

Processing of 3-dimensional flash lidar terrain images generated from an airborne platform

Alexander Bulyshev^a, Diego Pierrottet^b, Farzin Amzajerdian^c, George Busch^b, Michael Vanek^c, and Robert Reisse^c,

^aAMA-Inc, Hampton, Virginia

^bCoherent Applications, Inc, Hampton, Virginia

^cNASA, Langley Research Center, Hampton, Virginia

ABSTRACT

Data from the first Flight Test of the NASA Langley Flash Lidar system have been processed. Results of the analyses are presented and discussed. A digital elevation map of the test site is derived from the data, and is compared with the actual topography. The set of algorithms employed, starting from the initial data sorting, and continuing through to the final digital map classification is described. The accuracy, precision, and the spatial and angular resolution of the method are discussed.

Key words: LIDAR, Flash Ladar, 3-D imaging, Hazard Detection, Planetary Landing

1. INTRODUCTION

As part of its Autonomous Landing Hazard Avoidance Technology (ALHAT) program^{1,2} NASA is evaluating the applicability of a Laser Radar (LIDAR) sensor to detect, identify, and avoid hazards during descents to planetary surfaces. LIDAR uses the round-trip time of light pulses reflected from distant targets to measure the range or distance to the target, in direct analogy to standard RADAR, which uses Radio wave pulses. Three-dimensional (3-D) images of a target surface may then be obtained by rastering the laser over the surface, i.e., sending out and detecting separate pulses to each transverse (x,y) position to be resolved on the surface. However, this standard LIDAR approach requires a significant overhead in time and computational complexity to geo-locate a large number of objects in time and space, at high resolution. These limitations are particularly important when rapid decisions based on the measurements must be made in near real time from a moving platform.

In a new technique under development, called Flash Lidar, instead of rastering the laser beam, a single laser pulse (flash) is sent out, illuminating the entire target area and the light returning from each resolvable transverse position (pixel) is detected with a camera, in the form of an image. This single-pulse technique eliminates both the time it takes to gather rastered data, and the computational complexity of removing the platform motion to recover target surface detail. In order to get the desired 3-D information, each pixel must separately be able to record the round-trip time of light arriving at that pixel. Such a camera has been developed and demonstrated^{3,4} by Advanced Scientific Concepts (ASC), Inc.

Previously, static measurements of the ASC camera performance were reported² for measurements made from a fixed platform at the Sensor Test Range (STR) of NASA-Langley Research Center. More recently, dynamic airborne tests of the ASC Flash Lidar system have been performed by NASA, aboard an AS-350 helicopter. The current paper analyzes some of the Flash Lidar data obtained during these first flight tests. This analysis includes initial data conversion, signal data processing, noise analysis, image noise filtering, calculation of digital elevation maps, and procedures for recognizing and characterizing objects within a scene. Color-coded images are used to display either the intensity of signal returns for each pixel, or the range distance to the surface at each pixel position, and methods to derive and interpret these images are presented.

2. FLASH LIDAR

The laser transmitter of the ASC Flash Lidar system operates at the eye-safe wavelength of 1570 nm, with pulse lengths on the order of 10 ns FWHM. Detection of the returning spatially resolved reflected light pulses is via an InGaAs focal plane array camera with each pixel operated in avalanche mode. The required time dependence of the return, needed to yield a 3-D image, is achieved by starting a counter at the time the laser pulse is sent out by the transmitter, then reading the counter and storing that count at which the return pulse is detected. This is done independently for each pixel in the focal plane array, so each pixel must be a “smart” pixel to carry out the required independent signal processing. The processing is accomplished in each unit cell of the proprietary silicon read-out integrated circuit (ROIC) bonded to InGaAs focal plane array sensor. Additional waveform temporal information is also recorded and stored for each pixel, permitting post processing range resolution enhancements. More technical details about ASC Flash Lidar systems can be found in ASC publications³⁻⁶. A photograph of the Flash Lidar laser and camera head is shown in Fig.1.



Fig. 1. ASC Flash lidar with a moderate power laser (box with vertical slots, below) and camera mounted above, with a 250 mm, f/2 lens. Image courtesy of Advanced Scientific Concepts, Inc.

3. FIELD TEST DESCRIPTION

During the test event several test sites were used for observations and testing of the flash lidar. Below we will concentrate on the data obtained from one of them: Dryden Borrow Pit Rock and Crater Hazard Site. The general view of the site is presented in Figure 2.



Fig. 2 Dryden Borrow Pit Rock and Crater Hazard Site one of several sites used during this flight test. A number of rectangular and hemispherical objects with varying sizes from 30 cm to several meters placed and their location registered along with several man-made craters.

Flight tests were performed with a suite of instruments onboard the helicopter. In addition to the Flash Lidar system, two visible digital cameras were used to record the same test site, one with a narrow field of view for high resolution, the other with a wide field of view for increased scene identification. All measurements for the tests were made from 250m to 500m height above the ground with different viewing slant range angles..

4. DATA ANALYSIS

4.1 General description

Each Flash Lidar data set is composed of the 504 consecutive frames taken at a frame rate of 8 Hz. Depending on the helicopter velocity and trajectory, the degree of spatial overlap among consecutive frames varied from almost no overlap to almost super imposable. Signals are completely saturated for the frames taken close to the surface, whereas signals were so weak for ranges exceeding 500 m that only a few pixels received sufficient signal to trigger and record a return time. We chose to make most of the measurements in the intermediate range 200 – 300 m so that image intensities remained mostly in a linear or near linear regime.

Because the quantity of data is very large, the analyses were limited to frames of greatest interest. As the first step of the data analysis we sorted through the frames and eliminated from further analyses those frames that did not contain images that met the interest criteria, as described in more detail below. Two major criteria are used in the analysis of the all frames: a) to find frames which have close to 100 % of the pixels triggered, but for which the signal is not completely saturated; and b) to find frames which are close to the limit of detection for the current version of hardware. The fraction of triggered events for the low detection limit frames is between 50 and 70%. In all cases, only frames with visible test objects (boxes and craters) were selected for further analysis.

Each frame is associated with two types of information at each pixel. One of them is the intensity of the Lidar return signal. The other is the range to the surface giving the lidar return at that pixel. The intensity information can be displayed as a standard video image, and we refer to it here as an intensity image. The range information may also be displayed as a video image by employing a color-coded representation of the range. We call the latter representation a range image. Initial (coarse resolution) range images are obtained by simply counting the number of clock ticks to the received event trigger. This measurement of signal return time has an uncertainty of the order of one clock tick, with a clock period being about 2.4 ns. This resolution is sufficient for quick evaluation of large amounts of data. Frames with image data outside expected range limits were excluded from further analyses. A more accurate, but time consuming approach is then applied when individual images are processed with the objective of obtaining as much resolution as possible. More details about image formation are discussed below.

4.2 Data processing

Received pulse samples of each pixel are saved in a circular buffer that holds 20 samples at one time. This architecture causes the data to appear as if it is wrapped in time. The first step of the data processing is to unwrap the samples for each individual pixel into the correct temporal order. Fig 3 contains the waveforms before (A) and after (B) unwrapping the samples to the correct order in time. The dip in sample 10 is a marker corresponding to the last sample.

Since the laser pulse has significant time duration, the exact position on the signal pulse that should be associated with return time is important. As it was suggested in Ref. 2, we use the position of the signal maximum as the correct round-trip flight time of the pulse. The data is sampled every 2.4 nanoseconds, so simply finding the clock count position with maximum value of the returning signal would lead to an uncertainty of the order of the sampling period.

Using a functional fit to the signal waveform helps to find the signal maximum position with an accuracy that is a small fraction of the sample time. A 6 degree polynomial approximation was used to create a signal model from the data. In order to estimate the quality of this approach, the following computational experiment was performed.

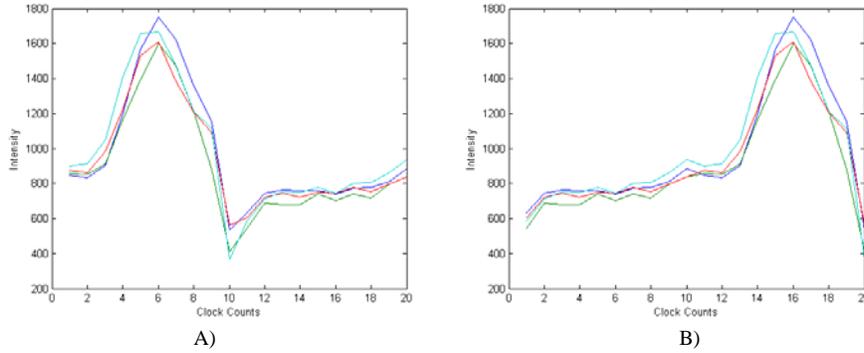


Fig 3 Signal waveforms before and after time reordering

We chose a Gaussian curve with the random additional noise as the signal model. The noise was normally distributed with zero mean value and for various levels of standard deviation. Two examples are presented in Fig 4.

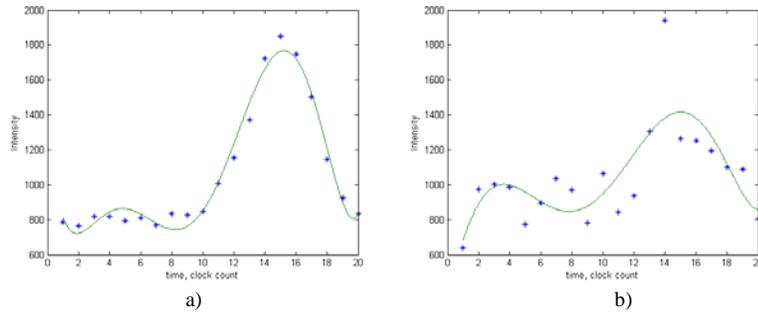


Fig 4 The results of the restoration of the signal maximum positions for different level of noise: a) the noise level is equal to 5% of the signal maximum; b) the noise level is equal to 30% of the signal maximum.

The results of the modeling show that the method is accurate and robust. It provides an accuracy of the time position of the pulse maximum that is a small fraction of the sample time for noise levels below 10%. The method gives reasonable improvement in timing accuracy for the case of extremely high levels of noise (up to 30%). The results show that the accuracy of the range detection is much better than the 36 cm distance corresponding to one sample time for all reasonable noise levels. The results for different levels of noise are summarized in the Table 1.

Table 1. Mean Error and Standard Deviation of measured range per pixel.

Noise Level	Mean, cm	STD cm
5%	2.2	1.5
10%	3.3	2.6
20%	6.9	4.8
30%	9.9	8.1

4.3 Images in intensity and range

After the time position and the maximum signal values for each pixel are determined from the polynomial fit parameters, the images in intensity and range are formed using those values. Examples of intensity and range images from flight tests over the test site shown in Fig. 3 are presented in Fig 5.

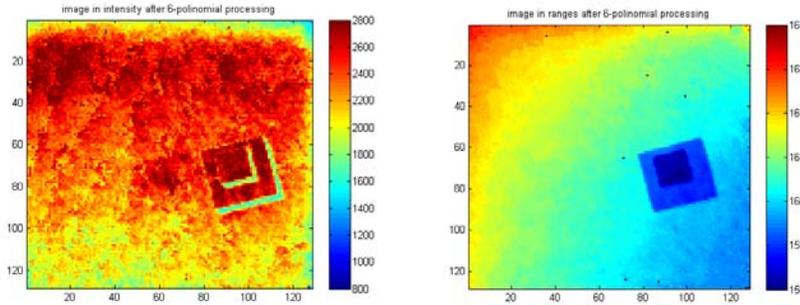


Fig 5 Images in Intensity and Range after fitting signal waveforms to a 6th order polynomial simulation of the return pulse

5. IMAGE RECONSTRUCTION

Data processing procedures were applied to each pixel channel individually without any knowledge or hypothesis about possible correlation between different pixels in the image. Image reconstruction procedures assume that pixels are not independent, but are the parts of a continuous image. So, the correlation between different pixels is expected. In this paper, the term “image reconstruction” refers to a procedure employing several image improvement techniques.

5.1 Filtering

The most obvious class of image reconstruction procedures is that of filtering⁵. The general purpose of filtering is to decrease the influence of “noise”, where noise is defined as any unwanted changes to the image, like random registration noise or errors in returned pulse peak detection. In this case, two major kinds of noise should be treated by filtering procedures; a) “missing” pixels; and b) random uncorrelated noise. Missing pixel noise occurs when some channels do not have any signal. This means that the intensity being received by a given pixel never crossed a preset threshold value to trigger the storage of the clock time and to record the received pulse waveform. In such cases, the clock times recorded are preset maximum range gate values corresponding to the expected range limits. The second source of noise is mostly random noise in the recording of waveforms which leads to an error in the position of the waveform maximum. To decrease the influence of both sources of noise we use the procedure of median filtering⁷. A 9-point median filtering scheme is a particularly useful procedure for range image correction. The algorithm is implemented using the following equations:

$$\tilde{r}_{ij} = \text{median}(r_{i-1,j} + r_{i,j} + r_{i+1,j} + r_{i,j-1} + r_{i,j+1} + r_{i+1,j+1} + r_{i-1,j+1} + r_{i-1,j-1} + r_{i+1,j-1}) \quad (1)$$

where (i,j) are the pixel coordinates, $r_{i,j}$ is the range value before applying the filtering procedure, and $\tilde{r}_{i,j}$ is the range value after filtering. While, equation (1) applies to all internal pixels, the filtering procedure reduces to 6-point or 4-point schemes for the boundary pixels. The choice of median filtering has been made mostly because this procedure provides good filtering against random noise, works very well against isolated missing points, and leaves

Comment [DFP1]: Is this information necessary? Its not clear to me how important this is for the current content without being too negative on ASC's work. I think Farzin will have a problem with “problems”

large gradients practically intact. Any linear filtering procedure (averaging for example) proved to be unsatisfactory for filtering isolated missing points or for handling large gradients.

5.2 Position and orientation correction algorithm

With knowledge of the position and orientation of the Flash Lidar, it is possible to transform the range image into an elevation map. If the position of the Flash Lidar at the time an image is recorded is (X_o, Y_o, Z_o) and the vector in the laser pointed direction is \vec{N} , the absolute coordinates of the pixel on the map can be calculated as:

$$\begin{aligned} X &= X_o + N_x r_{ij} \\ Y &= Y_o + N_y r_{ij} \\ Z &= Z_o + N_z r_{ij} \end{aligned} \tag{2}$$

Selecting the local position of the test site as the reference point, only the relative elevation coordinates were computed. Applying equation (2) to the range array, the elevation map in the Flash Lidar coordinate system is obtained. The results of the angle correction procedure are shown below in Fig 6:

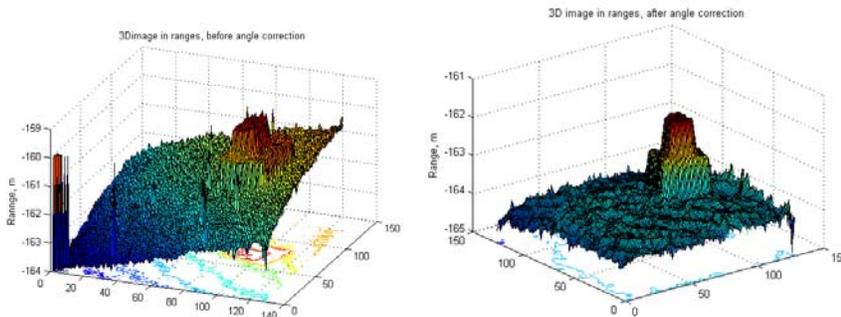


Fig. 6. 3D range images. a) Before, and b) after applying the angle correction procedure; Negative sign on the range means that the “0” level is at the helicopter level.

5.3 Digital Elevation Map (DEM)

After correcting the angle of view, it is straight forward to calculate a Digital Elevation Map. The average elevation is calculated from areas without objects present. A transformation of the range image into a Digital Elevation Map is made by subtracting the average elevation from the elevation at each point. The results are presented in Fig. 7:

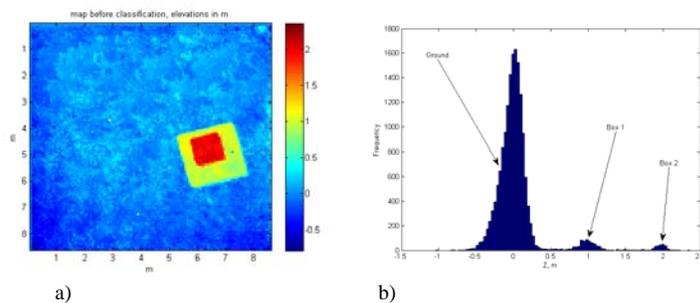


Fig. 7. a) Digital Elevation Map, and b) Histogram of elevations

The DEM shows positions and sizes of the boxes on the flat surface. The histogram has a very clear interpretation as there are only three distinct elevation levels: 1) ground level with the average value equal to zero, 2) box #1 with average elevation level equal to 1 m, and 3) box #2, with the average elevation level equal to 2 m.

6. SCENE EVALUATION

A DEM is essentially a 2D array of pixel elevations covering the area of the scene under investigation. To make an evaluation of whether a scene is hazard free; a second transformation of the DEM is executed after the position and orientation correction described above. The major idea consists of two steps: segmentation and features identification.

6.1 Optimal Segmentation Algorithm - Binary Approach

An Optimal Segmentation Algorithm (OSA) utilizing a Binary Approach divides a set of objects into subsets (segments) with statistically different values for the feature defining the segment. The method consists of the following steps:

1. The total area is considered as one segment. The average elevation and the standard deviation of elevations are calculated for this segment.

$$\begin{aligned}\bar{H} &= \frac{1}{N} \sum_{k=1}^{k=N} h_k \\ DH &= \frac{1}{N-1} \sum_{k=1}^{k=N} (h_k - \bar{H})^2\end{aligned}\quad (3)$$

In (3) N is the total number of pixels in 2D image, \bar{H} is the average elevation for the area, h_k is the elevation for the pixel number k , and DH is the standard deviation of elevations for the area.

2. Both the x and y axes are divided into two equal segments, so that four new segments are formed. The average elevations and standard deviations for each new segment are calculated:

$$\begin{aligned}\bar{H}_m &= \frac{1}{N_m} \sum_{k \in S_m} h_k \\ DH_m &= \frac{1}{N_m - 1} \sum_{k \in S_m} (h_k - \bar{H}_m)^2\end{aligned}\quad (4)$$

In (4) S_m is a segment number m , N_m is the number of pixel in the segment number m , \bar{H}_m is the average elevation for the segment number m , and DH_m is the standard deviation of elevations for the segment number m .

3. If the following inequality is valid for only one new segment

$$|\bar{H} - \bar{H}_m| \geq \xi \times (DH + DH_m) \quad (5)$$

then the segmentation is accepted, otherwise the segmentation for the segment is stopped. The parameter ξ is a tuning parameter which serves to increase or decrease the details of the object segments. Typically ξ is set equal to 2.

4. The process is stopped when no more segments are available for division.

The results of OSA applied to two-box objects are presented in Fig 8.

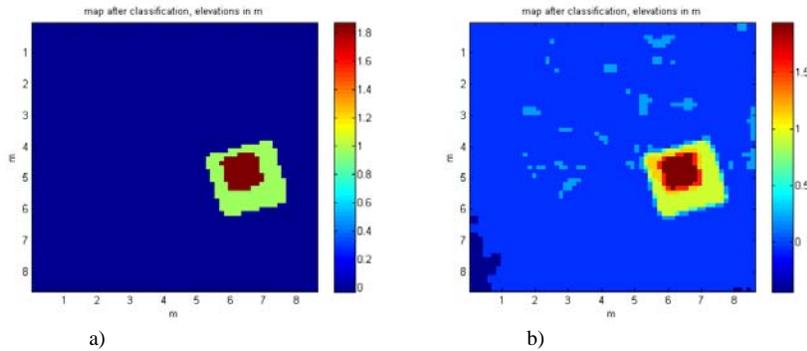


Fig 8. Classified Maps; a) $\xi = 6$; b) $\xi = 2$

Both pictures show two box objects very clearly. The more detailed map (b) also shows some roughness on the surface.

7. CONCLUSIONS

A 3-D imaging Flash lidar collected data of known targets from a helicopter at altitudes reaching 500 m. Digital Elevation Maps (DEMs) were generated from the raw lidar data by applying the appropriate filtering followed by position and orientation correction transformation. The DEMs were used to demonstrate scene evaluation and identification of potential hazards using an Optimal Segmentation Algorithm demonstrating the lidar capabilities in performing hazard detection and avoidance for future NASA landing missions to the Moon and Mars.

8. ACKNOWLEDGMENTS

The authors would like to acknowledge the contributions of NASA Jet Propulsion Laboratory ALHAT team for the helicopter flight tests support and producing the flight trajectories data from onboard inertial navigation units and GPS receivers. The authors are thankful to Advanced Scientific Concepts for providing the Flash lidar system and their assistance leading to this campaign. The authors are also thankful to the ALHAT project manager, Chirold Epp, NASA Johnson Space Center for his support.

REFERENCES

- [1] Epp, C.D., Robinson, E.A., and Brady, T., "Autonomous Landing and Hazard Avoidance Technology (ALHAT)", Aerospace Conference, 2008 IEEE, pp.1-7.
- [2] Pierrottet, D., Amzajerjian, F., Meadows, B., Estes, R., Noe, A., "Characterization of 3-D imaging lidar for hazard avoidance and autonomous landing on the Moon," Proc. of SPIE Vol. 6550 2007.
- [3] Stettner, R., Bailey, H., and Richmond R., " Eye Safe Laser Radar Focal Plane Array for Three-Dimensional Imaging", Proc. of the SPIE Vol. 4377, 2001.
- [4] Stettner, R., Bailey, H., and Silverman S., "Three Dimensional Flash Ladar Focal Planes and Time Dependent Imaging," International Symposium on Spectral Sensing Research, Bar Harbor, Maine, 2006.
- [5] Castleman, K., [Digital Image Processing], Prentice Hall, 1996.