An Integrated Vehicle Navigation System Utilizing Lane-Detection and Lateral Position Estimation Systems in Difficult Environments for GPS

Christopher Rose, Jordan Britt, John Allen, and David Bevly

***Abstract*—A navigation filter combines measurements from sensors currently available on vehicles—Global Positioning System (GPS), inertial measurement unit, inertial measurement unit (IMU), camera, and light detection and ranging (lidar)—for achieving lane-level positioning in environments where standalone GPS can suffer or fail. Measurements from the camera and lidar are used in two lane-detection systems, and the calculated lateral distance (to the lane markings) estimates of both lane-detection systems are compared with centimeter-level truth to show decimeter-level accuracy. The navigation filter uses the lateral distance measurements from the lidar- and camera-based systems with a known waypoint-based map to provide global measurements for use in a GPS/Inertial Navigation System (INS) system. Experimental results show that the inclusion of lateral distance measurements and a height constraint from the map creates a fully observable system even with only two satellite observations and, as such, greatly enhances the robustness of the integrated system over GPS/INS alone. Various scenarios are presented, which affect the navigation filter, including satellite geometry, number of satellites, and loss of lateral distance measurements from the camera and lidar systems.**

***Index Terms*—Camera, Global Navigation Satellite System, Global Positioning System (GPS), inertial measurement unit (IMU), Kalman filter, lane detection, light detection and ranging (lidar), outages, sensor fusion.**

I. Introduction

OVER 19 000 people died on U.S. highways due to road departures in 2008 [1]. Safety systems on vehicles, such as lane departure warning (LDW) systems, can reduce the number of these deaths by warning the driver of a lane departure. These safety systems determine the lateral distance to the lane markings and warn the driver when the vehicle has left or is about to leave the lane using a single sensor such as a camera. However, current LDW systems can fail due to varying conditions and environments. A multifaceted approach using multiple sensors provides a more robust solution when solutions are not available from a single sensor.

Manuscript received June 18, 2013; revised February 7, 2014; accepted April 14, 2014. Date of publication May 29, 2014; date of current version December 1, 2014. The Federal Highway Administration is funding this project and others across the range of issues that are critical to the transportation industry through the Exploratory Advanced Research (EAR) Program. For more information, see the EAR Web site at http://www.fhwa.dot.gov/advancedresearch/about.cfm#focus. The Associate Editor for this paper was J. V. Krogmeier.

C. Rose, J. Britt, and D. Bevly are with the GPS and Vehicle Dynamics Laboratory, Auburn University, Auburn, AL 36849 USA.

J. Allen is with Aurora Flight Sciences, Manassas, VA 20110 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TITS.2014.2321108

A navigation filter was developed to provide a more robust solution for lane-level localization even when one or more sensors fail. This positioning system uses the following four sensors: Global Positioning System (GPS), inertial measurement unit (IMU), light detection and ranging (lidar), and low-cost camera. Commercial automotive systems typically have these sensors already on the vehicle and employ them for applications of navigation and safety, such as LDW with cameras, map-aided navigation with GPS, stability control with an IMU, and active cruise control with a lidar. In addition, the use of precision survey maps of lane markings is increasingly available from companies such as Nokia and Google. A combination of all of these available sensors in a positioning system can result in a more robust solution for applications such as vehicle safety.

Many lidar systems used for LDW systems employ a host of different techniques for detecting and tracking lane markings. In [2], a histogram and a histogram gradient feature extraction algorithm are used to detect lane markings with a modified lidar in which two mirrors were added on either of its sides, which reflects the lidar pulses that would normally be reserved for scanning behind the lidar. Solid lane markings were detected about 95% of the time, and dashed lane markings were detected about 35% of the time, but lateral distance in the lane was not reported. The histogram feature detection approach is again explored in [3], where a region of interest is defined at a distance of 0–30 m in front of the vehicle and 12 m in the lateral direction. Like [2], lateral position was not reported, but a measurement of the lane width was determined with a high degree of accuracy [3]. In [4], a lidar lane-tracking algorithm is presented for lane detection in urban environments where two six-layer laser scanners were each used to detect a lane marking. The algorithm's performance was compared with a camera system, and the vehicle position in the lane differed between the two systems in the range of 0.02–0.01 m [4].

Camera systems for lane detection are already present in most commercial vehicles as an optional safety system. One such system, by Jung and Kelber [5], used a linear-parabolic model to create an LDW system using lateral offset based on the near and far ranges. Feng *et al.* [6] used an improved Hough transform to obtain the road edge in the binary image, followed by establishment of an area of interest based on the prediction result of a Kalman filter. Hsiao and Yeh [7] avoided the use of the Hough transform and instead relied on peak and edge finding, edge connection, line-segment combination, and lane boundary selection. In [8], an extended Kalman filter was used

to estimate the parameters of the model of the lane markings to search within a specified area in the image so that far lane boundaries are searched with a smaller area than closer lane boundaries, which reduced the impact of noise. Dickmanns and Mysliwetz [9] used clothoid model parameters to recognize horizontal and vertical road parameters in a recursive manner. Khosla [10] used two contiguous clothoid segments with different geometries but with continuous curvature across each clothoid. However, Schwartz [11] argued that the clothoid model is unsuitable for sensor fusion due to the "sloshing" effect of the estimated values between the clothoid parameters.

The navigation filter herein uses both the lidar and camera measurements as well as GPS to correct the drift of the IMU using an extended Kalman filter. Similar work has been done using lane-based maps and lateral lane distance measurements. Miller and Campbell [12] used a particle filter to augment GPS and inertial solutions with vision-based measurements and a map. Li *et al.* fused cameras, low-cost GPS, and map data to achieve centimeter-level performance [13]. Vu *et al.* used a tightly coupled filter along with a measurement of angle of arrival from a known location of a traffic light to correct an Inertial Measurement System (INS) [14]. Since the lidar and the camera both measure lateral distance with respect to the lane markings, a map of the lane markings is needed to relate the camera and lidar measurements to the global coordinate frame of the GPS. Much prior work with map matching primarily is concerned with relating GPS and maps for consumer automotive applications. A GPS navigation system that is derived from a combination map matching algorithm has been developed, which combines the shortest distance algorithm and the assistant map matching algorithm to account for GPS error, city environment error, and road complexity, which can lead to inaccurate matching in the shortest distance algorithms [15]. Another system fuses camera-based speed sign recognition and digitized speed limit maps combined with a GPS sensor for speed limit assistants [16].

This paper uses a system for fusing lateral distance measurements with a known map in a GPS/INS system to analyze the impact of various nonideal scenarios, including loss of lateral distance, loss of GPS, and various satellite geometries. Results show that, by using sensors and systems that are already on commercial vehicles, a positioning solution can be found even in situations where only two satellites are visible. Section II describes the lidar and camera lane-detection systems. Section III shows the experimental results for the lidar and camera lane-detection systems. Section IV describes the navigation filter. Section V analyzes the experimental results of the integrated system under various conditions.

II. Lane-detection Systems

Two lane-detection systems were designed to provide redundant lateral distance measurements from the lane markings to the center of the vehicle. The camera uses visual images to determine the location of the lane markings, whereas the lidar detection system uses reflectivity. Reflectivity is a measurement known as echo width. Reflectivity is often a function of distance from the lidar and a function of the surface of the object measured. Hence, an object that is highly reflective at a large distance may return the same echo width measurement by the lidar as an object that is less reflective, but closer. This reflectivity measurement is often exploited by LDW systems because lane markings are designed to be more reflective than the road's surface [2].

Several assumptions are made for lane-detection systems. The lane markings are assumed to be present on one or both sides of the vehicle in order for a lateral distance estimate to be found, and the vehicle is assumed to be within the lane markings. No assumption is made about the weather or lighting conditions—the time can vary from day through night, and the weather can range from rainy to blue sky. However, the performance of the systems significantly changes, as described in Section III. Vehicle speed is assumed to be typical highway speeds. The sensors themselves measure a few meters (1–3 m) in front of the vehicle. Finally, both the camera and the lidar are assumed to be perfectly aligned with the vehicle and the vehicle is traveling parallel with the lane markings.

A. Camera-Based Lane Detection

This section presents a brief description of the algorithm for determining lateral lane position in the lane using a camera. More information can be found in [17]. The camera is mounted on the roof of the vehicle facing forward and is angled down toward the road. The lateral position is measured just over the hood and a few meters in front of the vehicle. The vision algorithm begins with thresholding, followed by edge detection and extraction of lines in the image using the Hough transform. Left and right lane marking lines are selected using two criteria to be included in the interpolation of a second-order polynomial, which serves as the model for the lane markings in the image. The coefficients of the polynomial are used as measurements into a Kalman filter, whose estimated states represent the estimate for the lane markings in the image. The following sections describe these steps in more detail.

1) Thresholding: The threshold operation must extract the lane markings from the image and ignore extraneous features on the road. Since vehicles drive in varying environments, the threshold operation must take into account different levels of lighting and weather to extract the lane markings. Fig. 1(a) shows an image of a road at dusk.

A dynamic threshold, based on Niblack thresholding [18], takes into account the statistics of the image to better approximate the best value for thresholding. Using the mean and standard deviation of the image, the dynamic threshold is calculated as

$$T = \mu + K\sigma \tag{1}$$

where T is the grayscale threshold value, μ is the mean of the grayscale image, K is heuristically chosen in order to capture the lane markings whose intensities range in the upper grayscale range of the image, and σ is the standard deviation of the intensity of the image. Using $K = 1.5$, the resulting thresholded image is shown in Fig. 1(b). The extracted lane markings are clearly visible.

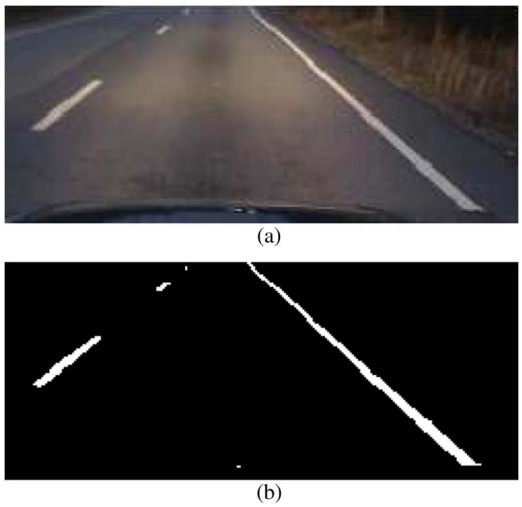

Fig. 1. Darker highway image and thresholded image. (a) Darker original image. (b) Dynamically thresholded image.

2) Edge Detection and the Hough Transform: The edges of the thresholded image are extracted using Canny edge detection [19], with a smaller threshold of 50 and a larger threshold of 200 for the hysteresis process and an aperture size of 3. Often, the edges are broken in various places along the lane marking's edge. The Hough transform [20] extracts the lines from the edge map while ignoring the gaps in the lines. For this work, the resolution for ρ was 5 pixels, and the resolution for θ was 1°. To capture the lane marking lines on curves, the minimum length allowed for a line from the Hough transform was set to be small (30 pixels). A long minimum line length would fail to capture the entire curve. Each line extracted from the Hough transform is considered either a left or a right lane marking depending on its slope. These two groups of lines, so-called line pools, are independently processed and can be made up of many lines. The following sections describe the processing on each of these line pools.

3) Line Selection: Due to erroneous features in the image, such as shadows or objects on the road, extracted lines from the Hough transform may not be lines from lane markings. Two criteria must be met before a line is considered a line from a lane marking edge—a spatial criterion and a slope criterion.

The spatial criterion requires the Hough lines to be close to the last estimated second-order polynomial lane marking model. Two additional second-order polynomial lines are created at equal distance from the last estimated second-order polynomial on both sides to serve as bounding curves. The distance to the polynomial lines changes depending on the success of tracking—consistent tracking of frames reduces the distance to ignore erroneous lines, whereas missed frames increase the distance to enlarge the allowable area for the lane markings. Hough lines outside of these bounds fail, whereas those that are fully within the bounds meet the criteria. Fig. 2 shows a lane marking with the polynomial boundary curves (black solid) around the last estimated lane model (dotted green). The valid blue Hough lines can faintly be seen around the yellow estimated lane model, and the rejected lines due to the spatial criterion can be seen along the edge of the black bounding polynomials as red lines.

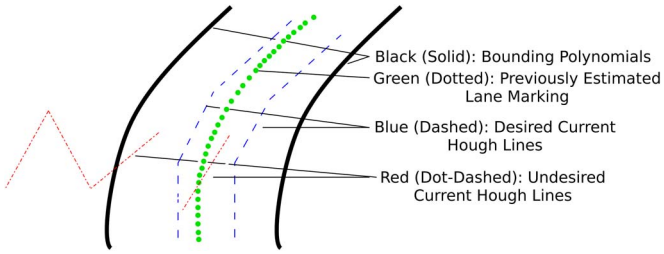

Fig. 2. Spatial and slope criteria for Hough lines from a lane marking.

The slope criterion requires the line to have a slope approximately equal to the slope of the last estimated second-order polynomial lane marking model at the same row in the image. The slope of the lines from frame to frame should not drastically change; thus, the slopes of the lane markings are expected to be close (within 0.3) to the slope of the second-order polynomial curve lane marking model near that lane marking line. The red (dot-dashed) line within the polynomial boundary lines in Fig. 2 passed the spatial criterion but failed the slope criterion. The slope criterion, then, is mostly used to remove erroneous lines near the real lane markings in the image.

4) Lane Modeling and the Kalman Filter: Each Hough line that passes both criteria have their endpoints and midpoints collected in point pools. These point pools then undergo the standard least squares polynomial interpolation to determine the coefficients of the polynomial model for the lane marking. The model for the estimated lane marking consists of a second-order polynomial in the form of

$$y = ax^2 + bx + c \tag{2}$$

where a, b, and c are the coefficients of the polynomial in image coordinates where the origin is located in the upper left corner of the image with the x-axis pointing to the right and the y-axis pointing down.

Despite the selection criteria, erroneous Hough lines may be considered valid lane marking edges, and the resulting polynomial may not align with the lane markings seen in the image. To reduce the impact of these erroneous model estimates, the coefficients of each lane marking line are used as measurements into a linear Kalman filter. This Kalman filter, whose states consist of a, b, and c coefficients for the left and right lane marking polynomial models, solves the least square problem with $A = I$ and the states changing only on the measurement update.

5) Lateral Distance Calculation: Using the coefficients of the polynomial lane model, the lateral lane distance can be calculated by employing known information from the lane markings. The general form of the quadratic equation is used to determine the horizontal point x on the lane model polynomial with respect to any height y within the image. The lateral distance x in pixels is the distance from the calibrated center of the vehicle in the image (in our case, close to the center of the image) to the lane marking. The lateral distance is then multiplied by a known calibration factor to obtain the lateral distance in meters. The calibration factor n is predetermined by measuring the number of pixels in an image that spans the bottom row of the image from the left lane marking to the right

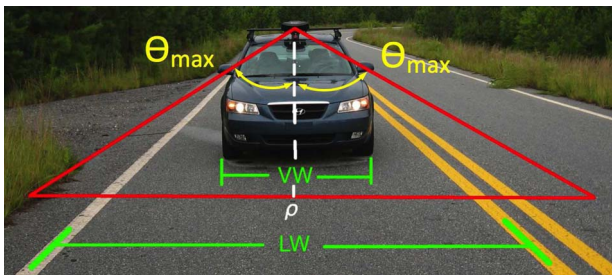

Fig. 3. Scanning for road markings.

lane marking. The bottom row of pixels is chosen as the row to measure the distance to either lane since it has the most resolution with respect to the road.

The hood of the vehicle is cropped from the image and the sky. Since the width of a lane is approximately 3.658 m, the calibration factor is calculated as

$$n = \frac{w_l}{p_c} \tag{3}$$

where w_l is the standard width of a lane, or 3.658 m, and p_c is the pixel count of the lane.

B. Lidar-Based Lane Detection

This section presents the algorithm for determining lateral lane position with a lidar using reflectivity. More information on this lidar-based method can be found in [21].

1) Detection: The detection of lane markings using the reflectivity of a lidar utilizes the principle that the lane markings are more reflective than the surface of the road. The detection algorithm will first bound the search area within the lidar scan in the region where the reflectivity of a lane marking is most likely to lie. Then, an ideal lane model is fit to the scanned data to detect the lane markings. If a lane exists, the position in the lane is determined and filtered to provide a final lateral distance solution.

2) Bounding: The detection algorithm will first bound the search area by assuming an ideal lane width so that, regardless of the position of the vehicle in the lane, a lane marking is likely to be scanned if it is present, as shown in Fig. 3. This bound is found by assuming that the rightmost tires of the vehicle are touching the lane markings. The angle that is required to scan the leftmost lane marker is determined to be the angle bound, i.e., $\theta_{\max}$, as shown in

$$\theta_{\max} = \arctan\left(\frac{\mathrm{LW} - \frac{\mathrm{VW}}{2}}{\rho}\right) \tag{4}$$

where LW represents the lane width, VW represents the vehicle's width, and ρ is the distance between the lidar and the ground measured with zero horizontal angle.

3) Scan Matching: The heart of the lidar lane-detection algorithm is the matching of an ideal lane scan to an actual lane scan. Fig. 4 shows 100 stationary lidar scans of a road averaged together. Note the distinction of the four labeled peaks and their correspondence to Fig. 5.

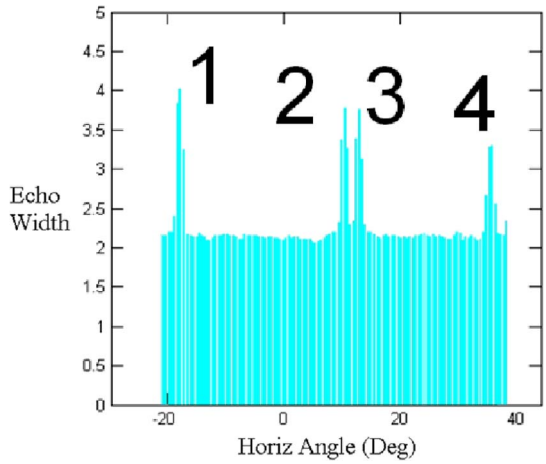

Fig. 4. Averaged lidar scans.

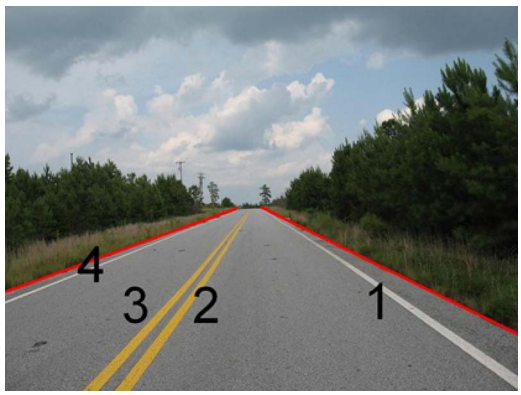

Fig. 5. Lane markings seen in lidar scan.

Fig. 4 shows that the region between the spikes representing the lane markings is a relatively constant area, and outside of the scan bounds, a highly noisy region is created by the environment bordering the roadway, where spikes 2 and 3 refer to the two center lane markings, and spikes 1 and 4 correspond to the outside lane markings. While this averaged scan looks quite ideal because the bulk of the sensor noise has been averaged out, a single scan of the environment can be significantly less intuitive, as shown in Fig. 6.

The algorithm must be capable of detecting lane markings in these types of nonideal conditions where the lane markings may not be distinct due to sensor noise or environmental effects. An ideal lane scan is therefore created to model the reflectivity profile of a lane. This ideal scan is shown in Fig. 7 and will be compared with the actual lidar data to detect the lane markings.

The ideal scan is created by first scanning a narrow area directly in front of the vehicle and averaging the reflectivity; this provides a metric of the reflectivity of the road surface and constitutes the constant area between the spikes in Fig. 7. The spikes are generated by increasing the averaged reflectivity of the road surface by 75% to represent the lane marking. Thus, this ideal lane model is dynamic and changes with each lidar scan in response to changes with the road and environmental conditions. Once this ideal scan is created, it is compared with the actual lidar scan. This comparison is done over the

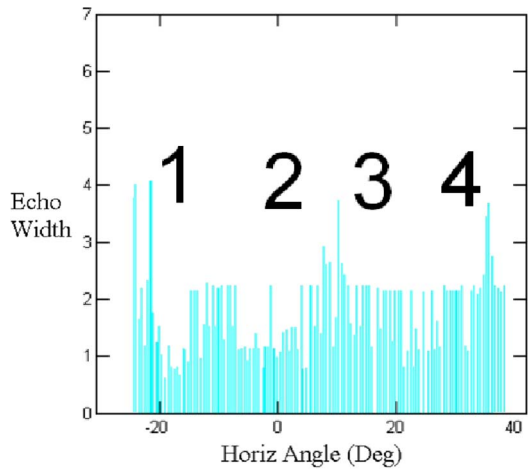

Fig. 6. Typical lidar scan.

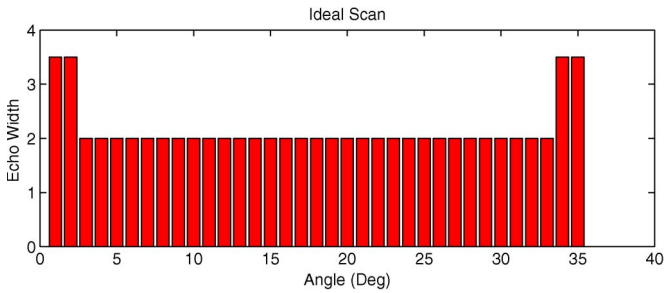

Fig. 7. Ideal scan.

entirety of the lidar scan, where the ideal lidar scan is varied in size from some minimum lane width to some maximum lane width until the search space is exhausted, thus allowing for varying changes in the true lane width. With each comparison, the mean square error between the ideal and actual scans is calculated. The region that provides the minimum mean square error between the ideal scan and the actual scan is determined to be the location of the true lane markings. A final comparison of the reflectivity of the actual lidar scan to where the ideal scan's lane markings are located is analyzed. If the area of the lane markings of the extracted lane is not at least 30% above the averaged reflectivity used to model the surface of the road, an assumption can be made that no lane marking exists. This process is repeated for any additional scan layers for a multilayer lidar. Finally, once a lane marking is detected, a window of approximately 4° is placed on the area from which the lane was detected, as shown in Fig. 8. This window allows for a narrowed search space for subsequent extractions and helps mitigate any sources of error in the scan. However, if no lane markings are detected in this narrowed search space for three consecutive scans, the search space is expanded back to the original bounds.

4) Filtering: Once the lane markings have been detected, the position in the lane and the lane width can be calculated using the distance to the left and right lane markings, which are denoted by d_L and d_R, respectively. Note that d_L will always be a negative number and d_R will always be a positive number in order to conform to the road coordinate frame used, as shown

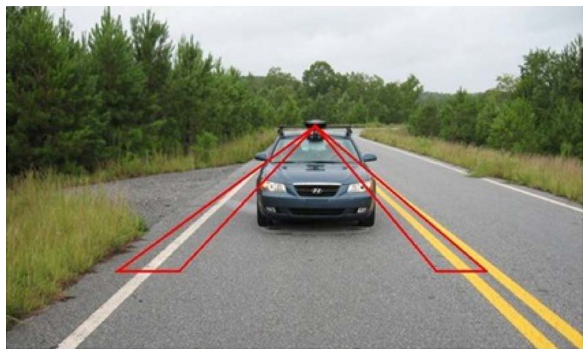

Fig. 8. Bounds for narrowed search area.

in Section IV. Thus, the lane width calculation is given in (5), and this value is stored for future processing, i.e.,

$$\mathrm{LW} = d_R - d_L. \tag{5}$$

The offset from the center of the lane, which is denoted by $\bar{\zeta}_{\mathrm{vision}}$, can be then computed via

$$\bar{\zeta}_{\mathrm{vision}} = -(d_L + d_R)/2. \tag{6}$$

However, if only the left or the right lane markings are detected, the offset from the center of the lane is then found in (7), where the most recent LW measurement is used, i.e.,

$$\bar{\zeta}_{\mathrm{vision}} = -\left(\frac{\mathrm{LW}}{2} + d_L\right) \quad \text{or} \quad \bar{\zeta}_{\mathrm{vision}} = \left(\frac{\mathrm{LW}}{2} - d_R\right). \tag{7}$$

If a simultaneous left and right lane marking measurement has yet to occur so that no lane width estimate yet exists, an ideal lane width of 3.36 m is used. Once this distance from the center of the lane is calculated, the result is low-pass filtered to smooth the data in an effort to limit any erroneous spikes. This filtered result is incorporated into the navigation filter, as shown in Section IV-B2.

III. Lane-Detection Experimental Results

Both lane-detection algorithms were tested at the National Center for Asphalt Technology (NCAT) test track in Opelika, AL, USA. The NCAT test track is a 1.7-mi oval track with 8° bank angles, where lane markings were surveyed at centimeter accuracy using a real-time kinematic (RTK) GPS. The track is divided into various sections, where each section contains a different type of asphalt. Thus, the track provides testing for multiple road types. The lateral lane distances estimated by each lane-detection system were compared with true lane distances using vehicle positions with centimeter accuracy levels using RTK GPS.

A. Camera Lane-Detection Experimental Results

The test vehicle for the camera lane-detection experiment was equipped with a low-cost forward-looking web camera (QuickCam Pro 9000) at 244 × 100 pixels (after cropping the image for region of interest) that was mounted on a conventional roof rack and is similar to cameras found on commercial

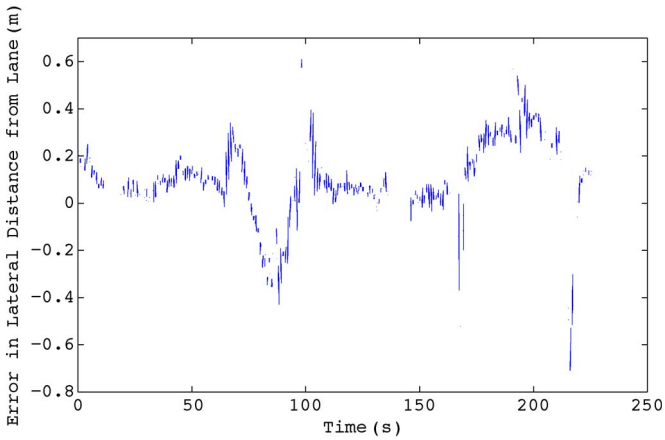

Fig. 9. Camera lateral distance to the right lane marking error.

vehicles. The vehicle was driven on the right lane around the full extent of the NCAT test track at approximately 40 mi/h to simultaneously acquire RTK GPS coordinates and frame grabs at approximately 10 Hz from the camera. Each lateral distance from the camera images was compared with the truth measurements from GPS to determine the accuracy of the system. The accuracy of the coefficient states can be seen by visual inspection of the images. Simply, if the polynomial models lie on the lane markings in the images, then the coefficients are being estimated correctly. Note that the testing metric of lateral distance relies on the polynomial coefficients; thus, the coefficients are indirectly verified. In the experimental test run, the quality of the lane markings is not ideal—the lane markings disappear on some segments of the track and are faint in others. As such, in several frames, no lane marking is detected. In addition, at one part of the track, the asphalt creates a distinct line between light-colored asphalt and dark-colored asphalt that is near the dashed left lane markings, and the left lane marking estimate tracks this line.

Fig. 9 shows the calculated lateral distance error for the test run. From about the 50-s mark to the 100-s mark and from the 150-s mark to the end of the test run, the lateral distance seems to slightly increase compared with the remaining sections of the graph. These sections consist of the curves of the track. While the driver drove closer to the center of the road in the turns, as shown in the truth data, the error between the truth data and the calculated lateral distance increased in this section as well. The lateral distance is measured at the lowest row of pixels in the image, which actually consists of a point in front of the vehicle. Due to the curvature in the lane markings, the actual lateral distance and the measured lateral distance at the point ahead of the vehicle are slightly different. Therefore, the error in the system increases due to the curvature of the lane markings. The maximum error in the system on a straightaway is about 20 cm.

Table I shows the quantitative results of the camera vision system in four test runs. The first test run consisted of a run at dusk, where the first half of the run was facing the Sun. During this run, the system failed to detect the lane markings due to the presence of the Sun in the image, which pushed the threshold above the lane marking values. This impact is reflected in the detection rate of 48%. The second test run was a night run,

TABLE I
CAMERA LATERAL DISTANCE ERROR IN NONIDEAL CONDITIONS IN METERS

	Dusk	Night	Rain Dusk	Rain Night
Average Absolute Error(m)	0.2379	0.0307	0.0327	0.0512
RMS Error(m)	0.4214	0.0401	0.0940	0.1253
std of Error(m)	0.3526	0.0402	0.0887	0.1149
% Detection	48.01	90.00	18.08	19.47

TABLE II
VARIOUS SCENARIOS FOR ROAD TYPE IN METERS

	Avg. Lane Width Error(m)	Std of Error(m)	Detection (%)
Highway Driving (dashed lines)	0.075	0.233	94.7
Yellow/White markings	0.042	0.272	81.7
Gravel on Road	0.129	0.215	97.4
Grass Bordering Road	.169	.329	76.86

where the headlights illuminated the lane marking lines. During the test runs at dusk and at night in the rain, the system failed to detect lane markings for the majority of the run, and the percent detected was less than 20%. For the night, rain and dusk, and rain and night test runs, the average absolute error was less than 6 cm, which signifies that the lateral distance estimates from the camera are appropriate as unbiased measurements in the navigation Kalman filter.

B. Lidar Lane-Detection Results

The hardware used for detecting the lane markings using lidar is an Ibeo ALASCA XT, which is a 3-D automotive-grade lidar. The lidar provides range measurements and reflectivity measurements (echo width). The Ibeo has four vertically diverging beams giving the Ibeo a 3.2° field of view directly in front of the vehicle. Data from the lidar were taken at 10 Hz with an angular resolution of approximately 0.25°.

The lidar was mounted to the roof rack of a vehicle. This location is ideal for the mounting of the lidar because it provides a lateral resolution of approximately 0.040 m at the lane markings while scanning less than 2 m in front of the vehicle. The shortened scan distance in front of the vehicle aids in the reduction of the effects of vehicle yaw and prevents the scanning of other vehicles, thereby reducing any processing time that would have been otherwise dedicated to the removal of these obstructions.

Tests were performed on a number of different scenarios of varied road types where no high-precision GPS measurement was available. Instead, an estimate of the road width was compared with the known road width as an estimate of algorithm performance. Detection rates were also recorded. Table II highlights these results.

Other testing included a comparison of the reported position of the lidar in the lane to a precision GPS survey of the test track. The results of these tests are shown in Figs. 10 and 11.

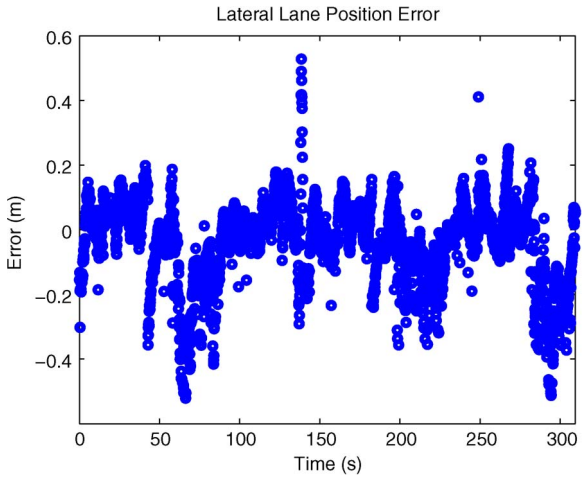

Fig. 10. Lane position lidar error (Test 1).

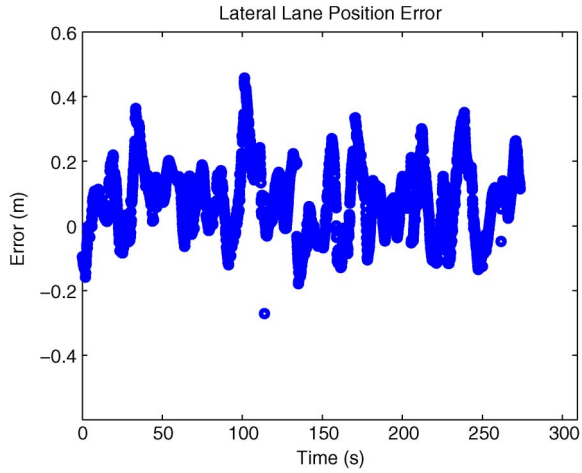

Fig. 11. Lane position lidar error (Test 2).

The first test (Test 1) had a mean error 0.0984 m and a variance of 0.1274 m, whereas the second test (Test 2) has a mean error 0.1074 m and a variance of 0.1115 m.

An additional series of tests were performed under various weather conditions, whose position results are compared with a high-accuracy GPS survey in Table III. The percent detection for the "Rain Dusk" and "Rain Night" sections are lower than the "Dusk" and "Night" data runs due to the presence of rain, which significantly degrades lidar performance, particularly when rain pools on the roadway. However, the lidar is not affected by direct light or low-light situations, as opposed to the camera system.

TABLE III
LIDAR LATERAL DISTANCE ERROR IN NONIDEAL CONDITIONS IN METERS

	Dusk	Night	Rain Dusk	Rain Night
Average Absolute Error(m)	0.2375	0.4974	1.0327	0.9232
RMS Error(m)	0.3559	0.5413	1.1187	1.0034
std of Error(m)	0.3436	0.2152	0.01858	0.1571
% Detection	99	99	18.9	32.57

IV. NAVIGATION FILTER

The six-degree-of-freedom extended Kalman filter estimates position, velocity, and attitude in the Earth Center Earth Fixed (ECEF) coordinate frame using GPS position and velocity or pseudorange and pseudorange rate measurements (depending on which is available) along with an IMU in a GPS/INS solution. The state vector is

$$\vec{x} = \begin{bmatrix} \vec{r}^{\,e}_{eb} & \vec{v}^{\,e} & \vec{\psi}^{e} & \vec{b}^{\,b}_{a} & \vec{b}^{\,b}_{g} & \vec{dt}_{\text{gps}} \end{bmatrix} \tag{8}$$

where $\vec{r}^{\,e}_{eb}$ is the 3-D estimate of position, $\vec{v}^{\,e}$ is the 3-D estimate of velocity, and $\vec{\psi}^{e}$ is the estimate for the three Euler angles that describe the attitude of the vehicle and is expressed in terms of the three necessary rotations to rotate the ECEF coordinate frame to align with the body coordinate frame. The bias states for the IMU consist of the three accelerometer biases, i.e., $\vec{b}^{\,b}_{a}$, and the three gyroscope biases, i.e., $\vec{b}^{\,b}_{g}$. Finally, $\vec{dt}_{\text{gps}}$ contains the clock bias and drift terms necessary to use GPS pseudorange and pseudorange rate measurements and is only necessary if the pseudorange and pseudorange rate measurements are used. The state vector, then, consists of 15–17 states depending on the type of GPS measurements being used.

A. Time Update Propagation Equations

The time update consists of the propagation of the states forward through time through inputs using inertial measurements from the IMU. The equations are given as

$$\dot{\vec{x}} = \begin{bmatrix} \vec{x}_{4:6} \\ C^{e}_{b}(\vec{u}_{1:3} - \vec{x}_{10:12}) + \omega_e \begin{bmatrix} \omega_e \vec{x}_1 + 2\vec{x}_5 \\ \omega_e \vec{x}_2 - 2\vec{x}_4 \\ 0 \end{bmatrix} - g_e \\ C_{\text{MECH}}(\vec{u}_{4:6} - \vec{x}_{13:15} - \begin{bmatrix} 0 \\ 0 \\ \omega_e \end{bmatrix} \\ \vec{0}_{5\times 1} \\ \vec{x}_{17} \\ 0 \end{bmatrix}. \tag{9}$$

Equation (10) splits the specific force caused by gravity into its vector components in the ECEF coordinate frame, i.e.,

$$g_e = \text{GM}\left(\vec{x}_1^2 + \vec{x}_2^2 + \vec{x}_3^2\right)^{-1.5} \vec{x}_{1:3}. \tag{10}$$

C_{MECH} is the IMU mechanization matrix, i.e.,

$$C_{\text{MECH}} = \begin{bmatrix} 1 & \frac{s_7 s_8}{c_8} & \frac{c_7 s_8}{c_8} \\ 0 & c_7 & -s_7 \\ 0 & \frac{s_7}{c_8} & \frac{c_7}{c_8} \end{bmatrix}. \tag{11}$$

The body frame's origin is located at the IMU's position in the center console of the vehicle, which is approximately the center of mass of the vehicle. The x-axis points in the forward direction (longitudinal) with respect to the vehicle, the y-axis points to the right (lateral), and the z-axis points down. The lidar and camera are mounted directly above the center console on the roof and measure lateral distance in the road coordinate frame. The GPS antenna is mounted behind the lidar and camera on the roof of the vehicle along the center line of

the vehicle and is approximately 1 m behind and 1 m above the IMU, which results in a lever arm for the GPS antenna with respect to the IMU of $l_a = [-1 \quad 0 \quad -1]$. The rotation matrix from the body frame into the ECEF coordinate frame is given by

$$C_b^e = \begin{bmatrix} c_8c_9 & s_7s_8c_9 - c_7s_9 & c_7s_8c_9 + s_7s_9 \\ c_8s_9 & s_7s_8s_9 + c_7c_9 & c_7s_8s_9 - s_7c_9 \\ -s_8 & s_7c_8 & c_7c_8 \end{bmatrix}. \quad (12)$$

For these equations, GM $= 3.896004418 \times 10^{14}$ m^3/s^2, s_i is the sine of the ith state, c_i is the cosine of the ith state, and $\vec{u}$ is the input vector consisting of the accelerations and angular rates from the IMU. The state matrix A for the extended Kalman filter, then, is the Jacobian matrix obtained from the partial derivative of each change in state equation of (9) evaluated at the current estimate.

The filter uses the lane position measurements from the camera and lidar in conjunction with a map. Much of the analysis of this paper will focus on the improvement of the filter due to this lane position measurement. Additional information on the navigation filter can be found in [22]. The next sections describe how the vision measurements are incorporated into the global navigation filter.

B. Vision Measurement Updates Using a Map

The vision measurements, either from the lidar or the camera, give an estimated lateral position in the current lane in the road coordinate frame. For this paper, the road coordinate frame is approximated using a waypoint map, and each lane is assumed to have its own associated lane map. The waypoints lie in the center of the mapped lane, and the distance between waypoints is defined by the road geometry. Complex road geometry requires waypoints to be close together. For example, the waypoints in a turn must be close enough to capture the geometry of the road. On a straightaway, the waypoints can be very spread out due to the lack of change in road geometry.

The road coordinate frame is a 3-D coordinate frame and is denoted by the subscript/superscript r. The x-axis of the road coordinate frame points from the last waypoint passed to the next waypoint. The y-axis of the road coordinate frame is perpendicular to the x-axis. If facing the direction of travel, the y-axis points to the right. The road coordinate frame is assumed to have no superelevation; therefore, the y-axis is always parallel with the plane tangent to the Earth's reference ellipsoid. The z-axis is perpendicular to the $x - y$ plane and points to the center of the Earth. This type of coordinate frame was chosen because the vision measurements give a direct measurement of the y-axis position. In addition, since this work was developed for ground vehicles, the position of the sensor in the z-axis of the road coordinate frame is known, can be measured once, and is included in the navigation filter as a measurement update. Note that this height is independent of the geometry of the road itself and is related to the global coordinate frame through the map. The change in vertical height on any point on the vehicle is a function of the roll angle of the vehicle and the height of the point above the roll axis; however, this angle and height above the roll axis are small for normal operating conditions. A similar idea is presented in [23].

In order to use the measurements given in the road coordinate frame, the estimated position in the navigation (ECEF) coordinate frame must be mapped into the road coordinate frame. Equation (13) is used to map position from the navigation coordinate frame to the road coordinate frame

$$\vec{r}_{rb}^{\,r} = C_e^r \left[\vec{r}_{eb}^{\,e} - \vec{r}_{er}^{\,e}\right] \quad (13)$$

where $\vec{r}_{eb}^{\,e}$ is the estimated position vector in the navigation coordinate frame, $\vec{r}_{er}^{\,e}$ is the position of the origin of the road coordinate frame, C_e^r is the rotation matrix from the ECEF coordinate frame to the road coordinate frame, as shown in (15), and $\vec{r}_{rb}^{\,r}$ is the estimated position vector in the road coordinate frame. In order to use this equation, the position of the origin of the road coordinate frame, which in practice is implemented as the last waypoint the vehicle passed in the waypoint map, must be known $(\vec{r}_{er}^{\,e})$. The rotation matrix from the navigation coordinate frame to the road coordinate frame (C_e^r) must be also known and is constructed with the known attitude of the road coordinate frame with respect to the navigation coordinate frame, which is represented by three Euler angles, which rotates the z-axis first, followed by the y-axis, and, finally, the x-axis. The next section describes the waypoint-based lane map.

1) Waypoint-Based Lane Map: The information needed for the lane map is the position of each waypoint in the ECEF coordinate frame, the three Euler angles necessary to rotate the ECEF coordinate frame to the road coordinate frame for each waypoint, and the distance to the next waypoint for each waypoint. Therefore, the map database takes the form

$$\text{Map Database} = \begin{bmatrix} \vec{r}_{er,1}^{\,e} & \vec{\phi}_1 & d_{r,1} \\ \vdots & \vdots & \vdots \\ \vec{r}_{er,m}^{\,e} & \vec{\phi}_m & d_{r,m} \end{bmatrix} \quad (14)$$

where $\vec{r}_{er,i}^{\,e}$ is the position vector of waypoint i in the ECEF coordinate frame, and $\vec{\phi}_i$ is the attitude vector of the road coordinate frame. $d_{r,i}$ is the distance from waypoint i to waypoint $i + 1$ and is used to check if the vehicle has passed the next waypoint in order to keep up with the location of the vehicle along the map.

The form of the rotation matrix from the ECEF coordinate frame to the road coordinate frame is given as

$$C_e^r = \begin{bmatrix} c_2c_3 & c_2s_3 & -s_2 \\ s_1s_2c_3 - c_1s_3 & s_1s_2s_3 + c_1c_3 & s_1c_2 \\ c_1s_2c_3 + s_1s_3 & c_1s_2s_3 - s_1c_3 & c_1c_2 \end{bmatrix}. \quad (15)$$

The elements from the road coordinate frame attitude $\vec{\phi}_i$ are used to construct the rotation matrix from ECEF to the road coordinate frame. The rotation matrix and the position of the road coordinate frame (position of the last waypoint passed) can be used to solve for estimated position in the road coordinate frame, where c_1 is the cosine of the first attitude angle in the attitude vector, and s_1 is the sine of the first attitude angle in the attitude vector. Similarly, c_2 and s_2 are the

trigonometric functions of the second angle, and c_3 and s_3 are the trigonometric functions of the third angle.

2) Vision Measurement Model: To use lateral lane position measurements from the vision sensors, the lateral lane position with respect to the lane map must be estimated using the current states of the navigation filter. Equation (16) is used to find the position estimates in the road coordinate frame, i.e.,

$$\begin{bmatrix} \hat{x} \\ \hat{y} \\ \hat{z} \end{bmatrix} = C_e^r \left[\vec{r}_{eb}^{\,e} - \vec{r}_{er,i}^{\,e} \right]. \tag{16}$$

$\vec{r}_{eb}^{\,e}$ denotes the current position estimate in the ECEF coordinate frame. $\vec{r}_{er,i}^{\,e}$ denotes the position of the last waypoint passed in the ECEF coordinate frame. $\hat{y}$ is the lateral lane position estimate, $\hat{x}$ is the distance into the current road coordinate frame, and $\hat{z}$ is the vertical position in the current road coordinate frame. The z-axis of the road coordinate frame points to the center of the Earth.

The three Euler angles that are substituted into C_e^r come from the map database $\vec{\phi}_i$ along with the position of the last waypoint passed $\vec{r}_{er,i}^{\,e}$. The first three states of the filter $\hat{x}_{1:3}$ are the position vector of the vehicle in the ECEF coordinate frame $\vec{r}_{eb,i}^{\,e}$. To perform a measurement update, three requirements must be known. The first requirement is the measurement vector $\vec{\zeta}$, as shown in

$$\vec{\zeta} = \begin{bmatrix} \tilde{\zeta}_{\text{vision}} \\ \tilde{\zeta}_{\text{height}} \end{bmatrix}. \tag{17}$$

The first measurement is the lane position measured by the lidar or the camera. The second measurement is the height above the lane.

The second requirement needed to perform a measurement update is the measurement equations. The measurement equations for the y- and z-axes of the road coordinate frame are given as

$$\begin{aligned} &\hat{\vec{\zeta}}\left(\vec{x}, \vec{r}_{er}^{\,e}, \vec{\phi}_i\right) \\ &\quad = \begin{bmatrix} \hat{\zeta} \\ \hat{h} \end{bmatrix} \\ &\quad = \Big[C_{e(2,1)}^r \left(\vec{r}_{eb,1}^{\,e} - \vec{r}_{er,i,1}^{\,e} \right) + C_{e(2,2)}^r \left(\vec{r}_{eb,2}^{\,e} - \vec{r}_{er,i,2}^{\,e} \right) \\ &\qquad - C_{e(3,1)}^r \left(\vec{r}_{eb,1}^{\,e} - \vec{r}_{er,i,1}^{\,e} \right) - C_{e(3,2)}^r \left(\vec{r}_{eb,2}^{\,e} - \vec{r}_{er,i,2}^{\,e} \right) \\ &\qquad \cdots + C_{e(2,3)}^r \left(\vec{r}_{eb,3}^{\,e} - \vec{r}_{er,i,3}^{\,e} \right) \\ &\qquad \cdots - C_{e(3,3)}^r \left(\vec{r}_{eb,3}^{\,e} - \vec{r}_{er,i,3}^{\,e} \right) \Big]. \end{aligned} \tag{18}$$

The measurement equations are a function of the states of the filter and the map parameters. The states and map parameters are substituted into these equations to estimate the measurement based on the current states. These equations are also used to determine the measurement model matrix, i.e., H. The first measurement equation is for the lateral lane position, which is provided by the camera or the lidar. The second measurement equation is the height above the road coordinate frame.

The last requirement needed to perform a measurement update is the linearized measurement matrix H. The linearized measurement model is created by taking the partial derivative of the measurement equations with respect to each state. The measurement matrix is

$$H = \begin{bmatrix} \hat{e}_1 & 0_{1\times3} & 0_{1\times3} & 0_{1\times3} & 0_{1\times3} & 0_{1\times3} & 0_{1\times2} \\ \hat{e}_2 & 0_{1\times3} & 0_{1\times3} & 0_{1\times3} & 0_{1\times3} & 0_{1\times3} & 0_{1\times2} \end{bmatrix} \tag{19}$$

where $\hat{e}_1$ is the unit vector in the navigation coordinate frame that points in the direction of the y-axis of the road coordinate frame, as shown in

$$\hat{e}_1 = \left[C_{e(2,1)}^r, C_{e(2,2)}^r, C_{e(2,3)}^r \right] \tag{20}$$

and $\hat{e}_2$ is the unit vector in the navigation coordinate frame that points in the direction of the negative z-axis of the road coordinate frame, as shown in

$$\hat{e}_2 = \left[-C_{e(3,1)}^r, -C_{e(3,2)}^r, -C_{e(3,3)}^r \right]. \tag{21}$$

Both unit vectors are directly taken from rows of the rotation matrix from the navigation coordinate frame to the road coordinate frame.

C. Observability Analysis

With the introduction of additional sensors to the full system, a discussion on the observability is necessary. The observability analysis conducted for this paper uses the time-varying linearized observability test [24], [25], i.e.,

$$O(k_f, k_0) = \sum_{k=0}^{k_f} \phi^T(k, k_0) H^T(k) H(k) \phi(k, k_0) \tag{22}$$

where

$$\phi(k, k_0) = J_d(k) J_d(k-1), \ldots, J_d(k_0) \tag{23}$$

where $J_d(k)$ is the discretized Jacobian of the system dynamics, k_f was chosen to be 17 (window size), the same as the number of states for the tightly coupled filter, and k_0 is the time for which the observability test is done. In the observability analysis for this nonlinear system, the system is linearized about the trajectory. Since this operation is time varying, integration must be done over short periods to see the observability at each time. The observability matrix value is determined by checking the rank of the observability matrix at each measurement epoch and is a function of the available measurements. The rank of the observability matrix in the linearized observability analysis case signifies the number of observable modes, rather than states. Modes are states or linear functions of states. The observability matrix's rank, then, describes a condition where states and linear combinations of states are observable, but the nonlinear filter may not be able to distinguish between the linear functions of states. The time-varying linearized observability test used in this paper is a necessary but not sufficient condition for observability; thus, this observability analysis shows conditions

that the filter is guaranteed to fail when the observability matrix is not full rank.

Nevertheless, knowing the conditions in which the filter is guaranteed to fail is important. The height and lane position measurements provide measurements in two of the three coordinate directions. The system is almost full rank just using these measurements. The missing piece of information needed for full rank is the position in the x-axis of the road coordinate frame. If this position is measured (using something other than GPS), then the navigation filter will remain full rank even without GPS measurements. The results, shown in the following sections, will show that, when using vision measurement updates, the navigation filter remains full rank even if only two GPS ranges are available (instead of the typical four needed to calculate a position estimate).

V. Navigation Filter Experimental Results

This section presents results from the navigation filter and an evaluation of the performance of the filter in various scenarios. Along with showing the filter with all measurements (GPS, lidar, and camera), various sensor outages are simulated to show how the loss of a particular sensor affects the estimates of the filter. The results also show how the filter reacts when no lane map is available. Finally, the results will show how the filter performs when less than four GPS satellites are available and the result of losing vision measurements or the lane map under a limited satellite constellation.

All the error plots where created using the unfiltered RTK GPS position and velocities from the Septentrio GPS receiver as ground truth. The RTK corrections were provided by an on-site base station. The close proximity to the base station ensures that RTK GPS position is accurate on a centimeter level. These RTK positions and velocities (taken at 10 Hz) were resampled to 100 Hz to match the output rate of the navigation filter. The error was estimated by differencing the resampled RTK GPS measurements and the navigation estimates of the filter.

The data for the following results were collected at the NCAT test track in Opelika. For this data set, the vehicle started in the outside lane at the end of the front (north) straight. The test vehicle drove down the front straight at 55 mi/h, through a 180° turn, then down the back straight at 70 mi/h. The vehicle came to a complete stop at the end of the back straight.

A. Vision-Aided Lane-Level Positioning

Fig. 12 shows the estimated position of the vehicle compared with the reference RTK GPS. The red dots represent the RTK reference solution (at 10 Hz). As shown in Fig. 12, the reference solution resides close to the center of the lane as expected. The green dots represent the filtered solution (including IMU at 100 Hz) with non-RTK GPS, vision, and map constraints as measurements.

Fig. 13 shows the estimated position of the filter when using stand-alone GPS measurements. The red dots represent the stand-alone (not RTK) GPS measurements (at 2 Hz). As can be seen, the GPS measurements no longer reside within the lane lines due to the fact that stand-alone GPS does not have lane-level accuracy. The green dots represent the filtered solution (at 100 Hz) with stand-alone GPS, vision, and map constraints. The filtered solution resides within the lane lines since the vision measurement constrains the position to the map. This comparison shows that GPS alone is not enough to determine lane position; however, fusing the GPS with vision and map constraints results in a solution that is lane level accurate. If the vision measurements are unavailable, the filtered solution will converge to the GPS measurements at a rate dependent on the measurement and process noise values in the filter, which results in a loss of lane-level accuracy, as expected.

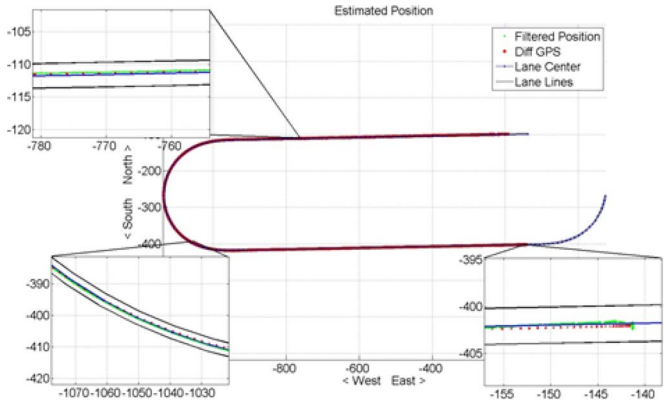

Fig. 12. Estimated vehicle position compared with RTK GPS.

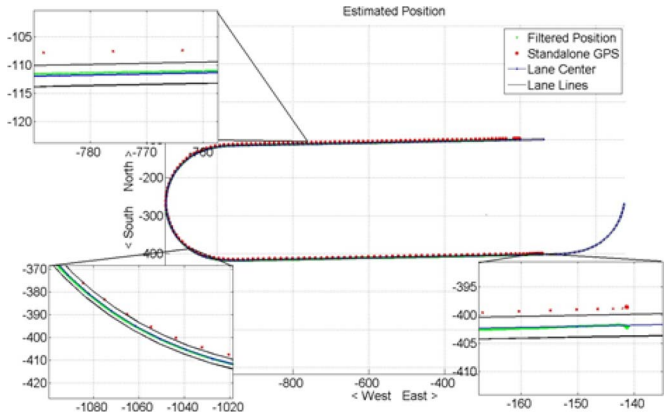

Fig. 13. Estimated vehicle position compared with stand-alone GPS.

B. Effects of Sensor Failures

The next set of figures shows the effects of sensor failures on the system. The dark gray horizontal lines represent the lane marking locations. At 30 s into the data run, different sensor outages are simulated to show what happens to the filter when certain sensors become unavailable. The simulated outages last for 1 min, after which the measurements are reinstated.

Fig. 14 shows three plots of the performance of the estimated lane position under several different failure conditions. The blue line (solid) in all three plots represents the lane position reported by the RTK GPS alone, which was used to represent true position in the lane. The cyan line (dashed, top plot) shows the lane position for the full integrated system—non-RTK GPS,

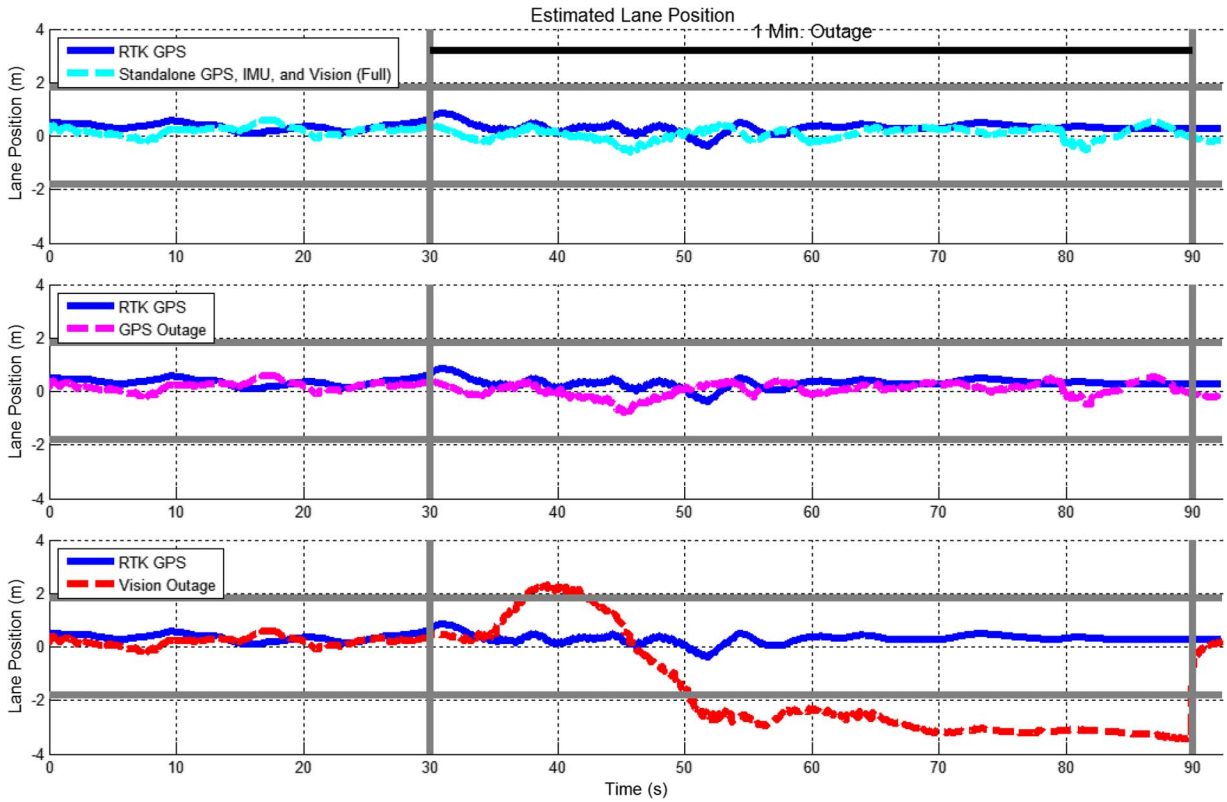

Fig. 14. Estimated lane position with various measurement outages.

IMU, and vision sensors. In this system, after the outage, the filter can still estimate position to lane-level accuracy because vision measurements are still available. When GPS is completely turned off (purple dashed) in the middle plot outage, lane position can be estimated correctly because vision measurements are again still available. The bottom plot (red dashed) shows the loss of lane-level positioning. For this situation, vision is not used; thus, the GPS outage allows divergence of the INS estimation error. Loss of lane position accuracy for this case occurs within 5 s.

Fig. 15 is a plot of the estimated error in longitudinal position—the error in the distance along the path of the map. Two failure conditions result in a loss of accuracy in the estimated distance along the map. The first failure condition is the GPS outage, shown as the dotted purple line. For this case, only vision measurements are available. As shown in Fig. 14, the lane position estimate still remains accurate for this case. However, error accumulates in the longitudinal direction because that direction is unobservable during the time interval without GPS. In the lateral axis, the lane position estimate is still available. The second failure condition is when both GPS and vision are not available. Fig. 15 shows that the error in estimated position grows quicker for the case when no vision measurements are available and that having lateral vision measurements helps in constraining drift in longitudinal position; however, the error for both the loss of GPS (dotted magenta line) and the loss of GPS and vision (dashed black line) results in a continual error growth. Like the estimated lane position, the estimated longitudinal position is corrected by the filter after measurements are reinstated.

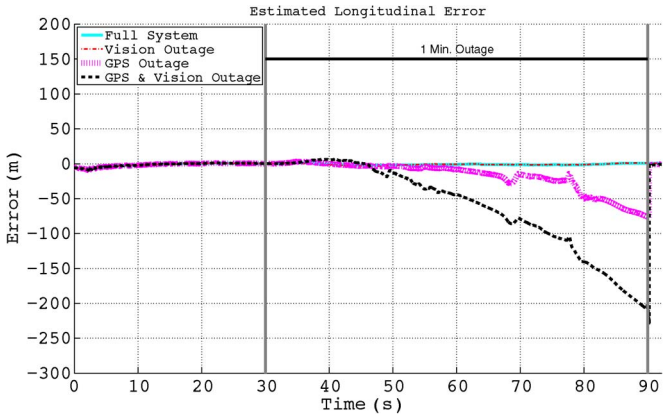

Fig. 15. Estimated longitudinal position error with measurement failures.

Fig. 16 provides the estimated error in each road axis when no GPS measurements are available. The blue line represents the error when no GPS, map (height measurement), or vision measurements are available. For this case, there are no measurement updates. Therefore, a continual error growth exists in all axes due to integration of the IMU errors, as expected. The black line represents the error when no GPS or vision measurements are available. For this case, the map is available, which constrains the error in the vertical road axis due to the height constraint. The magenta line shows the error when only GPS is not available. Vision measurements and the map height constraint are available and constrain the error in the lateral road axis. The error in the longitudinal axis is constrained with GPS.

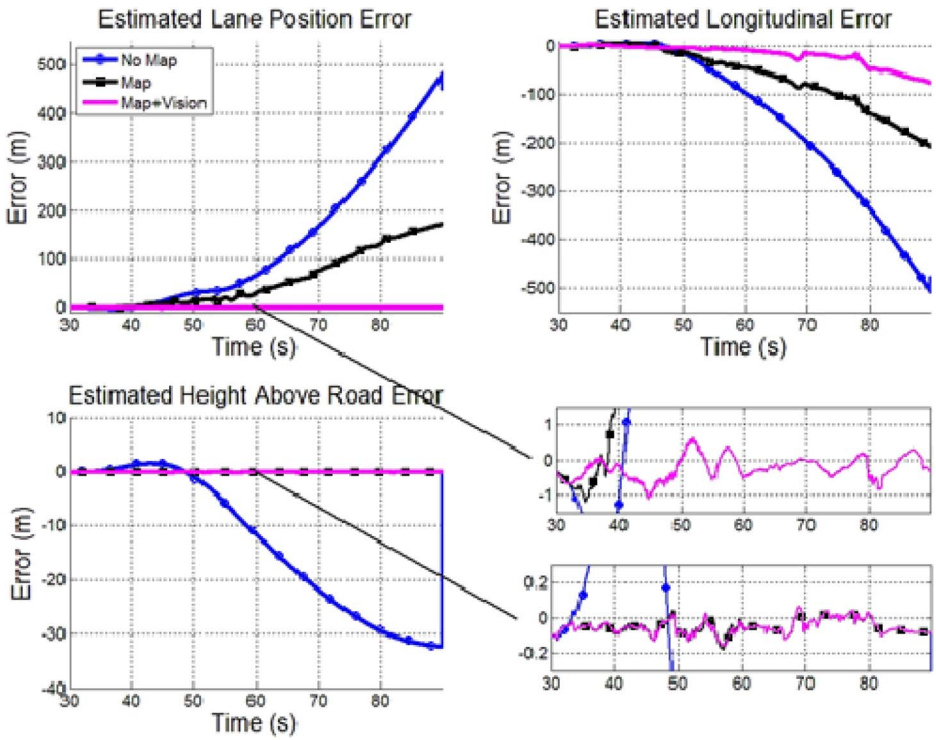

Fig. 16. Estimated error for various map and vision availability using no GPS.

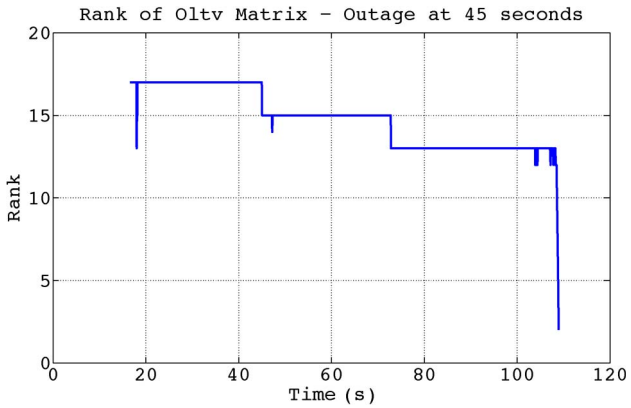

Fig. 17. Rank of observability matrix with vision and height measurements after outage.

Neither the vision measurements of lane lateral position nor the height constraint provides observability in the longitudinal road axis. Fig. 17 shows the observability matrix rank for the same test run where the GPS outage is simulated at 45 s into the test run. In this case, the pseudorange and pseudorange rate measurements and the vision lane position measurements and map are used before failure, resulting in full rank (17 states). As shown in the figure, the observability matrix rank drops to 15 after the GPS failure, which indicates a filter failure condition and results in the drift shown in Fig. 15. Around 72 s, the rank drops again to 13. This drop is most likely due to the lack of dynamics in the system after the 70-s mark, at which point the vehicle exits from the curved portion of the track. Another test (not shown here) where the outage is simulated earlier in the test run results in a drop to 13 and an increase to 15 as the vehicle enters the curve.

TABLE IV
ERROR STATISTICS OF ESTIMATED LANE POSITION GIVEN DIFFERENT MEASUREMENT AVAILABILITY

Failure Type	Average Absolute Error(m)	RMS Error(m)	Standard Deviation of Error(m)	Variance of Error (m^2)
None	0.1537	0.222	0.1601	0.0256
Vision	1.65	2.414	1.7591	3.0945
GPS	0.1502	0.2386	0.1854	0.0343
GPS and Vision	*	*	*	*

TABLE V
ERROR STATISTICS OF ESTIMATED LONGITUDINAL POSITION GIVEN DIFFERENT MEASUREMENT AVAILABILITY

Failure Type	Average Absolute Error(m)	RMS Error(m)	Standard Deviation of Error(m)	Variance of Error (m^2)
None	0.8571	1.2101	0.8541	0.7295
Vision	0.6193	1.1057	0.9159	0.8389
GPS	*	*	*	*
GPS and Vision	*	*	*	*

Table IV shows the standard deviation of the error for each combination of available measurements. For the case when only RTK GPS is used, the measurement of truth and the measurement used to update the filter are the same, resulting in a low standard deviation of error.

Table V gives the error statistics of the estimated longitudinal position for various available measurements. The error statistics of the longitudinal position match the error statistics of the GPS measurement used. When using RTK GPS, the error is on the centimeter level. When non-RTK GPS is used, the error is on

the meter level. There is a continual growth of longitudinal position error if no GPS measurements are available. The growth of the error is slowed if vision measurements of lane position are available; however, the error will continue to grow until a GPS measurement becomes available.

C. Limited Satellite Results

Urban and rural environments both pose different problems for each of the various algorithms. Navigation within a city has proved to be an interesting problem for GPS-aided navigation systems due to the blockage of satellites by buildings. Similarly, trees on rural roads can block GPS satellites. However, the use of multiple sensors can provide a solution that does not drift even with poor GPS satellite visibility. A typical navigation filter based on GPS measurement updates requires at least four GPS observations to calculate a GPS measurement. Each of the four observations is needed to solve an unknown when calculating a GPS position. Three-dimensional position makes up three of the unknowns. The fourth unknown is the clock bias (or the difference in the GPS receiver's clock and the GPS system time).

The filter's vision measurements can measure position in one axis—the axis perpendicular to the road. Since the road map contains vertical information, position in the axis perpendicular to the ground plane can be constrained. The height of the vehicle above the road is assumed to be constant and known. Therefore, the only unconstrained axis is the axis parallel to the road. Since position is only unconstrained in one axis, using only two GPS observations will result in a fully operational system. One GPS observation is needed for the unknown axis, and the other is needed for receiver clock errors. Note that the INS error growth in the ECEF frame is also constrained due to road direction changes from the vision-based measurement of lane position correcting different components of the ECEF error.

The plots presented here show the same data run from Section V-B; however, at 30 s, all but a few selected GPS observations are turned off. After 1 min, all the GPS observations are turned back on. The simulated satellite outages will show how the filter with IMU and vision measurements reacts when less than four GPS observations are available and the effects satellite geometry and measurement quality have on the filter. Like the measurement failure results, these results were created by looking at the effects of satellite failures on the same data set. Thirty seconds into the run, all satellite observations, except the satellite observations listed in the legend, are turned off.

Fig. 18 shows a plot of the estimated lane position for the various satellite constellations, where the numbers next to SV pertain to each satellite's ID number. The blue line is the estimated truth using the RTK GPS measurements. The horizontal gray lines represent the lane lines assuming a 12-ft lane. When vision measurements are used, the lane position estimate remains accurate for all limited satellite constellations. Even when no GPS observations are available, the vision measurements ensure that lateral lane position accuracy is maintained. The lane position estimate reflects a lateral position, and the component of any GPS observations that lie in the lateral road

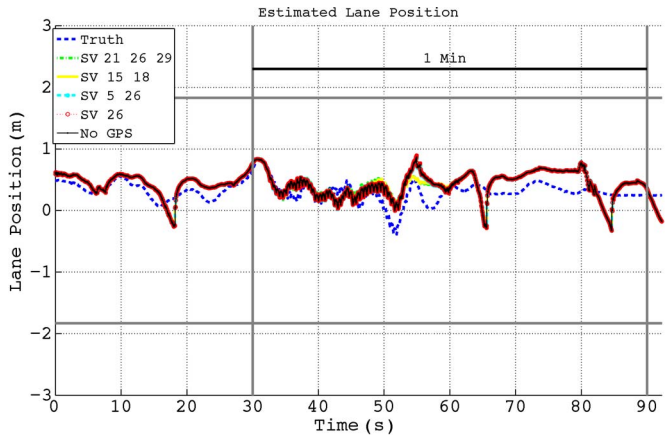

Fig. 18. Estimated lane position with vision and limited satellite visibility.

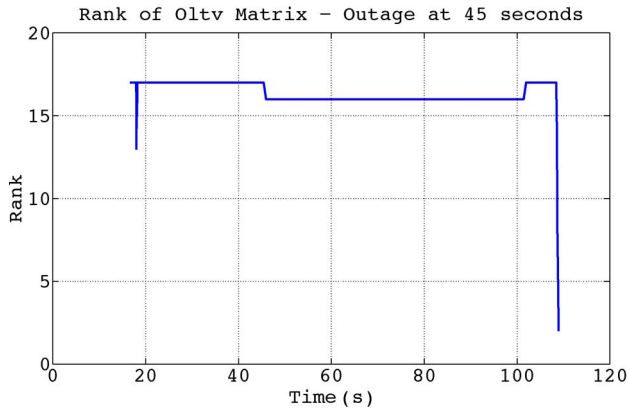

Fig. 19. Rank of observability matrix with two satellites.

axis will be automatically weighed less because the vision measurements have a much smaller associated variance.

If vision measurements are not available, the lack of observability in the filter allows the estimation error to grow, as expected. Fig. 19 shows the rank of the observability matrix for the data run with an outage at 45 s with two satellites. As shown in the figure, the rank after the outage remains at 14 until the end of the outage, which signifies a guaranteed filter failure.

Fig. 20 shows a plot of the estimated error in longitudinal position with only the satellites shown in the legend. The figure shows that the estimated longitudinal position contains a large bias for the case when only SV 5 and SV 26 are used. This large bias is likely due to the satellite geometry of the satellites—the azimuth angles in both of these satellites are close together. For the case when only SV 15 and SV 18 are used, the bias present is along the level seen when using stand-alone GPS with more than four observations. This bias is most likely due to the azimuth angles being farther apart. If only two observations are available, then the filter still remains full rank.

However, a bias can be present in the estimated longitudinal position. This bias is a function of the quality of the observations available and the constellation. The best results appear to come when the two observations are 180° apart in azimuth angle and when both of the azimuth angles are orthogonal to the road. Typically, in an urban environment, observable satellites will reside in the axis parallel with the direction of the road. The

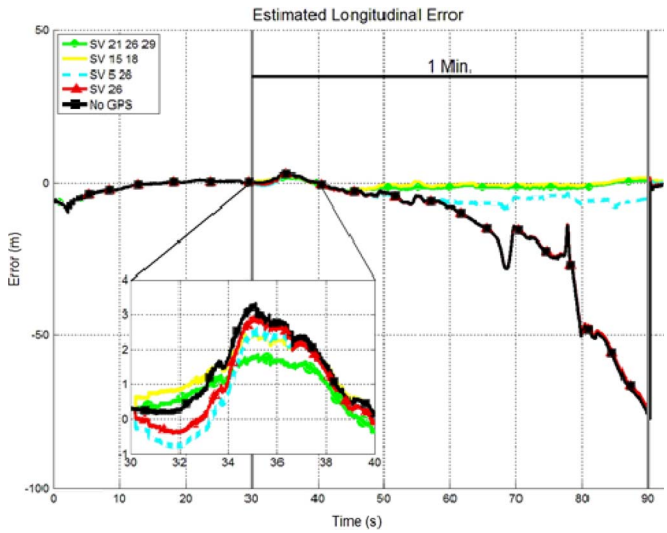

Fig. 20. Estimated longitudinal position error with vision and limited satellite visibility, only using the satellites shown in the legend.

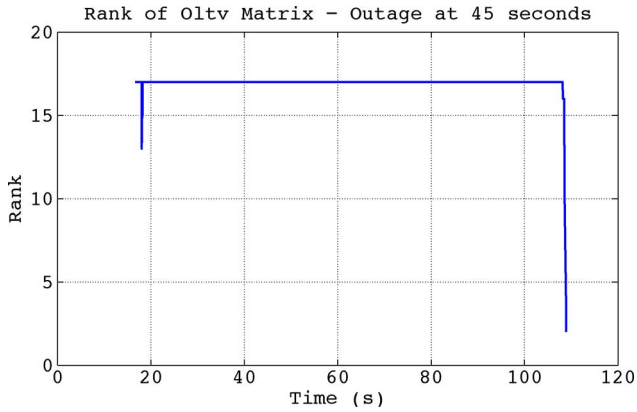

Fig. 21. Rank of observability matrix with two satellites, vision measurements, and height measurements.

estimated longitudinal position cannot be corrected with GPS if one or no GPS observations are available.

Fig. 21 shows the observability analysis for the experimental data where the satellite geometry results in only two satellite observations, vision measurements, and height measurements. The plot shows that the system is full rank before and after the failure with only two satellites in view. In this case, the filter is not guaranteed to diverge, and its observability is based on the vehicle dynamics.

VI. Conclusion

This paper has shown that vision and height measurements, along with a lane map, will greatly reduce the drift experienced when GPS is lost, such as under heavy foliage or in an urban canyon, using sensors already present in commercial vehicles and the inclusion of a lane map. The subsystems for lateral distance from both the camera and the lidar showed submeter accuracy and provided measurements for the integrated positioning system. A known lane marking map related the lateral distance measurements from the local coordinate frame to the global coordinate frame is used in the navigation filter.

TABLE VI
Impact of Sensor Loss on Navigation Filter

Failure Type	Impact
None	Lane level positioning
Vision	Position converges to GPS/INS solution
GPS (all but 2 satellites)	Position constrained depending on satellite geometry
GPS (Full)	Lateral position constrained to road, longitudinal position drifts

The navigation filter fused measurements from range data of satellites from GPS, inertial data from an IMU, lateral distance information from the camera and lidar, and known height data to constrain the position estimate in situations where less than four satellites were available for GPS. Experimental tests conducted at the NCAT test track showed that the use of vision data prevented drift due to an INS-only solution in situations with low satellite visibility.

Table VI shows a summary of the loss of various measurements for the filter. Vision and height measurements constrain IMU drift in two axes; thus, the solution only drifts in the road's longitudinal axis (assuming a straight road). One GPS range could be used to observe position in the road's longitudinal axis. However, using GPS observations adds another unobservable mode (receiver clock bias and drift). Two GPS observations are needed (along with vision measurements and a map) in order to maintain full rank of the observability matrix. At least one of the GPS observations must have a component in the road's longitudinal axis; therefore, as long as both GPS observations do not lie in the plane perpendicular to the road and vertical to the vehicle, full rank of the observability matrix is maintained. It is unlikely that two GPS observations will have the same azimuth angle or an exactly 180° azimuth angle difference. While the navigation filter remains fully observable with two GPS observations, map, and vision measurements, the performance of positioning in the road's longitudinal axis is dependent on the degree of observability, the constellation of the two GPS observations, and the error in the observations.

Future work involves further testing of the full system in various environments, such as different weather conditions and time of day. Other work includes resolving the detection of lane changes in multilane roads and testing of the system with actual production sensors present on the vehicle. Finally, an implementation of this system using global landmarks will eliminate the need for a road map and will create a more viable navigation strategy in the real world.

References

[1] "Traffic safety facts 2008 data," Nat. Highway Traffic Safety Admin., Washington, DC, USA, DOT HS 811 162, 2008. [Online]. Available: http://www-nrd.nhtsa.dot.gov/Pubs/811162.PDF

[2] J. Kibbel, W. Justus, and K. Furstenberg, "Lane estimation and departure warning using multilayer laserscanner," in *Proc. IEEE Intell. Transp. Syst.*, 2005, pp. 607–611.

[3] K. Dietmayer *et al.*, *Advanced Microsystems for Automotive Applications 2005*. Berlin, Germany: Springer-Verlag, 2005.

[4] A. von Reyher, A. Joos, and H. Winner, "A lidar-based approach for near range lane detection," in *Proc. IEEE Intell. Veh. Symp.*, 2005, pp. 147–152.

[5] C. R. Jung and C. R. Kelber, "A lane departure warning system based on a linear-parabolic lane model," in *Proc. IEEE Intell. Veh. Symp.*, 2004, pp. 891–895.
[6] Y. Feng, W. Rong-ben, and Z. Rong-hui, "Based on digital image lane edge detection and tracking under structure environment for autonomous vehicle," in *Proc. IEEE Int. Autom. Logist. Conf.*, 2007, pp. 1310–1314.
[7] P.-Y. Hsiao and C.-W. Yeh, "A portable real-time lane departure warning system based on embedded calculating technique," in *Proc. IEEE 63rd VTC-Spring*, 2006, vol. 6, pp. 2982–2986.
[8] A. Takahashi and Y. Ninomiya, "Model-based lane recognition," in *Proc. IEEE Intell. Veh. Symp.*, 1996, pp. 201–206.
[9] E. D. Dickmanns and B. D. Mysliwetz, "Recursive 3-D road and relative ego-state recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 2, pp. 199–213, Feb. 1992.
[10] D. Khosla, "Accurate estimation of forward path geometry using two-clothoid road model," in *Proc. IEEE Intell. Veh. Symp.*, 2002, vol. 1, pp. 154–159.
[11] D. A. Schwartz, "Clothoid road geometry unsuitable for sensor fusion clothoid parameter sloshing," in *Proc. IEEE Intell. Veh. Symp.*, 2003, pp. 484–488.
[12] I. Miller and M. Campbell, "Particle filtering for map-aided localization in sparse GPS environments," in *Proc. IEEE Int. Conf. Robotics and Automation ICRA 2008*, 2008, pp. 1834–1841.
[13] H. Li, F. Nashashibi, and G. Toulminet, "Localization for intelligent vehicle by fusing mono-camera, low-cost GPS and map data," in *Proc. 13th IEEE ITSC*, 2010, pp. 1657–1662.
[14] A. Vu, A. Ramanandan, A. Chen, J. Farrell, and M. Barth, "Real-time computer vision/DGPS-aided inertial navigation system for lane-level vehicle navigation," *IEEE Trans. Intell. Transp. Syst.*, vol. 13, no. 2, pp. 899–913, Jun. 2012.
[15] Y. Zhao, B. Song, and J. Li, "A map matching algorithm in GPS-based car navigation system," in *Proc. 3rd Int. Conf. IIHMSP*, 2007, vol. 1, pp. 77–80.
[16] C. Bahlmann, M. Pellkofer, J. Giebel, and G. Baratoff, "Multi-modal speed limit assistants: Combining camera and GPS maps," in *Proc. IEEE Intell. Veh. Symp.*, 2008, pp. 132–137.
[17] C. Rose and D. M. Bevly, "Lane level localization with camera and inertial measurement unit using an extended Kalman filter," M.S. thesis, Auburn Univ., Auburn, AL, USA, 2010.
[18] W. Niblack, *An Introduction to Digital Image Processing*. Englewood Cliffs, NJ, USA: Prentice Hall, 1986.
[19] J. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-8, no. 6, pp. 679–698, Nov. 1986.
[20] P. V. C. Hough, "Machine analysis of bubble chamber pictures," in *Proc. Int. Conf. High Energy Accel. Instrum.*, 1959, pp. 554–558.
[21] J. Britt, "Lane detection calibration and attitude determination with a multilayer Lidar for vehicle safety systems," M.S. thesis, Auburn Univ., Auburn, AL, USA, 2010.
[22] J. Allen, "Use of vision sensors and lane maps to aid GPS-INS navigation," M.S. thesis, Auburn Univ., Auburn, AL, USA, 2011.
[23] J. Cheng, J. Farrell, L. Yu, and E. Thomas, "Aided integer ambiguity resolution algorithm," in *Proc. PLANS*, 2004, pp. 740–745.
[24] R. Stengel, *Optimal Control and Estimation*. New York, NY, USA: Dover, 1994.
[25] I. Rhee, M. Abdel-Hafez, and J. Speyer, "Observability of an integrated GPS/INS during maneuvers," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 40, no. 2, pp. 526–535, Apr. 2004.

Christopher Rose received the B.S. and M.S. degrees in electrical engineering from Auburn University, Auburn, AL, USA, in 2007 and 2010, respectively. He is currently working toward the Ph.D. degree in the GPS and Vehicle Dynamics Laboratory, Auburn University.

His research interests include video processing applications and sensor fusion in ground vehicles.

Jordan Britt received the B.S. and M.S. degrees in electrical engineering from Auburn University, Auburn, AL, USA, in 2008 and 2010, respectively. He is currently working toward the Ph.D. degree in the GPS and Vehicle Dynamics Laboratory (GAVLAB), Auburn University.

He is also currently a Research Assistant with GAVLAB. His research interests include sensor fusion and autonomous vehicle navigation and control.

John Allen received the B.S. and M.S. degrees in mechanical engineering from Auburn University, Auburn, AL, USA, in 2007 and 2010, respectively.

He is a Software Engineer with Aurora Flight Sciences, Manassas, VA, USA, where he creates and maintains unmanned aerial vehicle simulators for software- and hardware-in-the-loop testing infrastructures. His research interests include vehicle dynamics, multisensor fusion for vehicle navigation, and vehicle control/autonomy.

David Bevly received the B.S. degree from Texas A&M University, College Station, TX, USA; the M.S. degree from Massachusetts Institute of Technology, Cambridge, MA, USA; and the Ph.D. degree in mechanical engineering from Stanford University, Stanford, CA, USA.

He is the Director of the GPS and Vehicle Dynamics Laboratory, Auburn University, Auburn, AL, USA, which focuses on modeling, navigation, and control of vehicles.