Terrain-Based Road Vehicle Localization on Multi-Lane Highways

Adam J. Dean and Sean N. Brennan*

Abstract— **This work develops an algorithm for estimating the lateral lane index of road vehicles on multi-lane roadways by correlating vehicle attitude measurements to terrain maps of the individual lanes of travel. To localize a vehicle, a Bayesian belief algorithm and a particle filter algorithm are described and applied off-line using data collected from two lanes along a local highway. Results demonstrate that terrain-based algorithms are capable of measuring lane index. Because these measurements are immune to lighting conditions, this solution is a good complement to existing lane-detection camera systems.**

I. Introduction

In order to enable several safety and efficiency features of vehicle travel and automated highways, accurate vehicle position is required. The Global Positioning System (GPS) is the primary means of vehicle localization today; however, it has several shortcomings including slow update rates, satellite signal loss, multipath errors, and is vulnerable to noise and malicious hacking; thus, GPS is not a good stand-alone solution for vehicle localization.

Alternative methods, including GPS fused with wheel odometry [1] or inertial measurements have been used to continue vehicle global positioning during satellite signal loss, however these methods require initialization from GPS, and suffer from wheel-slip errors or sensor drift that can accumulate without bound until corrected when the satellite signal is restored. Vision [2] and LIDAR (Light Detection And Ranging) [3] feature matching can be used independently of GPS; however, they are both financially and computationally expensive. Beacon networks [4] can be used to triangulate vehicle position, but the infrastructure to do so does not exist and would be very expensive to develop.

Vision algorithms have been used for global terrain-based vehicle localization [5], but vision processing is more commonly used for determining relative lateral position along the highway. Many GPS-free vision-based approaches exist for lane detection [6] and lane indexing [7] and are advantageous in that they do not require an external infrastructure or previously recorded map data; however, vision systems do not perform well in situations with varying or low lighting conditions [8].

Terrain-based localization is another alternate to GPS as a global reference method and consists of matching in-vehicle sensor measurements to an on-board terrain map. The method of correlating height measurements in aircraft [9], missile [10], and even underwater systems [11] has been used extensively and was the primary localization method before GPS was available. These results suggest that a vehicle can be globally localized by correlating in-vehicle attitude measurements with a terrain profile map assuming that the terrain along the vehicles path has been previously mapped and is available on-board the vehicle. This is the method employed herein.

This paper investigates the use of vehicle attitude measurements as a means of rapid detection of lane position, develops an algorithm to achieve this detection, and presents experimental results from a highway implementation. This paper is outlined as follows: Section 2 discusses the feasibility of our terrain-based approach to discriminate lane index. Section 3 describes the method of collecting the terrain data for off-line simulation. Section 4 uses a simple Bayesian belief algorithm to estimate lane index. Section 5 describes and implements a particle filter algorithm and presents results for lane indexing along a multi-lane highway. Conclusions summarize the main results and future opportunities of this study.

II. Lateral Terrain-Based Localization

Previous work [12] has shown the preliminary feasibility analysis of the terrain-based method and demonstrated the ability to localize a vehicle longitudinally along the Thomas D. Larson Transportation Institute (LTI) Test Track to within 10 cm, the accuracy of the terrain map. Further study [13] has shown the ability to localize a vehicle longitudinally along an interstate roadway terrain map over 60 km in length to meter resolution. While these results are promising, they did not include lateral position, or even the simpler task of detecting lane index. Hence, this study was initiated.

Our approach to lateral positioning uses a terrain-based method of matching the vehicle attitude changes to a previously recorded and lane-correlated dataset. Roadway variations that can easily be measured from within the vehicle include the roads lateral profile, or superelevation, and the road grade, or slope of the road fore/aft. Estimates of these two road features are obtainable from in-vehicle roll and pitch measurements, respectively.

As previously discussed, GPS suffers from several sources of error which make lateral positioning using GPS unreliable. An example of these errors is shown in Figure 1, which shows the GPS/INS-mapped position of the vehicle as it was traveling down the lanes of a highway. The GPS/INS system was a defense-grade Novatel Span system factory integrated with a Honeywell Ring Laser Gyro IMU (HG1700), a

* Corresponding author

A. Dean is with the Department of Mechanical Engineering, Brigham Young University-Idaho, 525 S. Center St., Austin 164, Rexburg, ID 83460, deana@byui.edu, (208) 496-3833

S. Brennan is with the Department of Mechanical Engineering, Penn State University, 318 Leonhard Building, University Park, PA 16802, sbrennan@psu.edu, (814) 863-2430

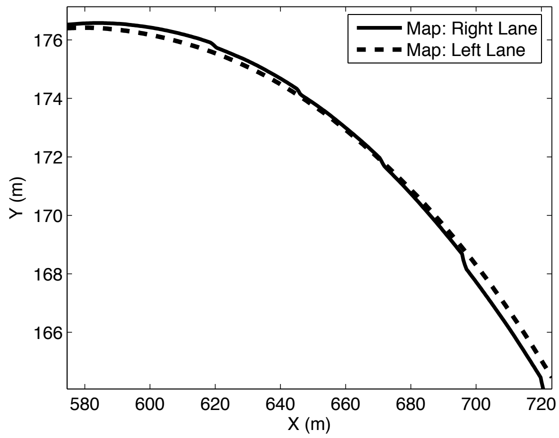

Fig. 1. Example of the difficulty of lane detection with GPS when the measured lane positions cross.

Fig. 2. Side view of the test vehicle shown with the wheel odometer and GPS antenna.

very high performance (and cost prohibitive) sensing system compared to commercial grade sensors. The errors in the orientation angles are 0.013, 0.013 and 0.04 degrees (one sigma) for the roll, pitch and the yaw angles respectively. With several satellites in view but without differential correction the accuracy of GPS is on the order of one meter, during low satellite visibility this accuracy will degenerate further. One can observe the effects of this resolution limit in practice: the resulting lane positions cross each other and there are signal discontinuity effects in the right-hand lane where the vehicle position seems to jump. Such errors are avoidable if differential GPS (DGPS) is available, but such capability cannot be assumed due to subscription costs of DGPS, and limited availability of such correction signals worldwide.

III. Description of Experiment

In order to test the feasibility of terrain-based lateral vehicle positioning, a 1999 Jeep Grand Cherokee was equipped with a wheel odometry sensor and the NovAtel SPAN GPS-IMU system, as shown in Figure 2. We drove the vehicle along one lane of highway 322 in State College, PA without changing lanes for several kilometers while recording the odometry, vehicle position, and vehicle attitude. We then repeated the route along the second lane of the highway, and then a third time but with lane changes about every kilometer. During the lane changing maneuvers, the lateral position, or lane index, was also recorded as truth and used to denote the right-hand lane as lane 1 and the left-hand lane as lane 2. Thus the true lateral vehicle position is equal to 1 or 2 while in the corresponding lane. A value of 1.5 is chosen to denote when the vehicle is moving between lanes in either direction. The true position measured from DGPS was also recorded.

The resulting terrain maps for each lane were decimated at 0.5 meters and filtered at 0.1 cycles/meter [12]. An overhead view of the map, as well as the pitch and roll measurements for each lane and the lane-changing data set is shown in Figure 3. The vehicle responses for each lane as well as the response during the lane change maneuver are shown. The dark shaded region depicts when the vehicle was in lane 2, the left-hand lane, during the lane-changing route. The lightly shaded region shows when the vehicle is in transition between lanes, and the unshaded region is when the vehicle is in lane 1, or the right-hand lane.

From Figure 3 it is evident that there is very little variation in the vehicle's pitch response between the lanes. While this may have an adverse affect in differentiating the lane position of the vehicle, it does show that longitudinal positioning may be possible along any lane using only the pitch response map from a single lane.

IV. Lane Indexing Using A Bayesian Belief Algorithm

In order to demonstrate a simple terrain-based lane index estimator, a Bayesian belief algorithm is used first to estimate lateral position. Because lane indexing is a discrete representation of the vehicle's lateral position, a Bayesian algorithm is a very simple approach to estimating the lane index. The Bayesian belief algorithm is essentially a method of estimating lane position from the likelihood that the vehicle is located within each lane. Under some general assumptions about the vehicle's likelihood of changing lanes, the belief is calculated by comparing the attitude measurements with the terrain map.

The lateral vehicle localization problem is simplified by assuming the vehicles longitudinal position along the roadway is known; this assumption is not restrictive: previous work ([12], [13]) have already demonstrated accurate longitudinal vehicle positioning. This assumption allows the lateral and longitudinal estimation problems to be decoupled by first estimating the longitudinal vehicle position until the vehicle is localized sufficiently, then the lateral vehicle position can be estimated.

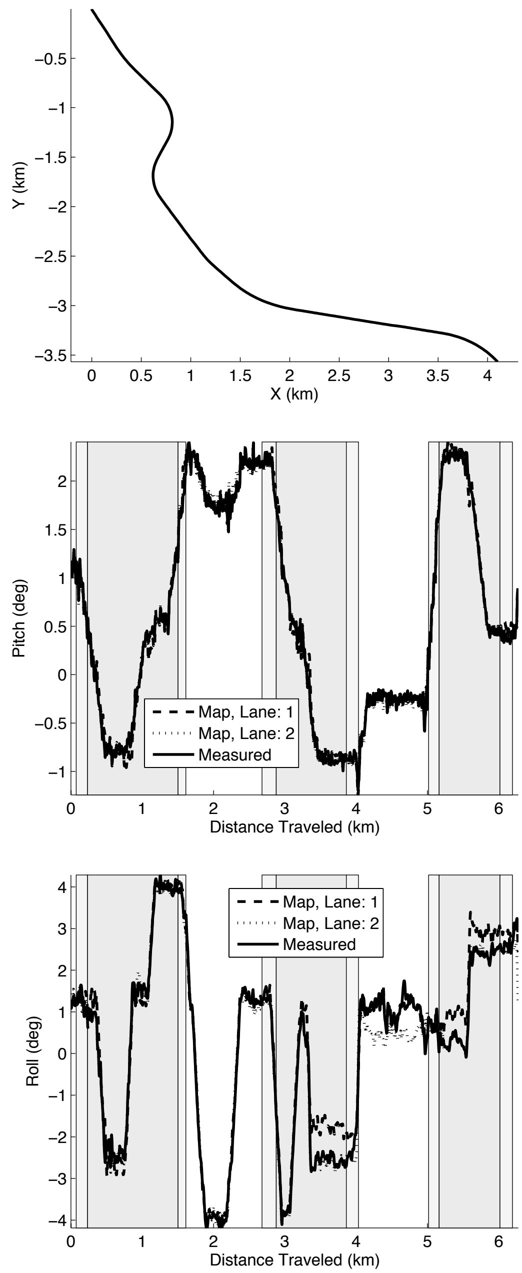

Fig. 3. Overhead view of the terrain data, as well as the measured pitch and roll maps of each lane and during the lane-changing maneuver.

The Bayesian approach is initialized by setting the belief of the vehicle's lane position to be equal for each lane: $Bel(y_{1,0}) = 0.5$ and $Bel(y_{2,0}) = 0.5$ where $Bel(y_{i,k})$ is the likelihood that the vehicle is located in lane index i at time k. The belief of the vehicle's lane position is updated every $dX = 5$ meters of vehicle travel, as measured by a wheel odometer, by repeating the following steps:

First, knowing that the vehicle is between lanes about only 10% of the route, as calculated from post-processing the actual lane-maneuver data, the probabilities of the motion model are estimated to be:

$$P(y_{1,k}|u_{k-1}, y_{1,k-1}) = 0.9 \tag{1}$$
$$P(y_{2,k}|u_{k-1}, y_{2,k-1}) = 0.9 \tag{2}$$
$$P(y_{1,k}|u_{k-1}, y_{2,k-1}) = 0.1 \tag{3}$$
$$P(y_{2,k}|u_{k-1}, y_{1,k-1}) = 0.1 \tag{4}$$

where $P(y_{2,k}|u_{k-1}, y_{1,k-1})$ is the probability that a vehicle positioned in lane 1 at time $k-1$ will be in lane 2 after dX meters of travel. Because these probabilities were calculated in post-processing from actual lane-maneuver data, the likelihood of changing lanes in normal, everyday driving may differ significantly.

It is beyond the scope of this study to determine "good" values for the lane-change probability, but there are several methods of estimating these numbers. For example, the likelihood can be based upon the vehicle's proximity to an exit, or the time since the last lane maneuver. Regardless of the method, the objective of this research is to demonstrate the feasibility of terrain-based localization using a Bayesian belief algorithm. Thus, these estimates are held constant, unlike the methods mentioned above which might update these probabilities at every iteration.

Subsequently, the belief of each state $y_{i,k}$ is updated according to the motion model by:

$$Bel(y^*_{i,k}) = \sum_{j=1}^{2} P(y_{i,k}|u_{k-1}, y_{j,k-1}) \cdot Bel(y_{j,k-1}) \tag{5}$$

which, with only two lanes of travel, is expanded to

$$\begin{aligned} Bel(y^*_{1,k}) &= P(y_{1,k}|u_{k-1}, y_{1,k-1}) \cdot Bel(y_{1,k-1}) \quad (6) \\ &+ P(y_{1,k}|u_{k-1}, y_{2,k-1}) \cdot Bel(y_{2,k-1}) \\ Bel(y^*_{2,k}) &= P(y_{2,k}|u_{k-1}, y_{1,k-1}) \cdot Bel(y_{1,k-1}) \quad (7) \\ &+ P(y_{2,k}|u_{k-1}, y_{2,k-1}) \cdot Bel(y_{2,k-1}) \end{aligned}$$

Second, the belief of each state $y^*_{i,k}$ is updated using one of two measurement models. The first uses only pitch measurements:

$$Bel(\hat{y}_{i,k}) = Bel(y^*_{i,k}) \cdot exp\left(-\frac{(\theta_a - \theta_y)^2}{2 \cdot R_p}\right) \tag{8}$$

the second uses only roll measurements:

$$Bel(\hat{y}_{i,k}) = Bel(y^*_{i,k}) \cdot exp\left(-\frac{(\phi_a - \phi_y)^2}{2 \cdot R_r}\right) \tag{9}$$

where R_p and R_r are the measurement noise variances in pitch and roll, θ_a and ϕ_a are the measured pitch and roll of the vehicle, and θ_y and ϕ_y are the pitch and roll of lane y according to the estimated position along the terrain map.

Third, the beliefs are normalized according to:

$$Bel(y_{i,k}) = \frac{Bel(\hat{y}_{i,k})}{\sum_{j=1}^{2} Bel(\hat{y}_{j,k})} \tag{10}$$

Fourth, the estimated lane position is then assumed to be the lane with the highest belief. The error of the lateral position estimate is calculated as the absolute difference

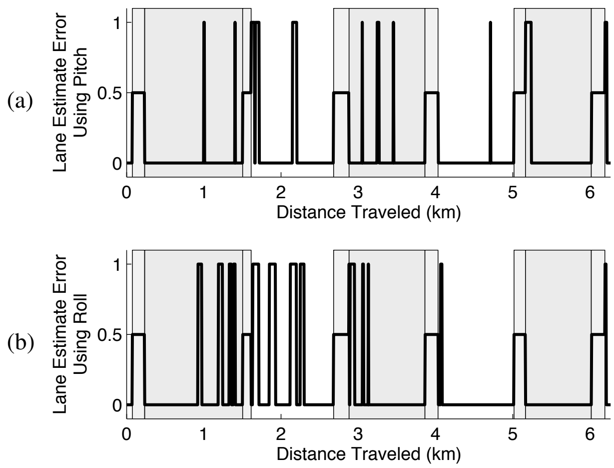

Fig. 4. Lane index estimate error from the Bayesian belief algorithm using only pitch (a) and only roll measurements (b).

TABLE I

BAYESIAN BELIEF RESULTS USING ONLY THE PITCH MEASUREMENT.

	Predicted Position		
Measured	Lane 1	Lane 2	Error
Lane 1	402	36	8.2%
Mid-Lane	115	78	100%
Lane 2	25	597	4.0%

between the estimated lateral position and the actual vehicle position of 1, 1.5, or 2 as discussed previously. Because we restrict the continuous lateral position to ordinal positions of either lanes 1, 1.5, or 2, while updating the belief of only lanes 1 and 2, some errors in this algorithm are expected, particularly during lane change events.

The Bayesian algorithm is implemented with $R_p = R_r = 0.1$ deg^2, is iterated every $dX = 5$ meters of travel, and uses only the pitch measurements (Eq. (8)) to correlate lane position. The resulting lateral estimate errors are shown in Figure 4a where it is evident that the Bayesian algorithm was capable of estimating the lane index for the majority of the route regardless of the little difference between the pitch maps of each lane.

The algorithm is repeated while using only the roll measurements to estimate lane index, Eq. (9), as shown in Figure 4b, demonstrating that the Bayesian algorithm was not as accurate when using the roll measurements as when correlating pitch. The results are also summarized in Tables I and II, demonstrating that the Bayesian algorithm exhibited only 8% error when using the pitch measurement.

From Figure 4 it is evident that regardless of which weighting function was used, the Bayesian approach is a simple, and fairly accurate, method of estimating lane index.

TABLE II

BAYESIAN BELIEF RESULTS USING ONLY THE ROLL MEASUREMENT.

	Predicted Position		
Measured	Lane 1	Lane 2	Error
Lane 1	373	65	14.8%
Mid-Lane	89	104	100%
Lane 2	49	573	7.9%

Although the algorithm required an assumed probability of lane change maneuvers, Eq. (1), these values can be derived from a variety of methods.

V. LANE INDEXING USING A PARTICLE FILTER

In contrast to the Bayesian belief algorithm, a particle filter algorithm can be a more robust approach to estimate lateral position and detect lane-change maneuvers from in-vehicle attitude measurements instead of using a pre-estimated probability. In this section, the particle filter algorithm is extended from the algorithm used in [12] to include a lateral position estimate and, as discussed previously, is simplified by assuming the vehicle's longitudinal position along the roadway is already known.

A. Detecting Lane-Change Maneuvers

As opposed to using an estimated likelihood of a lane change maneuver, as was used in the Bayesian Belief algorithm, it is assumed that the vehicle's yaw angle is available and measured in order to detect a lane change. This can be accomplished through a variety of sensors and methods; for example, using a digital compass or integrating steering input to estimate vehicle yaw angle. The latter method, as well as several others, were attempted in order to detect a lane-change maneuver using the steering input using the present data set, but the "play" in the steering column of the instrumented vehicle was too large to acquire an adequate yaw measurement. Instead of using the steering input, the yaw angle is collected using the GPS/IMU system already used to measure pitch and roll. The measured yaw angle of each lane is compared with the yaw angle of the lane-changing maneuver as shown in Figure 5. This method may be somewhat restrictive because GPS/IMU systems are not widespread, hence yaw angle measurements are not readily available in most production vehicles. Nonetheless, this method is still useful as lane index is generally not clear even with an accurate GPS/IMU system, as shown in Figure 1.

The variations in the yaw angle of the lane-changing maneuver from the yaw angle of the individual lanes can be seen in Figure 5, indicating a possible metric for determining a lane change. Subtracting the yaw measurements of the mapped lanes from the measured yaw angle during the lane-changing maneuver is shown in Figure 6. It can be seen that using the difference in yaw can be used to estimate when

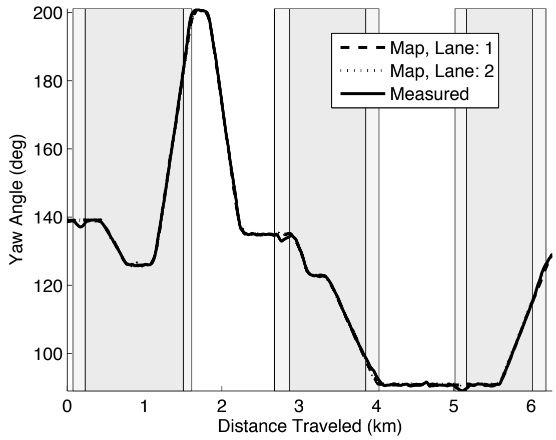

Fig. 5. In-vehicle yaw measurement of each lane and lane change maneuver.

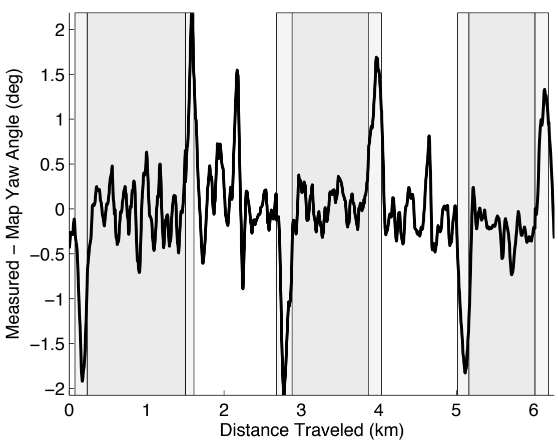

Fig. 6. The difference in yaw angle between the lane-changing maneuver and the mapped yaw angle of the true lane position.

a lane-change maneuver occurs. This ability of measuring a change in lane position can be used as additional information in the particle filter algorithm to shift the particles laterally between the parallel terrain maps of each travel lane.

B. Lane Indexing Particle Filter Algorithm

The lateral particle filter algorithm is modified from that described in [12] in order to include the lateral position estimate. The algorithm is initialized laterally (Y_p^0) by dividing N number of particle positions evenly between the two lanes of travel. The particles are then initialized longitudinally (X_p^0) about the estimated longitudinal position using a Gaussian distribution with a standard deviation of one meter. This assumes that the vehicle has been localized longitudinally with a standard error of one meter with a likelihood of existing in either lane. The particle filter is then implemented by iterating the following:

First, the longitudinal position estimates or particles, denoted by X_p, at iteration k, are updated from the previous estimate using:

$$X_p^k = X_p^{k-1} + dX + w \tag{11}$$

where dX is the longitudinal vehicle travel as measured by wheel odometry, and w is Gaussian white process noise of variance Q_x. The lateral position estimates, denoted by Y_p, are updated by

$$Y_p^k = Y_p^{k-1} + K \cdot (\psi_a - \psi_p) + w_y \tag{12}$$

where ψ_a is the actual yaw angle of the vehicle, ψ_p is the predicted yaw angle, K is a gain, and w_y is Gaussian white process noise of variance Q_y added to increase randomness and prevent degeneracy. The predicted yaw angle ψ_p is calculated every iteration from the yaw map at the current estimated lane position. After the particle positions are updated, the particle's lateral positions are rounded to the closest lane. Thus, the particles are free to roam in the longitudinal direction of the roadway according to the measured odometry, but restricted in the lateral direction to exist only along the mapped lanes; hence, large errors are expected to occur when the vehicle traverses between lanes.

Second, the weights of the position particles are updated by measuring the in-vehicle pitch and roll and comparing it to the particle's pitch and roll using a particle weighting function. The first weighting function uses only the pitch measurement to weight the particles:

$$q_i^k = exp\left(-0.5 \cdot R_p^{-1} \cdot \left(\theta_a^k - \theta_{p,i}^k\right)^2\right) \tag{13}$$

the second uses only the roll measurement:

$$q_i^k = exp\left(-0.5 \cdot R_r^{-1} \cdot \left(\phi_a^k - \phi_{p,i}^k\right)^2\right) \tag{14}$$

where R_p and R_r are the measurement noise variances in pitch and roll, θ_a and ϕ_a are the measured pitch and roll of the vehicle, $\theta_{p,i}$ and $\phi_{p,i}$ are the i^{th} particle's pitch and roll corresponding to its position along the terrain map. A normalizing factor η equal to the sum of the particle weights is then multiplied to q_i^k after the weights are calculated.

Third, the particles are re-sampled to remove particles with little weight and duplicate particles with high weight according to:

$$\begin{aligned}
&c = cumsum\left(q^k\right)\\
&u_1 = rand(1) \cdot N^{-1}\\
&i = 1\\
&\text{for}\ \ j = 1...N\\
&\qquad u_j = u_1 + (j-1) \cdot N^{-1}\\
&\qquad \text{while}\ \ u_j > c_i\\
&\qquad\qquad i = i + 1\\
&\qquad \text{end}\\
&\qquad X_{p,j}^k = X_{p,i}^k\\
&\qquad Y_{p,j}^k = Y_{p,i}^k\\
&\text{end}
\end{aligned} \tag{15}$$

where $rand(1)$ is an evenly distributed random number in

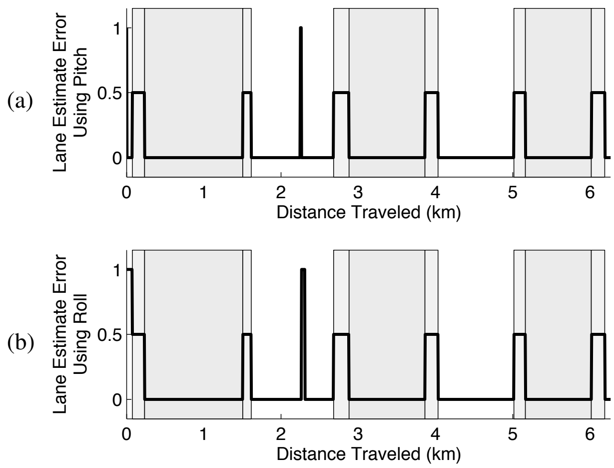

Fig. 7. Lane index estimate error from the PF algorithm using only pitch (a) and only roll measurements (b).

$[0, 1]$ and $cumsum$ is the cumulative sum.

Fourth, at every iteration the vehicle's position is estimated as the mean position of the particle estimates. The error of the lateral position estimate is calculated as the absolute difference between the estimated lateral position and the measured vehicle lane index of 1, 1.5, or 2 as discussed previously.

C. Lane Indexing Results

Using the terrain data shown in Figure 3, the lane indexing particle filter algorithm is implemented using $dX = 5$ meters, $R_p = R_r = 0.1$ deg^2, $Q_x = (0.01 \cdot dX)^2$ m^2, $Q_y = 0.01$ and $K = -0.5$. Because the algorithm displayed such accurate results, the number of particles could be reduced to only $N = 10$ particles, through trial-and-error to determine the minimal amount needed for accurate positioning. This is in contrast to findings in [12] where $N = 1000$ particles/mile was calculated to be necessary for accurate longitudinal positioning.

The results of the algorithm using $N = 10$ particles are shown in Figure 7 where the lane index is estimated using the pitch or roll measurement weighting functions, Eqs. (13) and (14), respectively. The results demonstrate a nearly perfect lane estimate for using either pitch or roll measurements. This method clearly indicates the capability of determining lane index along a multi-lane roadway, given the roll, pitch, and orientation or yaw maps of each lane.

The results also indicate that the particle filter algorithm was more accurate than the Bayesian belief algorithm and did not require an estimated probability of the vehicle performing a lane-change event. However, the particle filter algorithm required an additional sensor in order to detect lane change maneuvers.

VI. CONCLUSIONS

A terrain-based approach to lateral vehicle positioning along a multi-lane highway has been presented. The particle filter algorithm was shown to be an accurate method of estimating the lane position while in lane and, in contrast with vision systems, the terrain-based approach provides an accurate means of lane indexing that is independent of lighting conditions or lane markers. A Bayesian belief algorithm was also demonstrated to be a simple means of estimating lane index.

This work is ongoing, but additional items of study are obvious, including a method of accurately estimate lateral position within each lane. The vehicle's exact lateral position relative to a lane was not measurable with high accuracy in this study, due to the limitations of GPS as noted earlier and lack of access to lane detection camera systems. An interesting follow up to this study would be to measure relative lane position using lane marker detection to discern the true lateral accuracy of this lane estimation method, and to determine whether a significant portion of the error observed in lane index estimation is instead error in lane keeping by the driver.

Additionally, consideration of lower-grade vehicle sensors and testing across multiple vehicles are also important considerations that deserve further attention.

REFERENCES

[1] M. E. E. Najjar and P. Bonnifait, "A road-matching method for precise vehicle localization using belief theory and kalman filtering," *Auton. Robots*, vol. 19, no. 2, pp. 173–191, 2005.

[2] T. Oskiper, Z. Zhu, S. Samarasekera, and R. Kumar, "Visual odometry system using multiple stereo cameras and inertial measurement unit," in *Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on*, Minneapolis, MN, USA, Jun. 17–22, 2007, pp. 1–8.

[3] M. Jabbour, P. Bonnifait, and V. Cherfaoui, "Enhanced local maps in a GIS for a precise localisation in urban areas," in *Intelligent Transportation Systems, 2006. Proceedings. 2006 IEEE*, Toronto, Ont., 2006, pp. 468–473.

[4] C. D. McGillem and T. S.Rappaport, "A beacon navigation method for autonomous vehicles," *Vehicular Technology, IEEE Transactions on*, vol. 38, no. 3, pp. 132–139, 1989.

[5] V. Gupta and S. Brennan, "Terrain-based vehicle orientation estimation combining vision and inertial measurements," *J. Field Robot.*, vol. 25, no. 3, pp. 181–202, 2008.

[6] M. Bertozzi and A. Broggi, "GOLD: a parallel real-time stereo vision system for genericobstacle and lane detection," *IEEE Transactions on Image Processing*, vol. 7, no. 1, pp. 62–81, Jan. 1998.

[7] R. Aufrere, R. Chapuis, and F. Chausse, "A model-driven approach for real-time road recognition," *Machine Vision and Applications*, vol. 13, no. 2, pp. 95–107, November 2001. [Online]. Available: http://www.springerlink.com/content/pbgr6bcvbvg64yuh/

[8] J. McCall and M. Trivedi, "Video-based lane estimation and tracking for driver assistance: survey, system, and evaluation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 7, no. 1, pp. 20–37, 2006.

[9] F. Gustafsson, F. Gunnarsson, N. Bergman, U. Forssell, J. Jansson, R. Karlsson, and P. J. Nordlund, "Particle filters for positioning, navigation, and tracking," *Signal Processing, IEEE Transactions on [see also Acoustics, Speech, and Signal Processing, IEEE Transactions on]*, vol. 50, no. 2, pp. 425–437, Feb. 2002.

[10] J. P. Golden, "Terrain contour matching/TERCOM/- A cruise missile guidance aid," *Image processing for missile guidance*, pp. 10–18, 1980.

[11] S. Williams, G. Dissanayake, and H. Durrant-Whyte, "Towards terrain-aided navigation for underwater robotics," *Advanced Robotics*, vol. 15, no. 5, pp. 533–549, 2001.

[12] A. Dean, R. Martini, and S. Brennan, "Terrain-based road vehicle localization using particle filters," in *Proceedings. 2008 American Control Conference*, Jun. 11–13, 2008.

[13] A. Dean, P. Vemulapalli, and S. Brennan, "Highway evaluation of terrain-aided localization using particle filters," in *Proceedings. 2008 Dynamic Systems Control Conference, submitted*, Oct. 20–22, 2008.