

# Field Test Implementation to Evaluate a Flash Lidar as a Primary Sensor for Safe Lunar Landing

Jason A. Keim, Sohrab Mobasser, Da Kuang, Yang Cheng,  
 Tonislav Ivanov, Andrew E. Johnson, Hannah R. Goldberg,  
 Garen Khanoyan, and David B. Natzic  
 Jet Propulsion Laboratory  
 California Institute of Technology  
 Pasadena, CA 91109  
 Jason.A.Keim@jpl.nasa.gov

*Abstract*—From May 2 through May 7 of 2008, the Autonomous Landing and Hazard Avoidance Technology (ALHAT) Exploration Technology Development Program carried out a helicopter field test to assess the use of a flash LIDAR as a primary sensor during lunar landing. The field test data has been used to evaluate the performance of the LIDAR system and of algorithms for LIDAR Hazard Detection and Avoidance, Hazard Relative Navigation, and Passive Optical Terrain Relative Navigation. Reported here is a comprehensive description of the field test hardware, ground infrastructure and trajectory reconstruction methodologies<sup>1,2</sup>.

and identify key performance parameters and sensitivities, several field tests are planned.

The objectives of first ALHAT field test, which we refer to as FT1 (Field Test 1), are to assess the performance of a commercial LIDAR sensor and the performance of Hazard Detection (HD), HRN and TRN algorithms. The test results are already in the literature: [1] provides the LIDAR data processing and performance, [9] present the TRN results, and [5] gives the HD and HRN results. This paper documents the FT1 hardware, ground infrastructure and trajectory reconstruction methodologies.

## TABLE OF CONTENTS

<b>1. INTRODUCTION</b> .....	<b>1</b>
<b>2. SENSORS AND SUPPORT EQUIPMENT</b> .....	<b>1</b>
<b>3. DATA COLLECTION</b> .....	<b>3</b>
<b>4. TARGET SITES</b> .....	<b>4</b>
<b>5. TRAJECTORY RECONSTRUCTION</b> .....	<b>5</b>
<b>6. INSTRUMENT VERIFICATION</b> .....	<b>10</b>
<b>7. CONCLUSIONS</b> .....	<b>10</b>
<b>ACKNOWLEDGMENTS</b> .....	<b>10</b>
<b>REFERENCES</b> .....	<b>11</b>
<b>BIOGRAPHY</b> .....	<b>11</b>
<b>APPENDIX</b> .....	<b>13</b>

## 1. INTRODUCTION

The ALHAT Project is funded by NASA to develop an integrated AGNC (Autonomous Guidance, Navigation and Control) hardware and software system capable of detecting and avoiding surface hazards and guiding humans and cargo safely, precisely and repeatedly to designated lunar landing sites [2]. Achieving this necessitates advancing the state of the art in Terrain Relative Navigation (TRN), Hazard Detection and Avoidance, and Hazard Relative Navigation (HRN) algorithms. This will also require the investigation and advancement of sensor technologies. Currently, the technologies under consideration are passive and active optical systems and radio frequency systems. To characterize applicable sensor technologies and algorithms

## Table 1: List of Acronyms and Abbreviations

ALHAT	Autonomous Landing and Hazard Avoidance Technology
APL	Johns Hopkins University Applied Physics Laboratory
CDSU	Command and Data Storage Unit
DEM	Digital Elevation Model
ECI	Earth Centered Inertial
EKF	Extended Kalman Filter
FT1	Field Test One
GNSS	Global Navigation Satellite Systems
GPS	Global Positioning System
HD	Hazard Detection
HRN	Hazard Relative Navigation
IGS	International GNSS Service
IMU	Inertial Measurement Unit
JPL	NASA Jet Propulsion Laboratory
LaRC	NASA Langley Research Center
LIDAR	Light Detection And Ranging
NASA	National Aeronautics and Space Administration
PPS	Pulse Per Second
RANSAC	Random Sample Consensus
SIFT	Scale-Invariant Feature Transform
TRN	Terrain Relative Navigation

## 2. SENSORS AND SUPPORT EQUIPMENT

Figure 1 is a block diagram of the entire FT1 system. The FT1 test platform was the S.N.I.A.S. AS350D ASTAR helicopter N145BH. To carry the ALHAT payload, a two-axis gimbal was mounted to the aircraft as seen in Figure 2. The gimbal assembly is seen in Figure 3 with the enclosure

<sup>1</sup> 978-1-4244-3888-4/10/\$25.00 ©2010 IEEE.

<sup>2</sup> IEEEAC paper #1686, Version 5, Updated January 5th, 2010

removed. The sensors mounted to the gimbal are the flash LIDAR, two optical still shot cameras, an Inertial Measurement Unit (IMU), an orientation sensor, and a context (witness) video camera. The commercial flash LIDAR is the device under test. One of the still shot cameras provided the attitude of the platform relative to the surface when surveyed targets fell into its field of view. The other camera was co-aligned with the LIDAR and had a comparable field of view. This camera was used to identify what the LIDAR had imaged. The IMU provides velocity and angle rate information and the magnetometer provides heading information. An additional IMU is seen in Figure 3. This IMU is a member of the gimbal stabilization system and is not an ALHAT sensor.

The other ALHAT sensors were also mounted externally to the helicopter. Inside of the gimbal enclosure, but not on the gimbal, was an environment sensor that recorded the temperature and humidity of the instrument environment. Directly above the gimbal was a non-articulating platform to which the GPS antenna was mounted. Additionally, two digital cameras were mounted to the back of the ballast box at the rear of the helicopter. These cameras, seen in Figure 4, were the primary sensors for the APLNav [9] terrain relative navigation system.

Inside the helicopter was an equipment rack. As seen in Figure 5, it housed power supplies, a power distribution and switching box, the PXI chassis, and the LIDAR computer and computer switch. For pre- and post-flight checkout, the power supplies were plugged into a wall socket and powered the instruments. When helicopter power was available, the operator could use the power distribution box to switch from the power supplies to the helicopter power. The box also controlled the power of the individual devices. The PXI chassis contained a National Instruments PXI-6682 time synchronization card, a computer switch and a computer. This system commanded and collected data from the LIDAR and APLNav passive optical cameras.

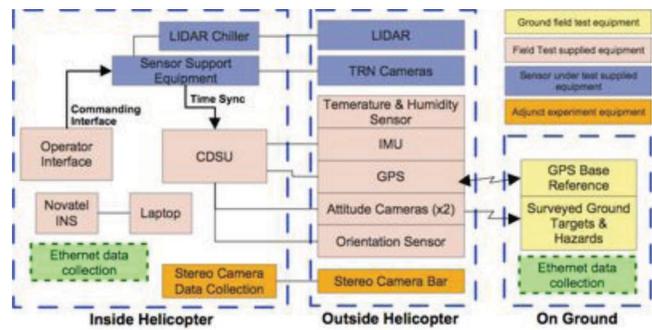
The remainder of this section documents each device on the Gimbal Assembly and the rack mounted equipment.

#### *Inertial Measurement Unit*

The IMU utilized during the test is a Litton LN-200. Three solid-state fiber optic gyros and three solid-state silicon accelerometers comprise the IMU and were sampled at a rate of 400 Hz.

#### *GPS Base Station*

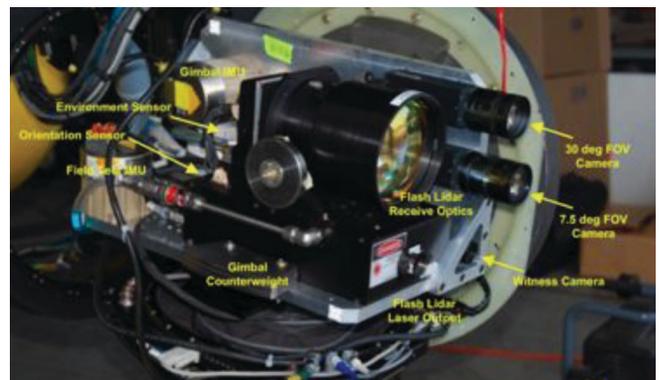
The GPS base station is a statically placed GPS antenna and receiver. The data is stored internal to the receiver and periodically downloaded to a computer. By recording the static location of the GPS, the position measurement accuracy of the helicopter can be increased via relative GPS. In order to support the long data sessions required, the receiver was set to sample data only once a second.



**Figure 1: Payload block diagram**



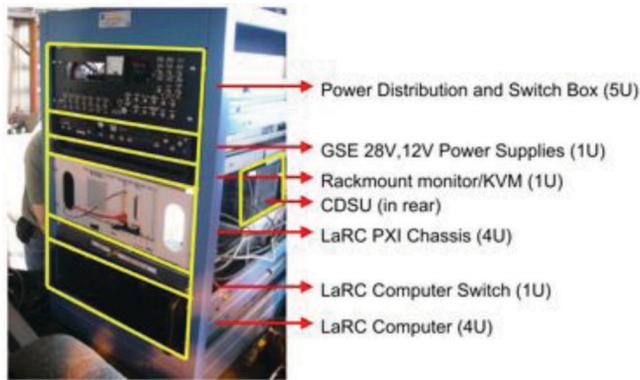
**Figure 2: Helicopter test platform**



**Figure 3: Gimbal assembly**



**Figure 4: Passive optical TRN cameras**



**Figure 5: Equipment rack**

### *Flight GPS Receiver*

The Ashtech MicroZ CGRS GPS receiver combined with a Magellan Dual-band L1/L2 antenna recorded the position of the helicopter. As stated above, the receiver was housed in the equipment rack and the antenna was mounted externally to the helicopter on a fixed platform above the gimbal. To protect against data loss, the GPS data was stored redundantly in the receiver and by the Command and Data Storage Unit (CDSU) described below. In addition to measuring position, the receiver provided a digital pulse per second (PPS). The PPS combined with a 33 MHz, 32-bit counter provided the GPS time tags for each instrument. Initially for FT1, the flight receiver operated at a rate of 10 Hz. However, during post-flight data quality analysis, it was observed that this rate caused gaps in the GPS data. Consequentially, the receiver sample rate was reduced to 5 Hz.

### *Backup GPS Receivers*

NovAtel Inc. ProPak-V3 GPS receivers were the FT1 backup receivers. Although, in our experience, it is relatively easier to operate and configure a NovAtel receiver than an Ashtech receiver, the ProPak-V3 receiver does not output the necessary PPS. Thus, a NovAtel receiver cannot be the primary flight receiver. Instead, the HG1700 SPAN<sub>58</sub> system (ProPak-V3 GPS receiver with IMU) was used only as a backup.

### *Gimbal Mounted Digital Cameras*

Two Sony XCD-SX910 digital cameras were mounted to the gimbal. One camera had a Fujinon HF12.5SA-1 lens that produced a thirty-degree field-of-view. The other camera had a Fujinon DF6HA-1B lens and seven-degree field-of-view. The wider-angled camera served as the attitude sensor of the trajectory reconstruction system. By observing surveyed ground targets, the camera attitude in the GPS frame could be determined. The narrow angle camera was co-aligned with the LIDAR. With a field of view similar to that of the LIDAR, it provided an image that roughly coincided with the LIDAR samples. For this test, both cameras produced monotone (grayscale) 1280×960

pixel images. Set to OneShot mode, the cameras captured a single Portable Graymap Graphic image when triggered to by the CDSU. The CDSU triggered the narrow-angle camera at a rate of 3 Hz and the wide-angle camera at 1 Hz. The exposure times of each camera were manually modified to account for changing lighting conditions.

### *Flash LIDAR*

FT1 sought to establish the feasibility of using a flash LIDAR as the primary sensor for Hazard Detection and Hazard Relative Navigation algorithms. To this end, NASA Langley Research Center (LaRC) provided the flash LIDAR seen in Figure 3. Manufactured by Advanced Scientific Concepts, the LIDAR was a combination of a 1.57 micron, diode-pumped Nd:YAG laser source and receiver optics. The laser source is diffused to provide a three-degree cone to actively illuminate the target site observed by the receiver optics. This receiver, with a field of view slightly less than three-degree, was the combination of a 128×128 pixel, InGaAs array and a 120-mm-diameter, 250-mm-focal-length aperture. By calibrating the focal plane, each pixel of the receiver provides a simultaneous range and bearing measurement. For this test, the sensor operated at eight Hertz.

### *Environment Sensor*

A HOBO U12-013 was attached inside of the gimbal enclosure. Set to sample every 30 seconds, both the temperature and the humidity inside the enclosure were recorded. This device operates independently of the data collection system and was not time synchronized. Instead, the device's internal clock was manually set at the beginning of the field test and allowed to free run. The HOBO logs data to its own internal drive. There it is stored until the gimbal enclosure could be opened, and the device removed. Once it was accessible, the HOBO was connected to a laptop, the data was downloaded, and the HOBO memory was cleared.

### *Command and Data Storage Unit*

The CDSU is a computer running the open-source, real time RTAI-LINUX operating system. It was the command and data collection unit for the trajectory reconstruction system. It directly interfaced with the GPS receiver, the IMU, both gimbal cameras, and the orientation sensor. In addition to these devices, the CDSU also received a synchronization pulse from the LIDAR data collection system. The purpose of this pulse was to inform the CDSU when the Flash LIDAR collected a data frame. When the CDSU collected a data sample from a device, it would generate the time tag for that sample.

## **3. DATA COLLECTION**

During FT1, different systems independently collect flight data. As noted above, the CDSU collected the reference system data. A separate computer operating LabView®

collected the flash LIDAR and Passive TRN data. In addition to flight data collection systems, other support-system data was recorded. The environmental sensor recorded the temperature and humidity within the gimbal enclosure. The GPS base station receivers recorded their position to act as a local, stationary GPS reference. The flight GPS receiver also recorded the flight data to its internal memory adding redundancy to GPS system. Lastly, a voice recorder captured all communications between the flight operator, the pilot and the ground crew.

At the end of each flight, data was downloaded from each system as illustrated by

. Initially, the CDSU-collected GPS data was parsed into navigation data and timing data. Then, all CDSU flight data was copied to an external, two-terabyte drive. Simultaneously, LaRC personnel retrieved a removable hard drive containing the flash LIDAR and passive optical TRN data. The data was copied into a data storage server and to the external two-terabyte drive. Between flights, the GPS base station receiver and flight receiver data was converted to a Rinex file and saved to a laptop running Ashtech® Micro-Manager. Once the receiver data was successfully downloaded and checked for quality, the receivers were erased in preparation for the next flight. This data was then copied from the laptop to the two-terabyte drive. Overnight, the data on the external drive was duplicated onto a second external drive to prevent data loss.

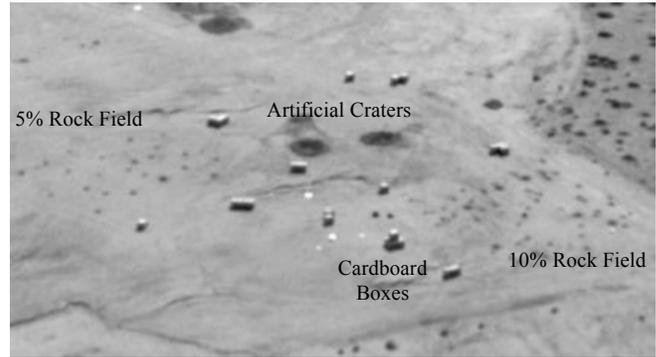
At the end of the field test, all data from all sources were copied from the external drives onto a JPL server.

#### 4. TARGET SITES

During FT1, data was collected at three target sites designated as the Borrow Pit, Lakebed and Mars Hill. The Borrow Pit and Lakebed sites were artificially created and surveyed for this experiment. The Mars Hill site was a

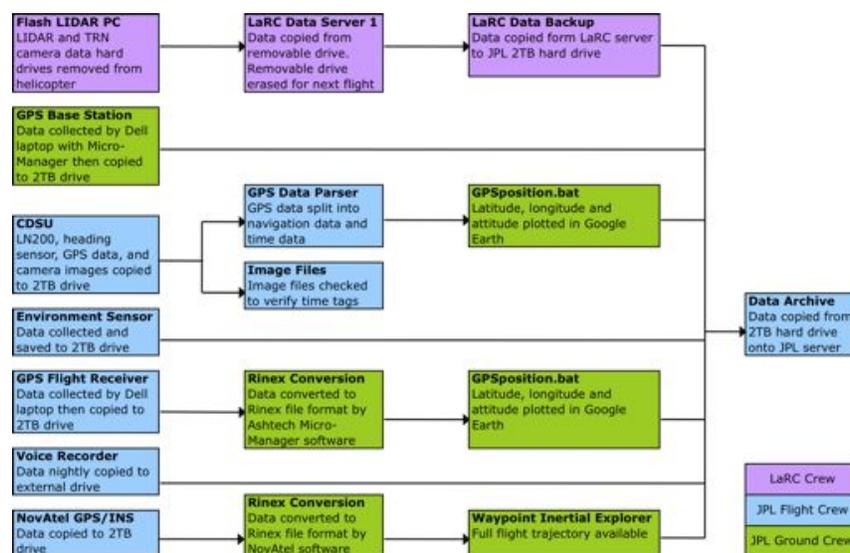
natural terrain sight selected for its varying slopes and rock collections.

#### Borrow Pit

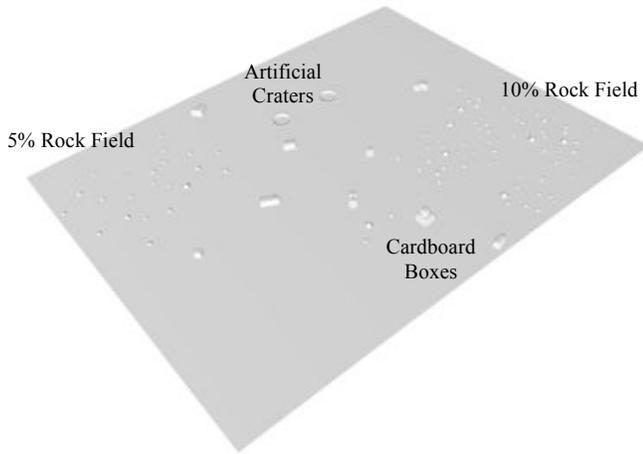


**Figure 7: Image of Borrow Pit captured by the wide-angle camera**

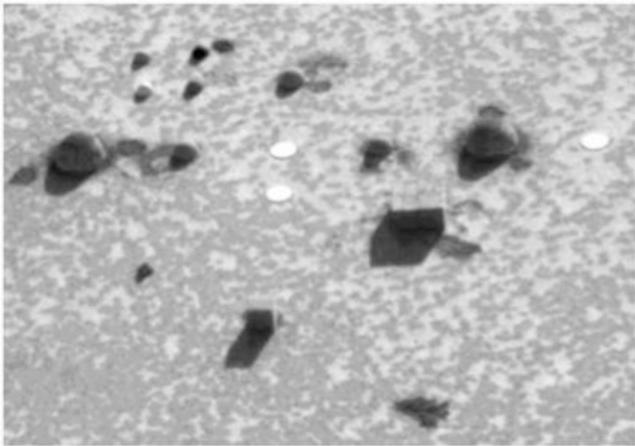
An artificially created target site located at NASA Dryden Flight Research Center, Borrow Pit seen in Figure 7 was constructed to test the flash LIDAR’s ability to detect hazards at various slant ranges and angles. To this end, several targets were located at the Borrow Pit. The targets included polystyrene and acrylic hemispheres, one-meter cardboard cubes, man-made craters, and plywood camera attitude targets. The hemispherical targets were placed into two areas: the 10% rock field and the 5% rock field. The 10% rock field contained forty-two 18” diameter hemispheres, eight 24” diameter hemispheres, and six 36” diameter hemispheres. The 5% rock field contained fifteen 18” diameter hemispheres, seven 24” diameter hemispheres, and one 36” diameter hemisphere. The locations of the hemispheres were surveyed so that the positions were known in the GPS frame. Using this information, the truth DEM (Digital Elevation Model) seen in Figure 8 was created. Also seen here are the man-made craters and the cardboard boxes.



**Figure 6: Data collection block diagram**



**Figure 8: Borrow Pit truth DEM**



**Figure 9: Image of the Lakebed captured by the thirty-degree fov camera**



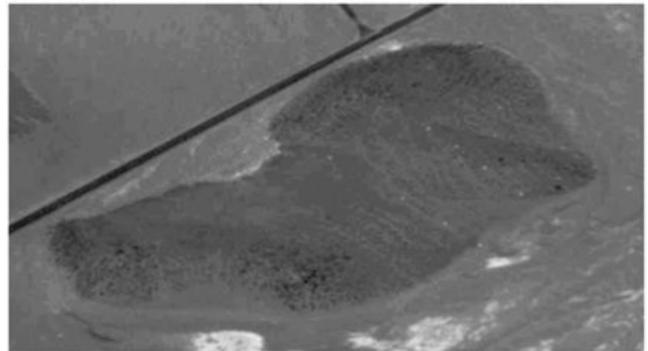
**Figure 10: Lakebed DEM**

### *Lakebed*

The Lakebed is the second man-made target site and is also located at NASA Dryden Flight Research Center. With its flat, featureless surface, the naturally occurring dry lakebed is an ideal site for determining the LIDAR's ability to detect targets. To this end, nine hemispheres of various sizes and albedos, a single one-meter cardboard cube, a cluster of four one-meter cardboard cubes and eleven attitude targets were placed there. The site was surveyed and a truth DEM was created. Figure 9 and Figure 10 show the Lakebed site as seen by the thirty-degree field-of-view camera and the truth DEM, respectively.

### *Mars Hill*

Seen in Figure 11, Mars Hill is the only naturally occurring target sight of FT1. Located within Death Valley National Park, this site is approximately 200 miles away from the Lakebed and Borrow Pit. With its natural rocks, slopes and lack of vegetation, this site is a perfect analog for lunar terrain. Plywood attitude targets, seen as white dots on the right of the hill, and cardboard boxes were placed at Mars Hill.



**Figure 11: Mars Hill captured by the thirty-degree fov camera**

## **5. TRAJECTORY RECONSTRUCTION**

### *Overview*

For FT1, the objective of the trajectory reconstruction process is to provide the position and attitude of the LIDAR for every LIDAR sample relative to a target. Initially, it was proposed to develop an Extended Kalman Filter (EKF) to produce the LIDAR attitude and position by combining the GPS position data, the IMU angular rate data and the camera-based attitude estimates. However, the development of the filter stalled. In order to provide results in a more timely fashion, a simpler method was adopted. The new process is as follows.

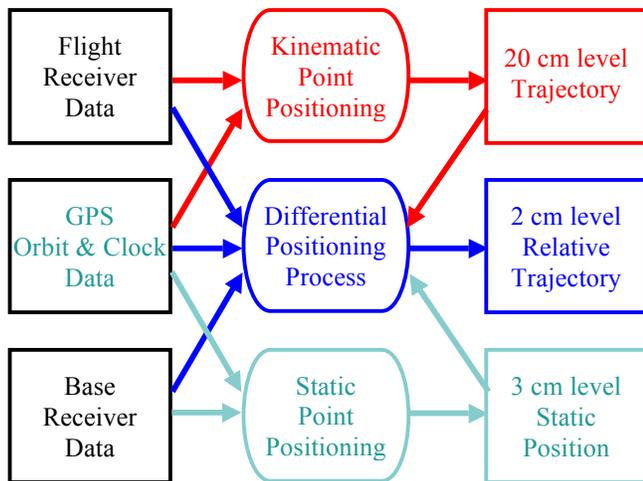
To begin, the raw GPS data is processed to produce the position measurements accurate to a standard deviation of two centimeters. Second, the images produced by the wide-angle camera are manually sorted through to find images

that contain six or more attitude targets. Once the images were identified, they are used to determine the camera attitude. Because the cameras infrequently captured images of the attitude targets, the camera-generated attitudes are used to initialize IMU-based attitude propagation. Finally, gyro propagation is combined with the GPS measurement to provide the full LIDAR pose (position and attitude).

This section details the reconstruction process. Initially, we present the GPS processing methodology followed by the imagery-based attitude estimation process. After that, the equations used to propagate the attitude estimate based on the IMU data are given. Finally, we provide how the camera estimates, gyroscope propagation and GPS position data are all combined to reproduce the LIDAR pose.

### GPS Data Processing

The GPS hardware configuration of FT1 consists of two receivers, one is a static base station on the ground, and the other mounted on the helicopter. The data recorded by both receivers is analyzed to derive the precise positions of each. To obtain the precise position of a receiver, the data logged by the receiver is processed together with satellite orbit position and clock information. The quality of the GPS orbit and clock information used in the data processing defines the accuracy of the receiver position solution. In the final processing of the GPS data for FT1, we use the precise GPS orbit and clock information called the FLINN product. This product is produced at JPL by processing tracking data from 80 globally distributed ground stations, with about 10 days latency and 3 cm accuracy. The product is routinely generated and submitted to International GNSS Service (IGS) to support precise applications in science and industry communities.



**Figure 12: GPS data post-processing for trajectory reconstruction**

Our solution approach is illustrated in Figure 12. At the left of the image is the data collected during a flight. The center column represents a GPS processing method and the right

hand side is the product and its solution accuracy. Each method is now described in detail.

*Base Station Static Point Positioning*—The base station receiver data is processed using the FLINN GPS orbit and clock product. The combination of GPS measurements at two different frequencies, commonly known as L1 and L2, are used to remove the effect of the ionosphere on the measurements. To further improve accuracy, data points with tracking elevation angle below seven degrees are excluded. Given two hours or longer hours of data, the typical accuracy of the static point position solution is three centimeters.

*Helicopter Kinematic Point Positioning*—Initially, the flight receiver data is processed to determine the helicopter’s position using a kinematic point positioning technique. Again, L1 and L2 GPS measurements are used to remove the effect of the ionosphere. Then, using the best FLINN GPS orbit and clock information, the absolute the kinematic point positioning is determined to within twenty centimeters. Here, the main error source for the solution is the correlation between the troposphere delay and the position height component. This solution is used as the initial trajectory for the relative positioning solution.

*Helicopter Kinematic Relative Positioning*—The base station receiver data and the flight receiver data are processed together to determine the helicopter’s position relative to the base station. In the relative positioning, those errors that are common to both the base station receiver data and the flight receiver data, such as the GPS orbit and clock errors and troposphere and ionosphere delay errors, cancel out over a short baseline. In our process, the base station position is fixed to the static point positioning solution and the helicopter’s position is solved iteratively starting with the kinematic point positioning solution. To do so, we use the GPS pseudorange and carrier phase measurement at frequency L1. Typically the helicopter’s position error relative to the base station has a two-centimeter standard deviation.

### Image Based Attitude Estimation

The objective of the image-based attitude estimation is to provide attitude measurements while over a test site. Because the image-based attitude estimation does not drift as the IMU does, it serves as an anchor point for the IMU data processing and the final trajectory reconstruction.

The initially proposed procedure for the image-based attitude estimation was the following. First, the image sequences containing a sufficient number of attitude targets were manually extracted. Then, using these images, the image pose was estimated. Last, we applied a bundle adjustment using the estimated pose and selected track features of an image series to improve attitude and position estimates. However, many image sequences contained a lack of distinguishable features, or the target field was

relatively small compared to the full image. For these, bundle adjustment produced bad results. Therefore, a method was adopted that relied on the kinematic relative GPS position rather than estimating it from imagery. This method is as follows.

*Image Rectification and Attitude Target Extraction*—Initially, sequences of images containing six or more attitude targets were found. Prior to inspection, we used the camera CAHVOR model [4] to rectify each image. Then, a manual scan was made of every eighth frame. If any attitude targets were seen, the skipped frames were examined to determine the exact frame in which the targets initially appeared. By comparing this frame to the surveyed locations of the attitude targets, each target was identified and its x- and y-pixel location were recorded. These locations are then used to initialize an automatic tracking algorithm.

To track the target locations, each image frame in a sequence containing the attitude targets was passed through a Scale-Invariant Feature Transform (SIFT) Keypoint Detector [7]. Comparing the SIFT points between consecutive images allowed for the computation of the Homography transform which mapped the pixels from one image into the next. By initializing the locations of the targets in the first frame with a location recorded earlier, the Homography transform is applied to propagate the x- and y-pixel locations forward through the entire series. When, however, there was poor matching between pictures, the propagation failed.

When the above failed, the algorithm automatically shifted the target grid of one image horizontally and vertically over the other. For every shift, the normalized sum of the pixel values under the targets is computed and the optimal shift was found. Since the targets are white, this is the shift at which the sum of pixels is the maximum. In order for automatic shifting to succeed, the targets must clearly contrast the background and the target grids in both frames must be relatively in the same configuration. When either condition was not met, the automatic shift algorithm failed. In this case, a target was hand-selected and the shift from one image to the next was manually determined. All targets were shifted by this amount and the x- and y-pixel location propagation continued. Starting with the manual shift, the local maximum was found to find the best local target location. When the propagation of the target locations completed, the results were visually inspected for correctness by plotting the predicted attitude target locations back onto the image frames. Figure 13 illustrates the entire target-tracking algorithm.

*Solving for the Camera Attitude*—Once a series of rectified images are produced and the pixel locations of the targets are known, the attitude of the camera can be found. For a CAHVOR calibrated camera, every x- and y-pixel location corresponds to a known three-dimensional ray expressed in the camera frame. Given the set of three-dimensional points

$P_i$  expressed in the GPS frame and a corresponding set of two-dimensional points  $v_i$  expressed in image coordinates (the x,y-pixel locations), it is possible to derive a unit length ray  $v_i$  in the camera frame using the camera calibration data. Let  $r$  be the camera location as measured by the GPS and  $\{w_i\} = \{w_1, \dots, w_n\}$  be a set of unit length rays from the camera to each  $P_i$ , i.e.,

$$w_i = \frac{(P_i - r)}{\|P_i - r\|} \quad (1)$$

where  $\| \cdot \|$  is the vector norm. Then, solving for the full camera pose amounts to computing the rotation between  $\{w_i\}$  and set  $\{v_i\} = \{v_1, \dots, v_n\}$ . To do so, we define  $V$  as the  $n \times 3$  matrix consisting of  $\{v_i\}$ ,  $W$  as the  $n \times 3$  matrix consisting of  $\{w_i\}$  and  $R$  as the rotation matrix that rotates  $V^T$  such that

$$W^T = RV^T \quad (2)$$

In practice, (2) does not have an exact solution and is computed either by  $R = AB^T$ , where  $A \Sigma B^T = W^T V$  is the singular value decomposition of  $W^T V$ , or by  $R = \text{Pinv}(W)V$  where  $\text{Pinv}$  is the Moore-Penrose pseudoinverse. Note that if  $W^T W$  is non-singular, then  $\text{Pinv}(W^T W) = (W^T W)^{-1} W^T$  and, therefore,  $R = (W^T W)^{-1} W^T V$ .

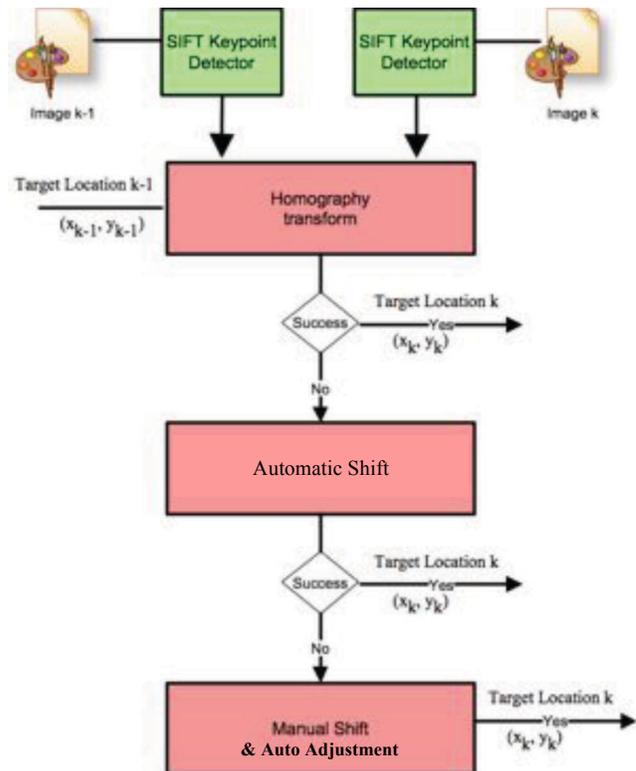


Figure 13: Block diagram of target tracking algorithm

$${}^M q_{k+1}^I \approx \left( 1_4 + \frac{1}{2} \frac{T\bar{\Omega}}{1!} + \frac{1}{4} \frac{(T\bar{\Omega})^2}{2!} + \frac{1}{8} \frac{(T\bar{\Omega})^3}{3!} \right) {}^M q_k^I \quad (4)$$

Next, we find the best attitude estimate using the RANdom SAMple Consensus (RANSAC) method [3]. From the set  $\{P_i\} = \{P_1, \dots, P_n\}$ , three points are selected and used to form an attitude estimate. Using the resulting R matrix, the attitude targets are projected into the image and the projection error is calculated. Using a predetermined threshold, erroneous members of  $\{P_i\}$  are identified and rejected. Then, the best attitude estimate is that formed from the remaining points. The resulting matrix R is the image-based attitude estimate.

Here,

$$T\bar{\Omega} = \begin{bmatrix} 0 & \theta_z & -\theta_y & \theta_x \\ -\theta_z & 0 & \theta_x & \theta_y \\ \theta_y & -\theta_x & 0 & \theta_z \\ -\theta_x & -\theta_y & -\theta_z & 0 \end{bmatrix}$$

and  $1_4$  is the  $4 \times 4$  identity matrix.

The angles  $\theta_x$ ,  $\theta_y$ , and  $\theta_z$ , are the angles provided by the x-, y- and z-gyroscope measurements, respectively. This equation is the result of truncating the infinite series expansion of the quaternion kinematic equation [8].

### IMU-Based Attitude Propagation

For a given flight, the resulting imagery-based attitude estimates exist only when the wide-angle camera captured a sufficient number of attitude targets, which occurred at a maximum rate of once a second. Given the high rotational velocities of the gimbal, this was insufficient for reconstructing the LIDAR trajectory. To this end, we determine the attitude every 2.5 milliseconds using the LN200 gyroscopes. To use the gyroscopes, we first define an inertial frame. Then, using the camera attitude estimates to initialize the attitude, we propagated the attitude over the time that the LIDAR was sampling data. The following sections provide the details of each step.

*The inertial frame*—We utilize an Earth Centered Inertial (ECI) frame. We define the inertial frame to be the location of the WGS84 frame at the initial time in a flight that the wide-angle camera captures enough attitude targets to produce an attitude estimate.

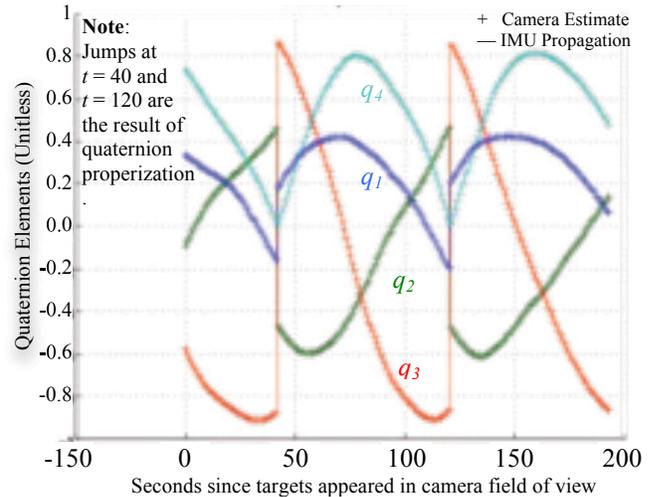
*Determining the initial attitude from a series of attitude estimates*—In this section, we find the initial attitude of the IMU relative to our ECI frame. The attitude of the IMU relative to the inertial frame is denoted  ${}^M R^I$ . This was determined for every imagery-based attitude estimate. The product of the image-based estimates was the attitude of the camera with respect to the GPS frame. Call it  ${}^E R^C$ . From the  $i^{\text{th}}$  estimate  ${}^E R^C_i$ , the IMU attitude relative to the initial frame was determined by the successive rotations

$${}^M R^I_i = {}^M R^C ({}^E R^C_i)^{-1} {}^E R^I(t). \quad (3)$$

Here,  ${}^M R^C$  is the known rotation between IMU frame and the camera frame, and  ${}^E R^I$  accounts for the rotation between the GPS frame and the ECI frame.

For propagating attitudes, quaternions are numerically better suited. Therefore, each  ${}^M R^I_i$  is converted to the equivalent quaternion  ${}^M q_i^I$ . Now that we expressed the attitude of the IMU relative to an inertial frame in quaternion form, we propagated it using the delta angles measured by the gyroscopes and using the (4).

Recall that image-based attitude estimates are only available when the camera images at least six attitude targets. Also, the estimates come from sequences of consecutive images. For each sequence, we attempt to find a good initial attitude to begin our propagation of the attitude over the entire flight. To start, we selected the initial estimate from a sequence and propagated the attitude forward until the end of it. Note that each sequence was a small portion of the entire flight. Figure 14 shows an example of  ${}^M q^I$  generated from a section of data collected over the Borrow Pit target site. As seen here, the image-based attitude estimates agree with the gyroscope propagation.



**Figure 14: Plot of attitude estimate based on camera versus attitude propagated with IMU**

Having both a propagated quaternion and camera-based estimates during the same time span, we found the mean error between them. For clarity, we now denote the propagated quaternion as  ${}^M p^I$  and retain the  ${}^M q^I$  notation for attitude estimates. The error quaternion mapping the

difference between the propagated attitude and the estimated attitude is given by

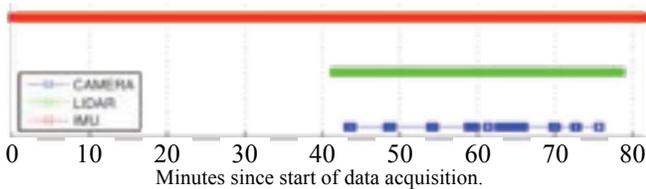
$$q_e = ({}^M p^I)^{-1} \otimes {}^M q^I. \quad (5)$$

Here,  $({}^M p^I)^{-1}$  is the inverse quaternion of  ${}^M p^I$  and  $\otimes$  is the quaternion product. From the error quaternion, the equivalent Euler angle of rotation  $\phi$  and unit Euler axis of rotation  $\underline{a}$  can be determined. Using  $\phi$  and  $\underline{a}$ , we define the scalar error quantity  $e_s = \phi$  and the vector error  $\underline{e}_v = e_s \underline{a}$ .

At this point, we have an error vector  $\underline{e}_v$  and scalar  $e_s$  for every attitude estimate. Taking  $\underline{u}_e$  to be the mean of all  $\underline{e}_v$  and  $u_s$  to be the mean off all  $e_s$ , we create the correction quaternion  $q_c = [\underline{u}_e^T \sin(u_s/2), \cos(u_s/2)]^T$ . Applying the correction quaternion to first camera-based estimate  ${}^M q^I_1$  provided the initial quaternion used to propagate the attitude for the entire flight, such that  ${}^M q^I_0 = {}^M q^I_1 \otimes q_c$ .

To illustrate the effects of this procedure, we include Figure 16, which, because of its size, has been placed in the Appendix. Here, frames (A) and (B) are the values of  $e_s$  and  $e_v$ , respectively, created when the attitude was propagated starting at  ${}^M q^I_1$  and (C) and (D) was propagated starting at  ${}^M q^I_0$ . By improving the starting point, the maximum scalar error  $e_s$  has dropped from 0.26 degrees to 0.18 degrees and the standard deviation has dropped from 0.06 degrees to 0.03 degrees. The vector error has the same shape, but now has zero mean.

*Determining the attitude of the IMU with respect to ECI—* During a flight, several passes at the target site are made. At each pass, the wide-angle camera takes a series of images used to estimate the attitude. See Figure 15. This figure shows the IMU, LIDAR and camera timestamps for the second Borrow Pit flight. Here nine sequences of attitude estimates are clearly visible. However, it can be seen that LIDAR data is available between each sequence. To provide attitude estimates between each series, we propagate the estimates using the IMU. As the propagation continues, the estimate accuracy degrades due to error sources such as the gyroscope bias and numerical inaccuracies. Here we attempt to compensate for these errors.



**Figure 15: Sample times of IMU, LIDAR and camera**

Initially, it was proposed to combine GPS, camera data and the IMU with an EKF. However, development of the filter stalled and remains the subject of future work. For FT1, we utilized the following ad hoc method. We began by finding the initial quaternion for each sequence of image-based

attitude estimates as determined by the previous section. In addition to  $q_0$ , the standard deviation of  $e_s$  for each sequence was calculated. Starting from the initial quaternion of each sequence, we propagate the attitude across the entire flight. While propagating, the accuracy of the attitude was also tracked via

$$\sigma_k(t) = \sigma_0 + |b(t_k - t_{q_0})| \quad (6)$$

where  $\sigma_0$  is the standard deviation of  $e_s$  for a given sequence,  $|\cdot|$  indicates the absolute value,  $b$  is an arbitrary scale factor,  $t_k$  is the timestamp of the  $k^{\text{th}}$  measurement and  $t_{q_0}$  is the timestamp of  $q_0$ . Take  $M$  to be the number of imaged-based attitude estimate. For the flight illustrated in Figure 15,  $M = 9$ . Since we propagated the attitude starting at each sequence's  $q_0$ , there were  $M$  different propagated attitudes at every time step. Each of these were combined using a weighted average

$$\underline{u} = \sum_{j=1}^M \frac{w_j \phi_j \underline{a}_j}{W} \quad (7)$$

where  $\phi_j$  and  $\underline{a}_j$  are the Euler angle and axis that define a propagated quaternion  $p_j$ ,  $w_j = \sigma_j^{-1}$  is the weight of  $p_j$  and  $W$  is the sum of all weights. Notice, that by using  $\phi$  and  $\underline{a}$ ,  $\underline{u}$  is the average of the attitudes axis angle form. Finally, for every time step,  $\underline{u}$  was converted to the quaternion  ${}^M q^I(\underline{u}) = [\underline{v}^T \sin(\theta), \cos(\theta)]^T$ , where  $\underline{v} = \underline{u}/\|\underline{u}\|$  and  $\theta = \|\underline{u}\|/2$ . Also note that this scheme favors samples where there is good camera and IMU agreement.

#### *Constructing the LIDAR trajectory*

The objective of the trajectory reconstruction effort was to determine the position and attitude of the LIDAR with respect to a target site. Finding the attitude requires the following successive rotations:

$${}^T q^L = {}^T q^E \otimes {}^E q^I(t) \otimes ({}^M q^I)^{-1} \otimes {}^M q^L. \quad (8)$$

Above,  ${}^T q^E$  and  ${}^M q^L$  are the known relative attitudes. The first is the attitude of the target frame relative to the GPS frame and the latter is the attitude of the LIDAR with respect to the IMU frame. The other term introduced in (8),  ${}^E q^I(t)$ , is the attitude of the GPS frame with respect to the ECI frame at time  $t$ .

To find the position of the LIDAR in the target frame, we first linearly interpolated the 1 Hz GPS position measurement,  $r_{L/E}$ , to the 8 Hz necessary for the LIDAR. Then, we removed the offset between the GPS and target frame,  $r_{E/T}$ , and rotated the interpolated measurements from the GPS frame into the target frame using the rotation matrix  ${}^T R^E$ , which is the rotation matrix equivalent to (8),

$$r^T = {}^T R^E (r_{L/E} - r_{E/T}) \quad (9)$$

The attitude provided by (8) and the GPS measured position via (9) gave the full pose of the LIDAR for each flight.

## 6. INSTRUMENT VERIFICATION

To verify the accuracy of the reconstructed trajectory, we used the LIDAR by shifting the LIDAR data until the minimum correlation error between the sensor data and truth DEM was found. The LIDAR camera boresite vector corresponds to the LIDAR z-axis and the sensor array rows and columns define the x- and y-axes. Therefore, any position shift in the z-direction can be interpreted as a range error. Shifts in the x- and y-axes correspond to yaw,  $\psi$ , and pitch,  $\rho$ , errors as follows:

$$\rho = \tan^{-1}(-\Delta_y (d + \Delta_z)^{-1}), \quad (10)$$

$$\psi = \tan^{-1}(\Delta_x (d + \Delta_z)^{-1}), \quad (11)$$

where  $\Delta_x$ ,  $\Delta_y$  and  $\Delta_z$  represent the x, y and z shifts, and  $d$  is the distance to the target. Notice that yaw is the rotation about the LIDAR y-axis, and pitch is the rotation about the LIDAR x-axis. Roll and pitch errors have been calculated using this method for Lakebed and Borrow Pit flights and plotted in Figure 17 (A) and Figure 18 (A), respectively.

Additionally, we compared the attitude determined by (7) to the image-based estimates and plotted the results in Figure 17 (B) and Figure 18 (B). The maximum observed errors observed by the LIDAR during the Borrow Pit flight were 0.25 and 0.16 degrees for pitch and yaw, respectively. These are comparable to the camera-generated values of 0.12 and 0.15 degrees (see Figure 17 (B)). For the Lakebed flight, seen in Figure 18 (A), the maximum LIDAR determined error was 0.37 and 0.58 degrees for pitch and yaw, respectively. However, the pitch and yaw errors seen in Figure 18 (B) for the cameras were significantly worse, 0.33 and 2.50 degrees, respectively.

In Figure 17 (B) and Figure 18 (B), a non-random process noise is apparent and is the likely cause of the larger than expected errors. Currently, the source of this error has not been identified, however, for the ad hoc method described here, gyroscope scale factor errors, sense axis misalignments, and bias have not been accounted for. Additional error sources for the camera and IMU include the numerical inaccuracies introduced by (4) and timing errors. For the LIDAR, the largest source of error is the timing uncertainty. Efforts have been made to identify the timing of the LIDAR sampling. Unfortunately, it changes from flight to flight and is not easily identifiable.

Despite the larger than expected error, the trajectory reconstruction method was sufficient enough to place the LIDAR hazards within one meter horizontally of the truth hazards. The combination of the mainly horizontal trajectory errors and the time-varying range bias inherent to the LIDAR caused the LIDAR data to be misaligned relative to the truth DEM. To eliminate this misalignment, the flash LIDAR data and truth DEM were correlated using a procedure based on the HRN algorithm [6] was used. The

end result was precise alignment to 1 DEM pixel (0.1m) or less.

Even though the results from the method presented here were sufficient to achieve the ALHAT objectives, it may be desirable in the future to try to achieve the expected attitude performance of 0.71-degrees. In that case, methods to identify the error sources will be developed. One such option is the use of an unconstrained nonlinear optimization, such as the MATLAB® function `fminsearch`. By formulating the effects that gyroscope misalignments, biases and timing offsets has on the propagated trajectory, the parameter values that minimize error between the camera estimates and the propagated trajectories can be found. Although this approach is simple, the time for the nonlinear optimizer to find the parameters will be great, and it will not compensate for the numerical inaccuracies.

A more complex alternative is to complete the EKF and extend it to an optimal smoother. The filter will directly solve for gyroscope biases and drift, due to the averaging between sensors, and reduce the impact of the numerical. By running the filter multiple times and observing the effects that shifting timing offsets has on the filter residuals, the best timing offset of the camera and IMU can be identified. Similarly by shifting the timing of the LIDAR samples and observing the effects on  $\Delta_x$ ,  $\Delta_y$  and  $\Delta_z$ , the best LIDAR offset for each flight can be found.

## 7. CONCLUSIONS

The expected performance of the reference system was .12 cm for position and 0.071 degrees for attitude. Although the maximum observed attitude error of 0.33 degrees was five times worse than expected, the error for the Borrow Pit flight never exceeded the FT1 requirement. However, this was not the case for the Lakebed flight. Examination of the error seen in Figure 17 (B) and Figure 18 (B) clearly demonstrates a yet unidentified, non-random process noise. By compensating for gyroscope misalignments, scale factor errors, gyro biases and timing, this error should dramatically reduce. Despite the larger than expected errors, this method was sufficient to achieve the FT1 ALHAT objectives. In case a better level of accuracy is required in the future, the sources of error will be analyzed and the methods to compensate for them will be developed. Two approaches are the use of an unconstrained nonlinear optimizer, or the development of an Extended Kalman Filter. The results of these efforts are expected to achieve the expected 0.071-degree performance for FT1.

## ACKNOWLEDGMENTS

We thank Asif Ahmed of the Jet Propulsion Laboratory for his guidance, patience and tutelage during the trajectory reconstruction effort.

The work described in this publication was performed at the Jet Propulsion Laboratory, California Institute of

Technology, under contract from the National Aeronautics and Space Administration. The work was funded by the NASA Exploration Technology Development Program and would not have been possible without the flash LIDAR sensor provide by NASA Langley and the field test system provided by JPL.

## REFERENCES

- [1] Bulyshev, Alexander; Pierrottet, Diego; Amzajerdian, Farzin; Busch, George; Vanek, Michael; Reisse, Robert "Processing of three-dimensional flash lidar terrain images generating from an airborne platform" roc. SPIE, Vol. 7329 April 2009.
- [2] Epp, Chiroid and Tom Smith, "Autonomous Precision Landing and Hazard Detection and Avoidance Technology (ALHAT)," Proc. IEEE Aerospace Conf., Big Sky, MT, March 2007.
- [3] Fischler, Martin A. and Robert C. Bolles. "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography". Comm. Of the ACM 24, June 1981: 381-395
- [4] Gennery, D.B. "Least-Squares Camera Calibration Including Lens Distortion and Automatic Editing of Calibration Points," Workshop Calibration and Orientation of Cameras in Computer Vision, XVII Congress of the International Society of Photogrammetry and Remote Sensing, Washington, DC, August 2, 1992.
- [5] Johnson, Andrew E., Jason A. Keim and Tonislav Ivanov. "Analysis of Flash Lidar Field Test Data for Safe Lunar Landing". IEEE Aerospace Conference, March 6-13, 2010, Big Sky, Montana. (expected)
- [6] Johnson, Andrew E. and Miguel SanMartin, "Motion Estimation from Laser Ranging for Autonomous Comet Landing," Proc. Int'l Conf. Robotics and Automation, pp. 132-138, April 2000.
- [7] Lowe, David G. "Distinctive image features from scale-invariant keypoints." International Journal of Computer Vision, vol. 60.2 (2004): 91-110.
- [8] Wertz, James R. Spacecraft Attitude Determination and Control. Boston: Kluwer Academic Publishers, 1978: 558-566.
- [9] Melvin (Jay) White II, Tom Criss, and Dewey Adams. "APLNav Terrain Relative Navigation Helicopter Field Testing", AIAA Guidance, Navigation, and Control Conference, August 2009, Chicago, Illinois

## BIOGRAPHY



**Jason Keim** received his B.S.B.M.E. from the University of Southern California and M.S.M.E. from California State University, Los Angeles. Since 2002, Jason has been a member of the Guidance and Control Analysis Group at the NASA Jet Propulsion Laboratory. His primary focus has been the development and validation of formation flight algorithms and technologies for missions such as NASA Starlight, TPF and DARPA F6. Additionally, he has contributed to the autonomous surface operations of the Mars Science Laboratory, data processing and trajectory reconstruction for the Autonomous Landing and Hazard Avoidance Technology program, and other research and technology development programs.



**Dr. Sohrab Mobasser** is a Senior Member of the Engineering Staff at the National Aeronautics and Space Administration's (NASA) Jet Propulsion Laboratory (JPL). Sohrab has more than 26 years of aerospace industry experience, most of it in spacecraft attitude determination. His work can be found on many planetary missions, from the Galileo mission to Jupiter to the successful Pathfinder mission to Mars and the Cassini mission to Saturn. His current interests are new technology and applications for autonomous spacecraft attitude determination. Sohrab is the Field Test Lead for ALHAT project.



**Dr. Yang Cheng** Dr Yang Cheng is senior staff member at JPL and he has been involved in many NASA and reimbursable robotic projects for many years. He is the key algorithm developer for the Descent Image Motion Estimation System (DIMES), which played a critical role for Mars Exploration Rovers landing safely on Mars. He was also the key software developer for the MER onboard visual odometry, which has been used during critical rover maneuvers and traverses. In addition, he was also involved in novel vision algorithm development for future Mars and Lunar safe and pinpoint landing



**Tonislav Ivanov** is an Associate Member of Technical Staff in the Computer Vision Group at JPL. He works on lidar data processing, field test setup, and is helping develop the hazard detection algorithm for the Autonomous Landing and Hazard Avoidance Project. He also works on recognition using stereo data for robot human awareness and lunar terrain characterization for future missions to the Moon.



**Dr. Da Kuang** received his Ph.D. in Aerospace Engineering from The University of Texas at Austin in 1995. He joined the Orbiter and Radio Metric Systems Group at JPL in 1996. His work has been focused on analyzing GPS and GPS-like tracking data for precise orbit determination and precise relative positioning.



**Dr. Andrew Johnson** is a Principal Member of Technical Staff in the Optical Navigation Group at JPL. He is the JPL Project Manager and terrain sensing algorithm lead for the Autonomous Landing and Hazard Avoidance Project which is developing technology for safe and precise landing for the next generation manned lunar lander. At JPL he works on development, validation and flight implementation of computer vision systems for planetary landers and Mars rovers.



**Hannah R. Goldberg** received her M.S.E.E and B.S.E from the Department of Electrical Engineering and Computer Science at the University of Michigan in 2004 and 2003 respectively. She has been employed at the Jet Propulsion Laboratory, California Institute of Technology since 2004 as a member of the technical staff in the Precision Motion Control and Celestial Sensors group. Her research interests include the development of nano-class spacecraft and microsystems.

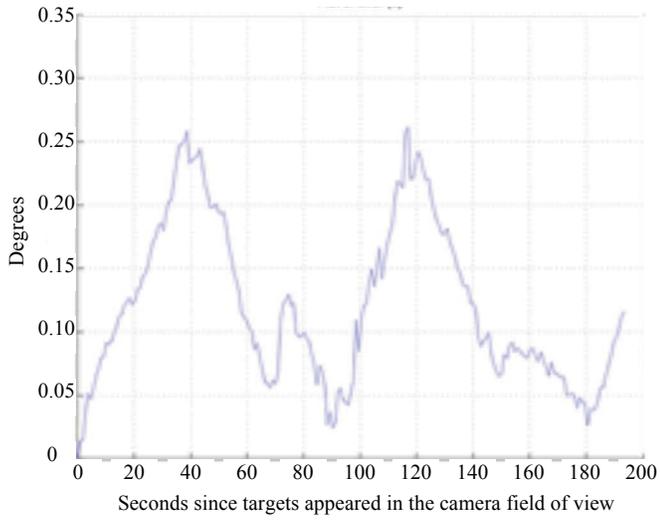


**Garen Khanoyan** is a member of the Advanced Computer Systems & Technologies group at JPL. He has been involved with field testing activities and the development of the Command & Data Storage Unit since 2008 for ALHAT and MSL projects. Garen received his B.S.E.E. and M.S.C.S. from the University of Southern California.

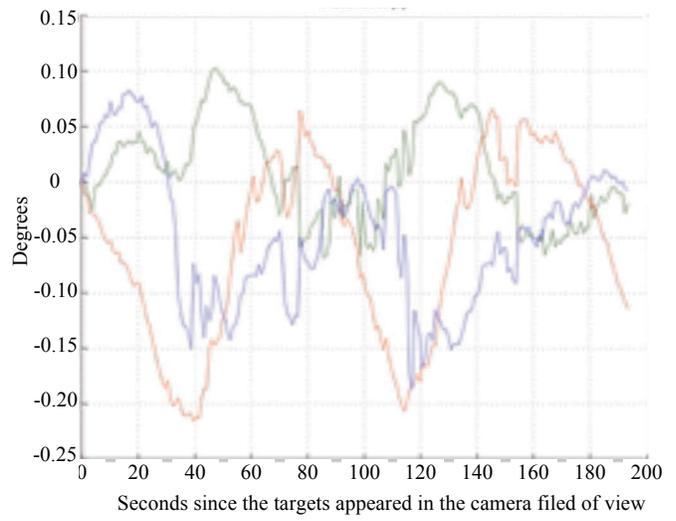


**David B. Natzic** received his B.S.C.S. from the Department of Computer Science at the University of Management and Technology. He has been employed as an Associate Member of Technical Staff in the Guidance, Navigation and Control Group at JPL. Dave Joined JPL in 1992 and currently serves as an essential ALHAT team member focusing on the design, integration and field testing of Flash LIDAR instruments onboard aerial platforms.

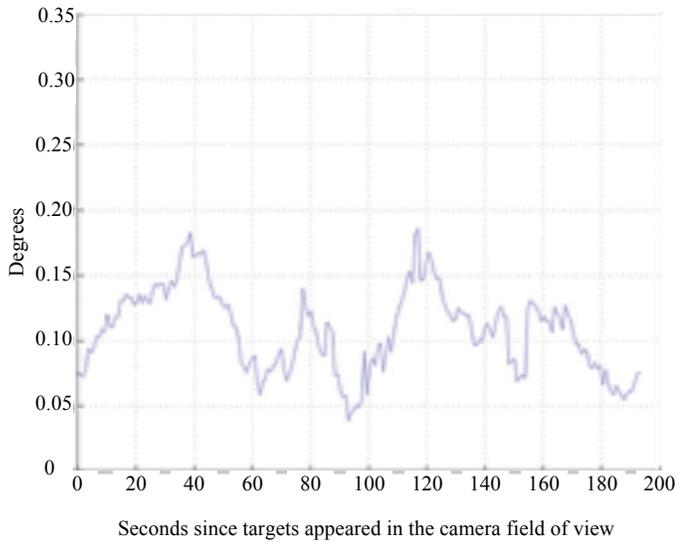
## APPENDIX



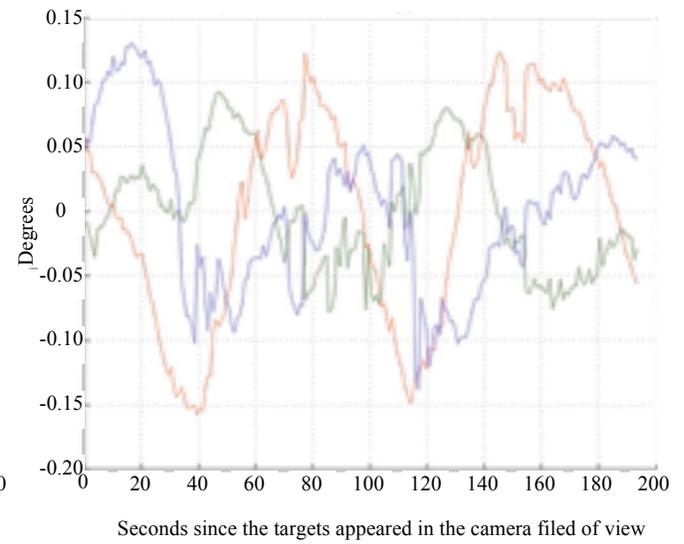
(A)



(B)

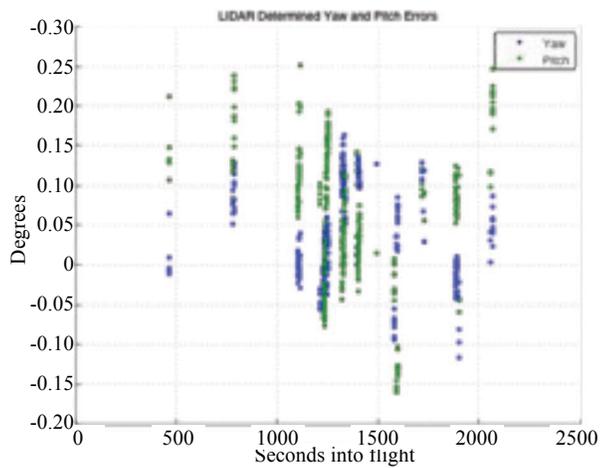


(C)

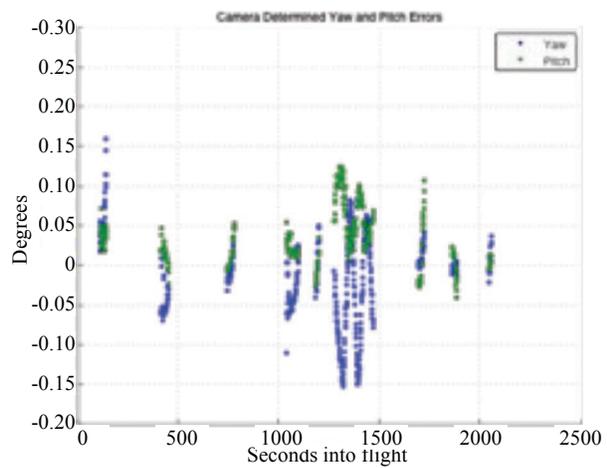


(D)

**Figure 16: The error between the propagated and estimated attitude. (A) is the scalar error and (B) is the vector error before the corrections. (C) is the scalar error and (D) is the vector error after the correction.**

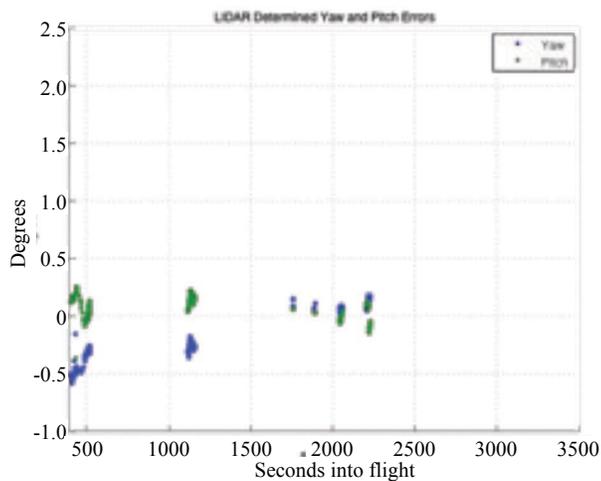


(A)

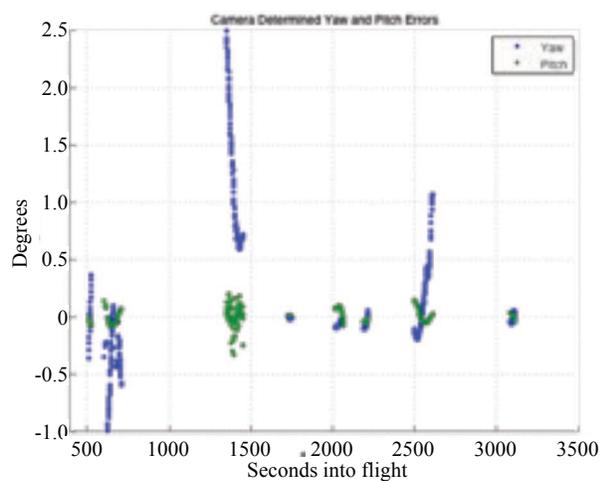


(B)

**Figure 17: Yaw and pitch errors from second Borrow Pit flight for (A) the LIDAR and (B) the camera.**



(A)



(B)

**Figure 18: Yaw and pitch errors from the third Lakebed flight for (A) the LIDAR and (B) the camera.**