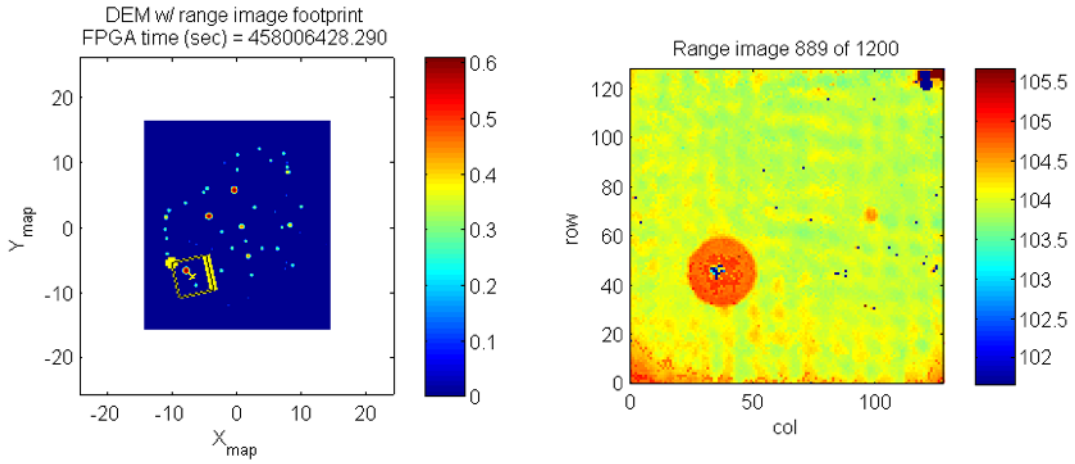

(a)

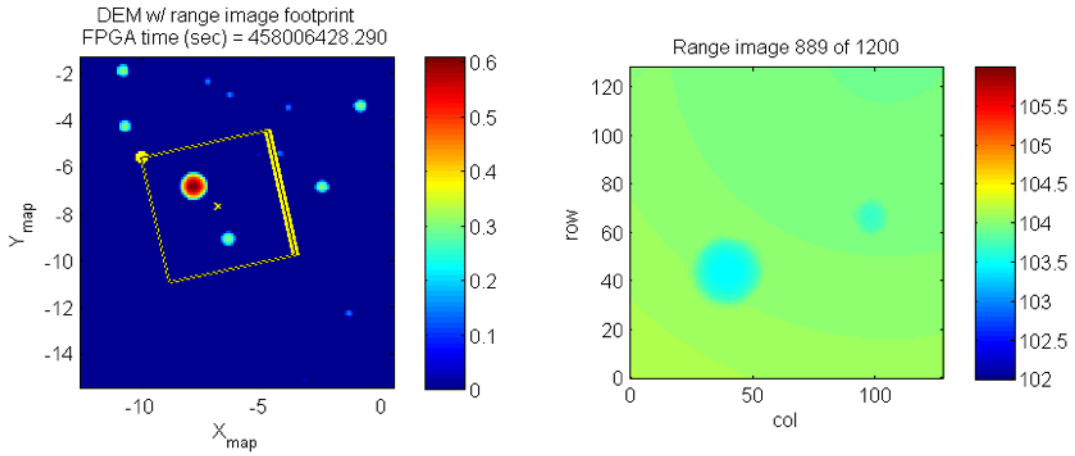

(b)

Figure 5. Example of fine alignment using feature matching between range image and DEM. (a) Raw range image (right) and footprint in DEM (left). (b) Zoomed-in footprint after fine alignment and synthetic range image generated from aligned pose show excellent agreement with the sensor measurements.

B. Fixed Pattern Noise Calibration

No a-priori or post flight LIDAR calibration data was collected against reference surfaces. Calibration of the fixed pattern noise was therefore conducted using data from the DV03.2 descent over the flat Death Valley airport tarmac. 200 frames covering ranges from the surface of 105 m – 58 m were used for the calibration. No camera images were available for this flight to confirm a lack of features in every frame and so the statistics of the standard deviation of the range in each image were used to select the images to use for fixed pattern extraction. Initial processing eliminated stuck pixels and those pixels with ranges > 110 m and < 50 m associated with range bias and noise associated range artifacts. A plane was then fitted to each frame using an M-estimator Sample Consensus (MSAC) algorithm.[32] It should be noted that all of the fitted target planes were significantly less than 30 deg from perpendicular and so errors in range at the edge of the field are less than 0.02 m at 100m range. The fitted plane was then subtracted from its associated range frame, thereby correcting each individual frame for any tilt of the LIDAR with respect to the ground surface and the corrected planes averaged to obtain the mean image. This gave a range map of the fixed pattern noise averaged over the frames (see Fig. 6), which was subsequently subtracted from all LIDAR images in a preprocessing step.

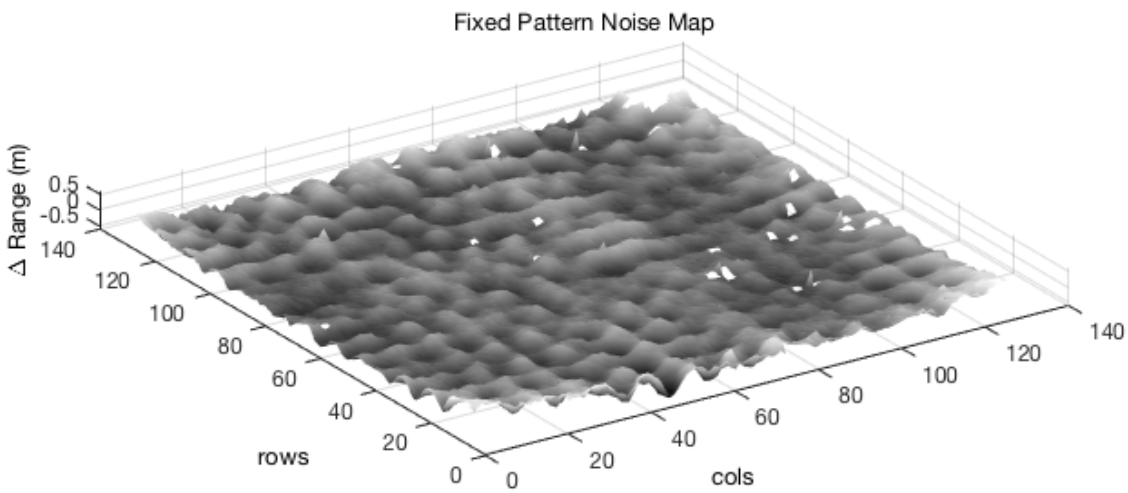

Figure 6. The average fixed pattern map obtained after subtracting the ground plane from 200 descent images over flat terrain.

VI. S4 Algorithm

Traditional hazard detection algorithms first create a digital elevation model (DEM) of the landing site, determine slope and roughness of the terrain, and then evaluate slope and roughness against the hazard tolerance of the vehicle to compute a terrain safety map.[30] The objective of the S4 algorithm is to significantly reduce the computational complexity associated with DEM creation and terrain evaluation by operating directly on a single flash LIDAR range image without creating a DEM.[29] This enables rapid execution within the limited computational resources and short descent timelines typical of robotic landers. Note that the S4 algorithm only evaluates landing safety based on vehicle dimensions, and does not consider sensitivity to payload deployment or access to science targets. S4 data processing consists of 8 principal steps (see Fig. 8):

1. Range image acquisition Flash LIDAR image centered on the desired landing site
2. Range image flattening and attitude correction Accounts for spherical-to-Cartesian projection and off-nadir viewing angle
3. Median filtering Noise reduction and outlier removal
4. Site Delta Range Image (SDR) – Slope over Lander Footprint Local approximate slope evaluation from maximum elevation difference over lander footprint (see Fig. 7)

5. Rock Delta Range Image (RDR) – Roughness over Range Image Approximate roughness evaluation from maximum elevation difference over expected rock footprint (see Fig. 7)

6. Site Rock Delta Range Image (SRDR) – Roughness over Lander Footprint Local roughness evaluation from maximum RDR over lander footprint.

7. Cost Image Creation and Navigation Uncertainty Smoothing Cost as thresholded sum of slope and roughness ratios Cost(i,j) = SRDR(i,j)/Rock Height Threshold + max(0, SDR(i,j)/Slope Threshold if >1)

8. Safe Site Selection Select pixel with lowest cost and convert to vector in vehicle frame.

Processing the flash LIDAR range images from the field test resulted in several bug fixes and small changes to the S4 algorithm, for example adding a pre-processing step to apply the fixed pattern noise correction, and modifications to properly account for invalid pixel data.

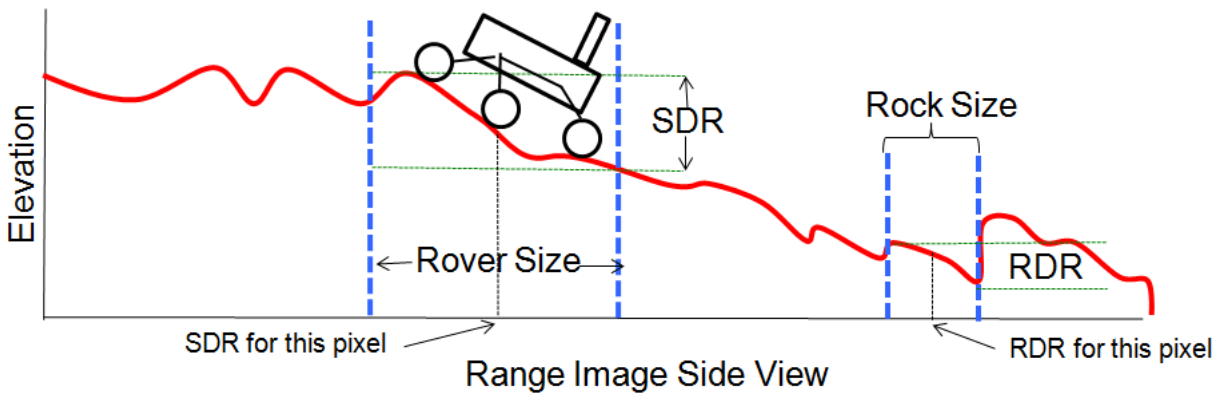

Figure 7. Site Delta Range image (SDR) and Rock Delta Range image (RDR) definition.[29]

VII. Results

Five selected LIDAR range images from the Death Valley flight DV06.2 were processed through S4. The S4 parameters shown in Table 1 were chosen to reflect the sensor FOV and a small generic robotic lander with typical hazard tolerances. Notice that the LIDAR FOV of 3° resulted in a 5.2 m × 5.2 m image footprint for a nadir pointing sensor at 100 m altitude. While achieving a ground sample distance of 4 cm/pixel, the overall footprint size is quite small for a lander of this size, resulting in few candidate safe landing sites per image that fit the entire vehicle, and few hazards imaged at the same time. This combination would not have been the preferred sensing geometry for an actual lander mission. For these flight tests, the altitude and sensor FOV were a compromise between maximum sensor range and cost of readily available optics. Intermediate results for range image 889 are shown in Figs. 9 and 10. Comparing the range image in Fig. 9 to the footprint in the DEM computed using the precision-aligned true sensor pose visually confirms good correspondence of the imaged and mapped hazards. The S4 intermediate products in Fig. 10 clearly show the contributions of the two hemispheres to roughness and slope hazards. RDR, SDR, and SRDR also demonstrate the rectangular keep-out zones at the image borders of one rock hazard radius and one lander radius, respectively. The final S4 cost image shown in Fig. 11 and topologically agrees very well with the ground truth safety map. Notice that the keep-out borders in SDR and SRDR push the cost minimum towards the image center. In addition, the masked arc pixel in the top right image corner is creating a local safety cost maximum. For these reasons, the location of the recommended safe site is not exactly at the location of maximum true safety probability, which is at the top right corner, at the maximum distance from the hazard. Still, the recommended site has an acceptable true safety probability of 0.833. The results for the remaining four images are shown in Fig. 12 and summarized in Table 2. For all four images, S4 selected a safe site in the same region as the ground truth safe site. The topology of cost image of range image 3782 differs somewhat from the ground truth, since the hazard on the left side of the image

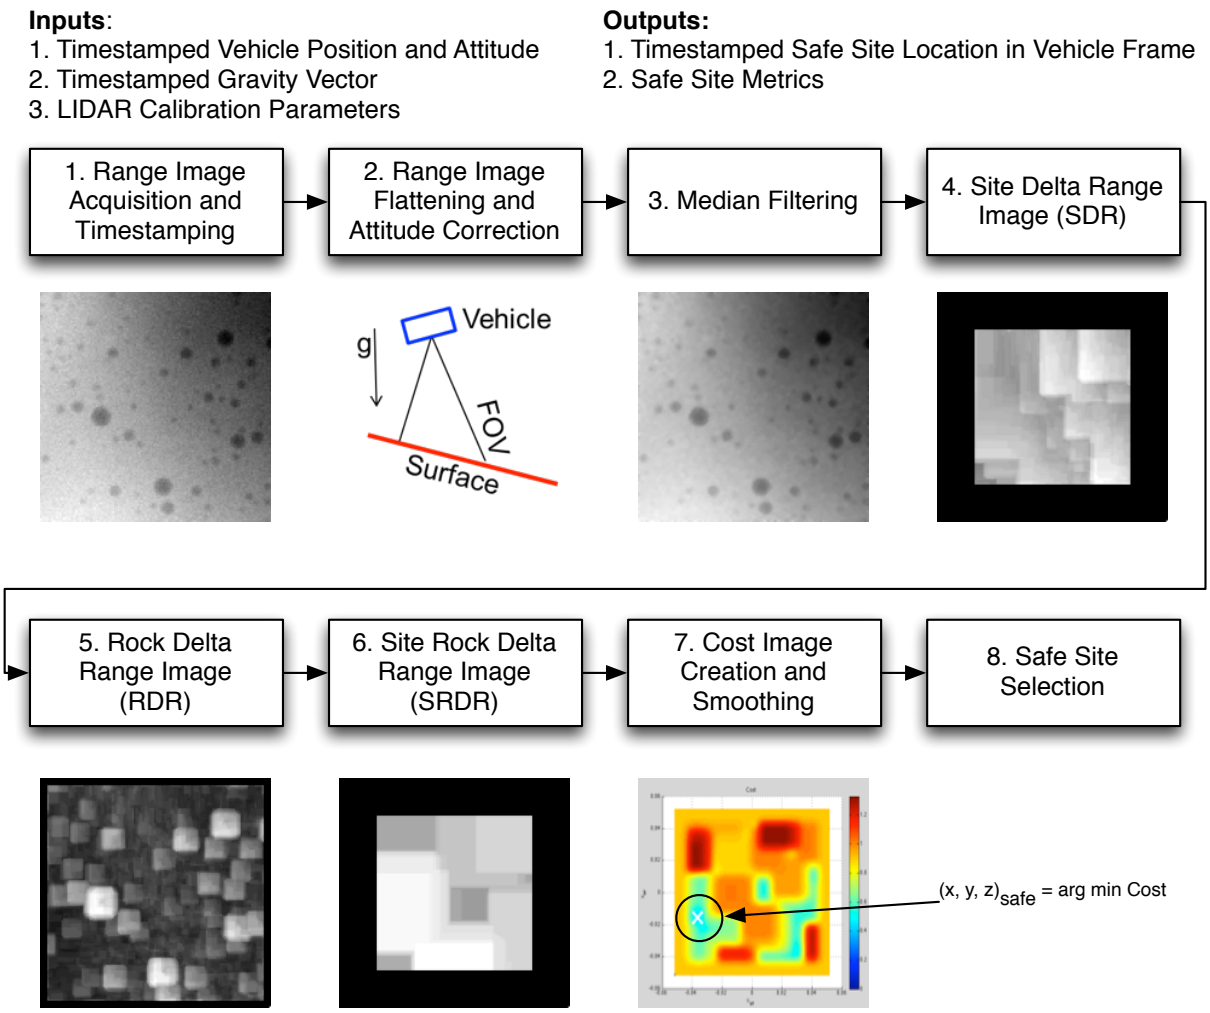

Figure 8. Simple Safe Site Selection Algorithm block diagram.

was within the keep-out borders and hence did not contribute to the S4 cost. Also note the small hazard in the bottom center of the range image that pushed the safe site upward. This hazard was just below the vehicle roughness tolerance and therefore did not contribute to the ground truth safety image. The selected safe site is correctly placed between the hazards.

Table 1. S4 parameters

S4 Parameter	Value
LIDAR field of view	3°
Lander radius	1.11 m
Rock hazard radius	0.35 m
Rock delta range threshold	0.35 m
Slope tolerance	10°
GNC touchdown error 1-σ (i.e. the navigation uncertainty)	1.0 m

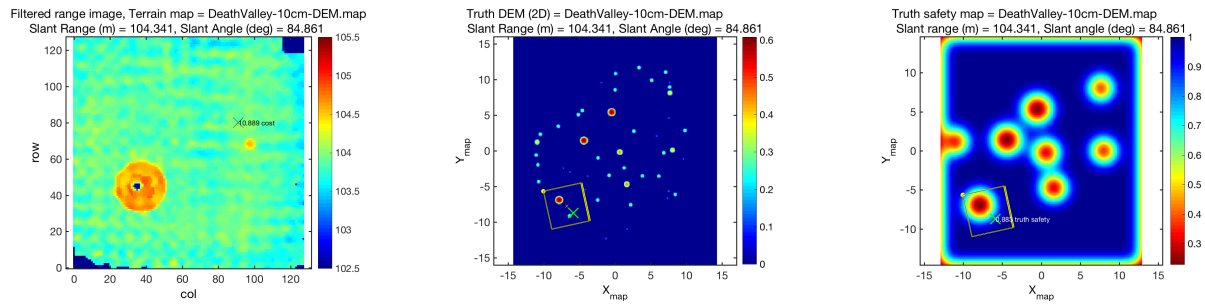

(a) (b) (c)

Figure 9. (a) Pre-processed range image 889, (b) its footprint in the truth DEM, and (c) its footprint in the truth safety map.

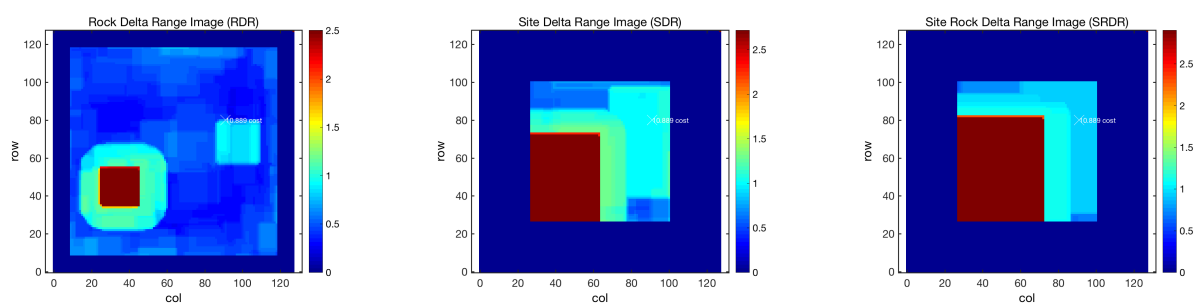

(a) (b) (c)

Figure 10. S4 intermediate products for range image 889: (a) RDR, (b) SDR, (c) SRDR.

VIII. Conclusion

Processing of the flash LIDAR data acquired during a 2014 helicopter field test at Death Valley, CA, demonstrated the performance of the S4 algorithm on actual range images subject to invalid pixels and sensor noise. Absence of time-synchronized ground truth data posed a particular challenge for performance analysis, which was overcome by a two-step precision alignment in post-processing. The difficulties encountered during this process underscored again the importance of careful pre-flight intrinsic, extrinsic, and temporal sensor calibration as well as ground testing, which will be emphasized prior to future, dedicated HD field tests. A ground truth DEM and safety map generated from survey data served to validate the S4 performance, and

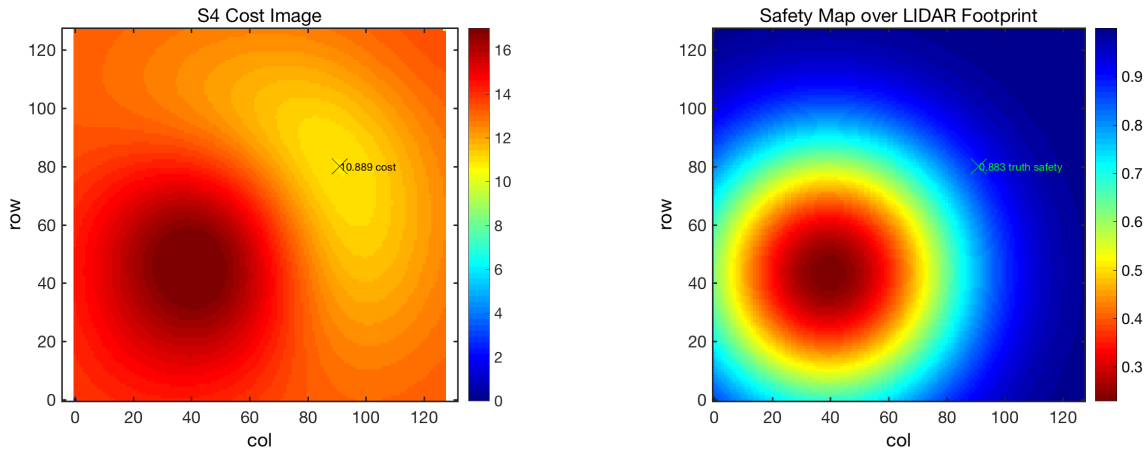

(a) Image 889 S4 cost image

(b) Image 889 truth safety image

Figure 11. S4 final products for range image 889: S4 cost image (center), and ground truth safety image (right).

Table 2. S4 Safe Site Metrics

Range Image #	Min Cost	Selected Safe Site Coordinate (ENU)	True Safe Landing Probability	Max. True Safe Landing Prob. in Footprint
889	10.889	-5.858, -8.731, 0.061	0.883	1.000
3529	5.120	-7.426, -4.359, 0.260	0.850	0.999
3565	5.935	-7.155, 1.310, -0.119	0.912	0.997
3773	3.843	-1.454, 2.204, 0.094	0.876	0.999
3782	5.069	1.021, -2.265, -0.131	0.716	0.997

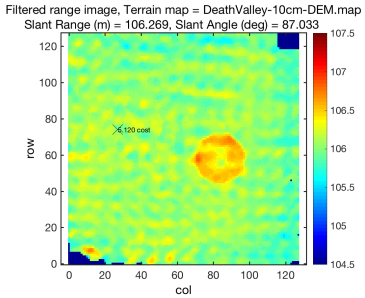

(a) Pre-processed range image 3529

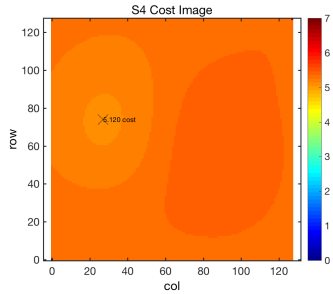

(b) Image 3529 S4 cost image

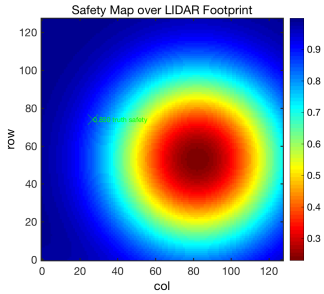

(c) Image 3529 truth safety image

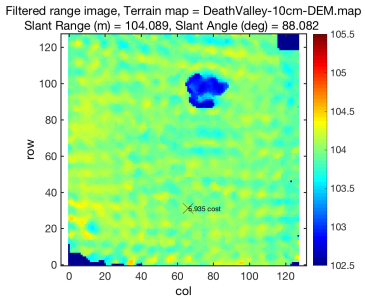

(d) Pre-processed range image 3565

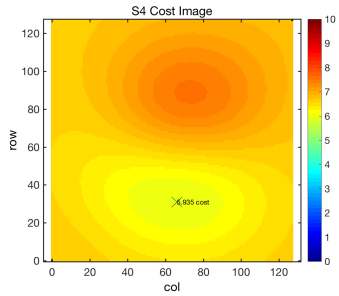

(e) Image 3565 S4 cost image

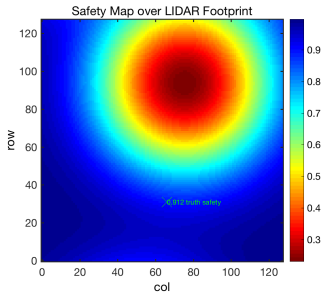

(f) Image 3565 truth safety image

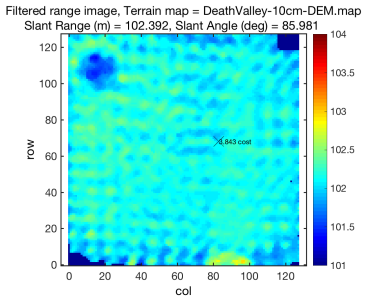

(g) Pre-processed range image 3773

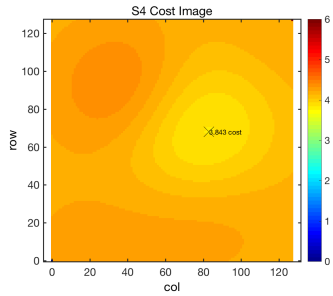

(h) Image 3773 S4 cost image

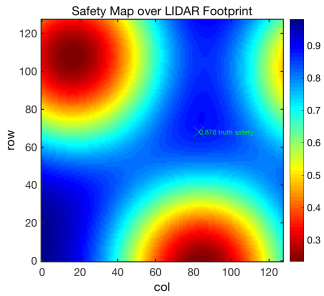

(i) Image 3773 truth safety image

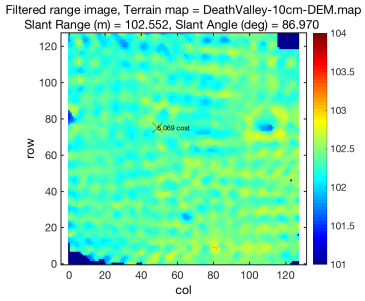

(j) Pre-processed range image 3782

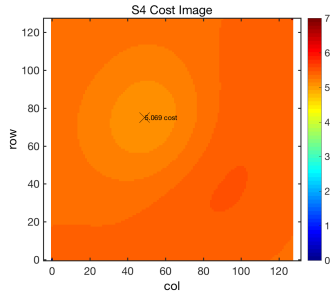

(k) Image 3782 S4 cost image

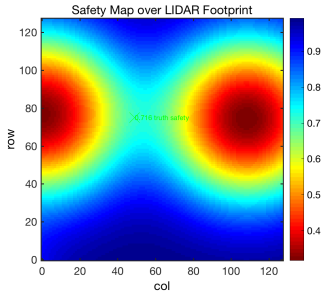

(l) Image 3782 truth safety image

Figure 12. Filtered LIDAR range image (left), S4 cost image (center), and ground truth safety image (right).

confirmed that the cost images created by S4 visually matched the ground truth safety map topography. For five selected range images containing considerable landing hazards, S4 recommended safe sites with true safety probabilities between 72% - 91%. At the same time, S4 can run in 1.3 seconds on a 60MHz flight processor,[29] underscoring its applicability for small robotic lander missions.

Acknowledgments

The work described herein was performed within the Lander Technology project at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. Government sponsorship acknowledged.

References

[1]Brady, T., Robertson, E., Epp, C., Paschall, S., and Zimpfer, D., "Hazard Detection Methods for Lunar Landing," *2009 IEEE Aerospace conference*, March 2009, pp. 1–8.

[2]Neveu, D., Mercier, G., Hamel, J.-F., Simard Bilodeau, V., Woicke, S., Alger, M., and Beaudette, D., "Passive versus active hazard detection and avoidance systems," *CEAS Space Journal*, Vol. 7, No. 2, 2015, pp. 159–185.

[3]Bajracharya, M., "Single image based hazard detection for a planetary lander," *2002 Proceedings of the 5th Biannual World Automation Congress*, Vol. 14, 2002, pp. 585–590.

[4]Cohanim, B. E., Hoffman, J. A., and Brady, T., "Hazard detection for small robotic landers and hoppers," *2013 IEEE Aerospace Conference*, March 2013, pp. 1–13.

[5]Crane, E. S., *Vision-based hazard estimation during autonomous lunar landing*, Ph.D. thesis, Stanford University, Stanford, CA, June 2014.

[6]Cuseo, J., Fleming, P., Hoff, W., Sklair, C., Eshera, M., and Whitten, G., "Machine Vision Techniques for Planetary Terminal Descent Hazard Avoidance and Landmark Tracking," *American Control Conference*, June 1991, pp. 1406–1407.

[7]Huertas, A., Cheng, Y., and Madison, R., "Passive imaging based multi-cue hazard detection for spacecraft safe landing," *2006 IEEE Aerospace Conference*, 2006, pp. 1–14.

[8]Mahmood, W. and Saaj, C. M., "Pinpoint planetary landers: Vision based hazard detection and obstacle avoidance," *2009 European Control Conference (ECC)*, Aug 2009, pp. 3045–3051.

[9]Matthies, L., Huertas, A., Cheng, Y., and Johnson, A., "Stereo vision and shadow analysis for landing hazard detection," *IEEE International Conference on Robotics and Automation (ICRA)*, May 2008, pp. 2735–2742.

[10]Woicke, S. and Mooij, E., "Passive Hazard Detection for Planetary Landing," *AIAA Guidance, Navigation, and Control Conference*, American Institute of Aeronautics and Astronautics, Jan. 4–8, 2016.

[11]Pollard, B. D., Sadowy, G., Moller, D., and Rodriguez, E., "A millimeter-wave phased array radar for hazard detection and avoidance on planetary landers," *Proc. 2003 IEEE Aerospace Conference*, Vol. 2, March 2003, pp. 1115–1122.

[12]Chakroborty, S., Aboutalib, O., and Meade, C., "Autonomous Lunar Lander Hazard Detection and Avoidance System," *AIAA SPACE 2009 Conference & Exposition*, American Institute of Aeronautics and Astronautics, Sept. 14–17, 2009.

[13]de Lafontaine, J., Neveu, D., and Lebel, K., "Autonomous Planetary Landing Using a LIDAR Sensor: the Closed-Loop System," *Guidance, Navigation and Control Systems*, Vol. 606 of *ESA Special Publication*, Jan. 2006.

[14]Frampton, R., Ball, J., Oittinen, K., Tandon, S., Schwab, D., Bishun, M., Richards, B., and Tripp, J., "GN&C, and Precision Landing and Hazard Avoidance Technology Demonstration," *AIAA SPACE 2007 Conference & Exposition*, American Institute of Aeronautics and Astronautics, Sept. 18–20, 2007.

[15]Amzajerdian, F., Pierrottet, D., Petway, L. B., Hines, G. D., Roback, V. E., and Reisse, R. A., "Lidar Sensors for Autonomous Landing and Hazard Avoidance," *AIAA SPACE 2013 Conference and Exposition*, American Institute of Aeronautics and Astronautics, Sept. 10–12, 2013.

[16]Johnson, A. E., Klumpp, A. R., Collier, J. B., and Wolf, A. A., "Lidar-Based Hazard Avoidance for Safe Landing on Mars," *Journal of Guidance, Control, and Dynamics*, Vol. 25, No. 6, Nov.–Dec. 2002, pp. 1091–1099.

[17]Camara, F., Rogata, P., Di Sotto, E., Caramagno, A., Rebordao, J. M., Correia, B., Duarte, P., and Mancuso, S., "Design and Performance Assessment of Hazard Avoidance Techniques for Vision Based Landing," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, American Institute of Aeronautics and Astronautics, Keystone, CO, Aug. 21–24, 2006.

[18]Kolb, F. M., Ahrns, I., and Moebius, B., "3D Imaging LIDAR for Autonomous Planetary Landing," *Proc. i-SAIRAS*, Sapporo, Japan, Aug. 29 – Sept. 1, 2010, pp. 152–156.

[19]Parreira, B., Rogata, P., Di Sotto, E., Caramagno, A., Rebordao, J., Motrena, P., and Mancuso, S., "Consolidated Performance Assessment of Hazard Avoidance Techniques for Vision Based Landing," *AIAA Guidance, Navigation and Control Conference and Exhibit*, American Institute of Aeronautics and Astronautics, Hilton Head, SC, Aug. 20–23, 2007.

[20]Parreira, B., Vasconcelos, J., Montaño, J., Ramón, J., and Penin, L. F., "Hazard Detection and Avoidance in ESA Lunar Lander: Concept and Performance," *AIAA Guidance, Navigation, and Control (GNC) Conference*, American Institute of Aeronautics and Astronautics, Boston, MA, Aug. 19–22, 2013.

[21]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," *AIAA SPACE 2013 Conference and Exposition*, American Institute of Aeronautics and Astronautics, San Diego, CA, Sept. 10–12, 2013.

[22]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.

[23]Carson, J. M., Robertson, E., Trawny, N., and Amzajerdian, F., "Flight Testing ALHAT Precision Landing Technologies Integrated Onboard the Morpheus Rocket Vehicle," *AIAA SPACE 2015 Conference and Exposition*, American Institute of Aeronautics and Astronautics, Aug. 31– Sept. 2, 2015.

[24]Epp, C. D. and Smith, T. B., "Autonomous Precision Landing and Hazard Detection and Avoidance Technology (ALHAT)," *2007 IEEE Aerospace Conference*, March 2007, pp. 1–7.

[25]Epp, C., Robertson, E., and Carson, J. M., "Real-Time Hazard Detection and Avoidance Demonstration for a Planetary Lander," *AIAA SPACE 2014 Conference and Exposition*, American Institute of Aeronautics and Astronautics, San Diego, CA, Aug. 4–7, 2014.

[26]Trawny, N., Huertas, A., Luna, M. E., Villalpando, C. Y., Martin, K., Carson, J. M., Johnson, A. E., Restrepo, C., and Roback, V. E., "Flight testing a Real-Time Hazard Detection System for Safe Lunar Landing on the Rocket-Powered Morpheus Vehicle," *AIAA Guidance, Navigation, and Control Conference*, American Institute of Aeronautics and Astronautics, Kissimmee, FL, Jan. 5–9, 2015.

[27]Villalpando, C. Y., Werner, R. A., Carson, J. M., Khanoyan, G., Stern, R. A., and Trawny, N., "A hybrid FPGA/Tilera compute element for autonomous hazard detection and navigation," *Aerospace Conference, 2013 IEEE*, March 2013, pp. 1–9.

[28]Jiang, X., Li, S., and Tao, T., "Innovative hazard detection and avoidance strategy for autonomous safe planetary landing," *Acta Astronautica*, Vol. 126, 2016, pp. 66 – 76.

[29]Johnson, A. E. and Mandalia, A. B., "Simple Safe Site Selection: Hazard Avoidance Algorithm Performance at Mars," *Proc. 37th AAS Rocky Mountain Guidance and Control Conference*, Breckenridge, CO, Jan. 31 – Feb. 5, 2014.

[30]Ivanov, T., Huertas, A., and Carson, J. M., "Probabilistic Hazard Detection for Autonomous Safe Landing," *AIAA Guidance, Navigation, and Control (GNC) Conference*, American Institute of Aeronautics and Astronautics, Boston, MA, Aug. 19–23, 2013.

[31]Johnson, A. E., Cheng, Y., Montgomery, J. F., Trawny, N., Tweddle, B., and Zheng, J. X., "Real-Time Terrain Relative Navigation Test Results from a Relevant Environment for Mars Landing," *Proc. AIAA Guidance, Navigation, and Control Conference*, Kissimmee, FL, Jan. 05 - 09, 2015.

[32]Torr, P. H. S. and Zisserman, A., "MLESAC: A New Robust Estimator with Application to Estimating Image Geometry," *Computer Vision and Image Understanding*, Vol. 78, 2000, pp. 138–156.