2008 American Control Conference
Westin Seattle Hotel, Seattle, Washington, USA
June 11-13, 2008

WeA07.4

Terrain-Based Road Vehicle Localization Using Particle Filters

Adam J. Dean, Ryan D. Martini, and Sean N. Brennan

***Abstract*— This work develops a real-time algorithm to localize a vehicle in the direction of travel without the use of GPS. The inputs to the algorithm include a terrain map of road grade and pitch measurements from an in-vehicle pitch sensor. Localization is achieved in real-time using a particle filter described in detail in this work. Simulations and experiments at The Pennsylvania Transportation Institute test track are used to demonstrate the algorithm, observe the speed of convergence, and to determine key parameters for practical implementation. The results indicate that the method can quickly localize a vehicle with one-meter accuracy or better.**

I. Introduction

For reasons of safety and efficiency, there is a great deal of interest in localizing road vehicles. Today, the Global Positioning System (GPS) serves as the primary means to determine vehicle position. However, due to poor GPS signal reception in some locations, the ease of jamming a GPS signal in battlefield operation, and the need for sensor redundancy in vehicle automation and driver assist applications, there has been great interest in localization technologies independent of GPS.

Several methods have been used to localize a vehicle without GPS or during short GPS outages including fusion of GPS with odometry [1], inertial measurements, vision [2], laser scans [3], or using a network of beacons [4]. Of particular relevance to this work is the fusion of GPS with map data. A demonstration of this capability was accomplished in real-time [1] where a Kalman filter was used to combine GPS data with odometry measurements. A methodology called Belief Theory was used to correlate an estimated vehicle position to a given digital road map, and this method is able to correctly map the vehicle's position in the absence of GPS if the position was first correlated correctly before the GPS outage. Many other approaches to vehicle localization exist, but like this example, most are designed to improve upon GPS accuracy or maintain localization during GPS outages and are thus not completely independent of GPS. The purpose of this research is to study methods of localizing a vehicle completely without a GPS device.

Similar to the work done in [5] where an aircraft's elevation profile is matched to a digital elevation map, this work demonstrates the use of a terrain map for real-time vehicle localization with the goal to obtain sub-meter position resolution of a vehicle's longitudinal position. Similar

A. Dean and S. Brennan are with the Department of Mechanical Engineering, Penn State University, 318 Leonhard Building, University Park, PA 16802 ajdean,sbrennan@psu.edu

R. Martini is a Hybrid Calibration Specialist with GM Powertrain R&D, 3300 General Motors Road, Milford MI, 48380, Bldg 16: 16D-2426

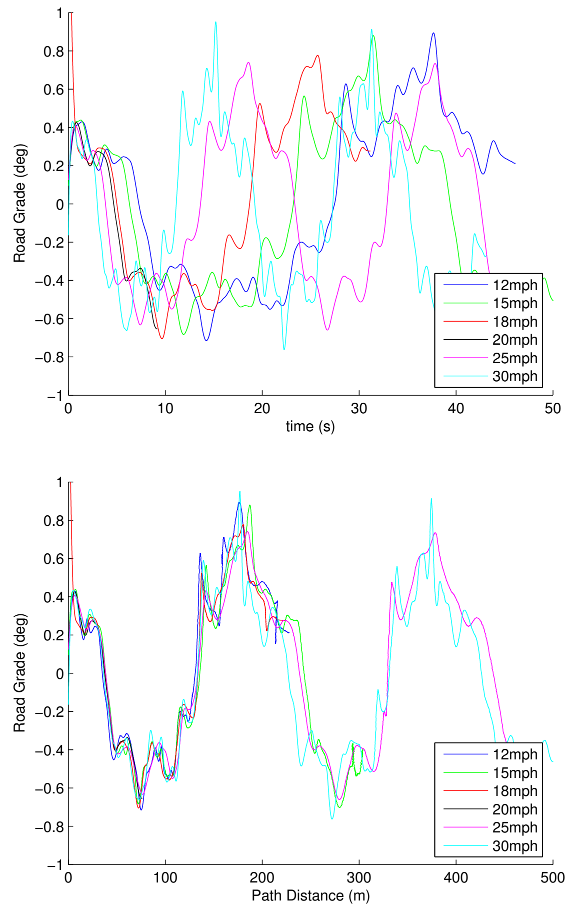

Fig. 1. Measured road grade as a function of time and distance at various vehicle speeds on a circular handling area [8].

terrain-aided applications include missile guidance systems [6] and underwater robotics [7].

It is assumed that the lane of travel has been previously mapped and that on-vehicle storage of the resulting terrain information is available. The first assumption is quite realistic given the large number of ongoing research projects focused on mapping terrain, whereas the second assumption is increasingly valid given the exponentially decreasing costs of data storage and recent integration of similar on-vehicle map databases into commercial products, for example the "TomTom" navigation devices.

In this study, several terrain characteristics were considered for vehicle localization: road height changes, road grade changes (derivative of height), and superelevation changes. The last two are the primary focus of this work

978-1-4244-2079-7/08/$25.00 ©2008 AACC.

because road grade correlates very strongly to in-vehicle pitch measurements [9], and superelevation correlates to vehicle roll. Both measurements are not perfect due to vehicle dynamics; for example, the vehicle's pitch response acts as a low-pass filter to road grade changes [8]. However, as this work demonstrates, this filtering effect is not significant enough to filter out measurements of roadway disturbances useful for determining vehicle position.

To demonstrate that a vehicle's pitch measurement gives repeatable terrain-correlated responses, a test vehicle was driven in a circular trajectory on a handling area of the Pennsylvania Transportation Institute (PTI) test track, one of the flattest portions of this vehicle testing facility. In Figure 1, one can see two plots: first, the pitch measurements versus time for different velocities, and second, the same data versus distance traveled using the same starting point. One can clearly see from the second plot that pitch correlates to position. Examining the response closely, the changes in observed pitch are due to the very small slope added to the handling area to allow water drainage. Also, because similar responses are measured for a vehicle traveling at various speeds, one can conclude that the "wheelbase filtering" effect [10] of the vehicle's pitch dynamics is relatively minor.

These plots suggest that a vehicle can be localized by correlating a previously-mapped roadway with a vehicle's pitch response history transformed into a spatial pitch measurement. This study tests this localization methodology and suggests an algorithm for fast localization. The remainder of this paper is organized as follows: Section 2 presents a preliminary analysis of feasibility including an analysis of the terrain features being correlated. Section 3 introduces a particle filter algorithm to achieve real-time position estimation from terrain disturbances. Section 4 presents the results of this filter tested using test-track data. Finally, a Conclusions section summarizes the main findings of this work.

II. Preliminary Feasibility Analysis

Because the pitch of a vehicle responds differently with respect to speed over the same terrain, the so-called wheelbase filtering effect, it was initially unclear whether terrain disturbances in pitch would be repeatable across different speeds and for roadways other than a circular track section. To test repeatability, we drove a test vehicle - a 1992 Mercury Tracer station wagon - at 2.2, 17.9, and 29 m/s (5, 40, and 65 mph) on the roadway section of the test track, recording pitch. By comparing the Power Spectral Density (PSD) of the pitch responses in spatial frequencies, shown in Figure 2, one observes the similarity of low-frequency content with frequency measured in cycles-per-distance. It is clear that the low-speed data has a higher power density than high frequency data. Further, the correlation between signals is poor at oscillations faster than 0.1 cycles per meter, but matches quite well for lower frequencies. This not only demonstrates the vehicle's speed-dependant filtering effects, but also shows that speed-independent correlation between multiple traversals of a path can be achieved if low-frequency pitch data is used.

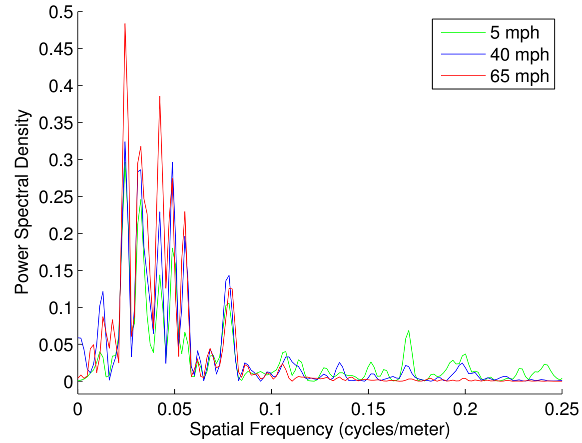

Fig. 2. Power Spectral Density (PSD) of the vehicle's pitch response at various speeds [8].

When explaining this research, we are often asked what road "features" are used for correlation in the algorithm. Possible pitch disturbance sources can be differentiated by their different spatial distance and hence frequency range: roadway surface texture has variations on the order of centimeters (100 cycles/meter); potholes on the order of 10 centimeters (10 cycles per meter); step changes surface elevation would cause pitch changes on the order of 1 meter (1 cycle per meter); imperfect surface leveling during roadway construction creates undulations between 10 and 100 meters in length (0.1 and 0.01 cycles per meter); and, due to sighting-distance requirements, roadway elevation changes are designed to change on the order of 100 meters or longer (lower than 0.01 cycles per meter). The PSD in Figure 2 shows that the most likely source of correlation is the low-frequency undulation caused by uneven road-surfacing during construction. Thus, if one were finding vehicle position by correlation of a pre-mapped profile, one would thus not expect any fundamental failure in localization from a simple pothole. Only a major roadway resurfacing event or rerouting would produce errors sufficiently large as to require reconstruction of the roadway map.

Further, the large-scale undulations also explain why the correlation is largely insensitive to lateral position, a fact implied by Figure 2 where no attempts were made to maintain exact lane position. Road construction and surface finishing tend to produce undulations in the longitudinal direction that are invariant to lateral position changes on the size scale of the vehicle width, e.g. the width of a steam roller or concrete slab pour. Additional experiments, discussed in [8], show that the longitudinal position estimation error indeed increases with larger errors in lateral lane-keeping error, and that the relationship is approximately linear for small lane deviations and modern roadways. While formal relationships were not established, a 0.5 meter RMS lane-keeping error resulted in approximately a 1 meter additional longitudinal localization error.

Hereafter, the pitch response of the map and in-vehicle pitch measurement are both filtered using a second-order low-pass spatial filter with a cutoff frequency of 0.1 cycles/meter. The roadway map features are recorded at 50 Hz using a vehicle traveling at 5 m/s to produce pitch readings at 0.1 meter intervals. This combination of extreme low speeds and high sampling rates at 100 times the period of correlation is intended to avoid any aliasing effects that could later bias the estimator.

To test whether terrain-based localization is feasible, experiments were conducted comparing a low-speed (5 m/s) pitch-map of the PTI test track to samples of pitch data taken at high speeds (20 m/s). Truth estimates of vehicle position during all runs were determined by a NovAtel "Span" Differential Global Positioning System (DGPS) that is factory-integrated with a Honeywell HG1700 ring-laser gyro Inertial Measurement Unit (IMU), with positioning accuracy of 2 cm. Off-line correlation was attempted using correlation windows of approximately 300 data points and a Pearson-product correlation metric. Comparison across 5 test trials with DGPS showed that the roadway correlation method was successful in localizing the vehicle along the mapped track in longitudinal position. The resulting localization accuracy ranged from 1 meter to 10 cm [8]!

Although this estimation accuracy is excellent, the off-line correlation algorithm is not fast enough for real-time, in-vehicle applications. Further, there are no obvious mechanisms to include vehicle dynamics or multiple, simultaneous sources of roadway terrain measurements, for example correlation from simultaneous in-vehicle pitch and roll measurements. This work solves these issues using advanced filtering methods, specifically the use of a particle filter. The roll correlation work is not shown because roll disturbance measurements are strongly coupled with vehicle steering inputs, and thus a lengthy discussion of vehicle roll dynamic models would be necessary to include such analysis. The focus hereafter on correlation just using pitch measurements is adequate to succinctly demonstrate the localization algorithm with little loss in generality.

III. Particle Filter Estimation

Particle filters are Monte-Carlo estimators that are known to be quite robust to non-Gaussian variance distributions similar to what would occur in this work due to similarities in the road profile along different segments. A Kalman filter couldn't be used to estimate the vehicle position because the vehicle can start anywhere along the map, hence the initial probability distribution is uniform whereas a Kalman filter requires a gaussian probability distribution. Due to the rapid advances in computing power, particle filters have recently been demonstrated to be fast enough for real-time applications [11] [12]. As a result, this method is gaining wide use for localization [13] [14] [15] [5], tracking [16], and even vehicle localization during GPS outages [17].

The particle filter algorithm used in this work was implemented off-line using data previously recorded using an instrumented vehicle equipped with the previously described

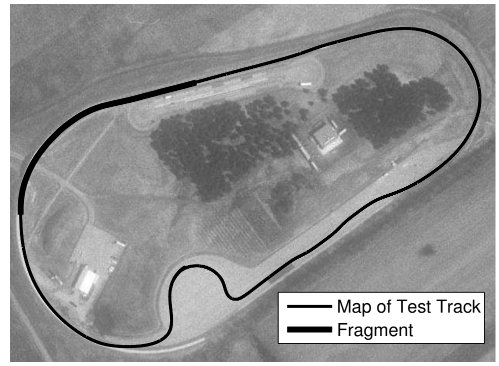

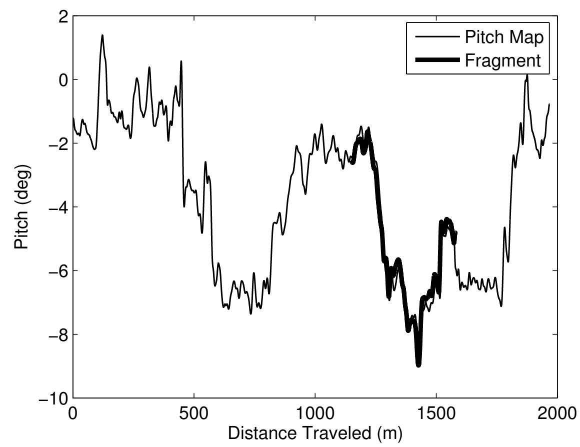

Fig. 3. Overhead view and pitch data of the terrain map and highlighted test fragment.

IMU and DGPS system at the test track at the Pennsylvania Transportation Institute. At various speeds, pitch and position data sets were recorded separately over small fragments of the test track and a complete map was recorded separately over the entire track at a constant speed. Figure 3 shows the map and fragment data demonstrating that there are visible variations in pitch between the data sets due to differences in speed and inexact path tracking. The pitch data were filtered using a low-pass filter at the cutoff frequency of 0.1 cycles/meter as discussed above.

The algorithm begins by converting the time-dependant data to the spatial domain, or more plainly as a function of distance from the starting point. Other than wheelbase filtering which is dependant on velocity, this removes velocity dependence on the pitch data. A set of N equally weighted and randomly distributed particles are located along the terrain map. The pitch estimate of each particle location is determined from the pitch map; particles that lie between the discretely mapped locations are determined via nearest-neighbor linear interpolation.

The particle filter algorithm is based off of Algorithm 3 in [16] and begins to iterate through the fragment data by repeating the following: First, the position estimates, denoted by X, at time interval k are updated from the previous estimate by

$$X^k = X^{k-1} + dX + dO, \tag{1}$$

where dX is the distance the vehicle travels between time steps as inferred from odometry and dO is gaussian white noise of variance R_O, equal to the variance of the odometry measurement.

Second, the weights of the position particles are updated by measuring the actual vehicle pitch and comparing it to the particle's pitch estimate using a standard particle weighting function. The importance density is assumed to be the prior density and the pdf is assumed to be gaussian:

$$q_i^k = \frac{exp\left(-\frac{1}{2\cdot R}\cdot(\phi_a - \phi_{p,i})^2\right)}{\sum_{i=1}^{N}\left(exp\left(-\frac{1}{2\cdot R}\cdot(\phi_a - \phi_{p,i})^2\right)\right)}, \tag{2}$$

Here R is the measurement noise variance on pitch, ϕ_a is the measured pitch, and $\phi_{p,i}$ is the i^{th} particle's pitch corresponding to its position along the terrain map.

Third, the particles are resampled following Algorithm 2 described in [16] where the number of effective position particles N_{eff} is calculated as

$$N_{eff} = \frac{1}{\sum_{i=1}^{N}\left(q_i^k\right)^2} \tag{3}$$

and when N_{eff} is below a threshold of N_T, the position particles are re-sampled by

$$\begin{array}{l} c = cumsum\left(q^k\right) \\ u_1 = rand(1)\cdot N^{-1} \\ i = 1 \\ \text{for}\ \ j = 1...N \\ \quad u_j = u_1 + (j-1)\cdot N^{-1} \\ \quad \text{while}\ \ u_j > c_i \\ \quad\quad i = i + 1 \\ \quad \text{end} \\ \quad X_j^k = X_i^k \\ \quad q_j^k = N^{-1} \\ \text{end} \end{array} \tag{4}$$

where $rand(1)$ is an evenly distributed random number in $[0, 1]$ and $cumsum$ is the cumulative sum.

Fourth, at every time step the vehicle's position is estimated as the mean of the position particles. The position error estimate is also calculated as the standard deviations of the position particles. This use of the entire population to characterize the estimate is fairly conservative since the position estimate of the "best" particle is in general far better than that of the population mean. However, for this study on the feasibility of the algorithm itself, convergence of the population to the correct solution is a far better indicator of algorithm performance than is analysis of the best particle estimate.

IV. RESULTS

The particle filter described above was implemented with N equal to 1000 particles, $N_T = 0.9\cdot N$, and the vehicle traveling about 15 m/s. The value of the pitch noise variance, R, was relaxed to a value of 0.1 degrees^2, much greater then the variance in the IMU pitch measurement of 0.000169 degrees^2 [8]. R enters Eq. 2 as a variance term, but it is actually an indicator of the amount of confidence placed in the accuracy of the measurement; a small R means the measurement is trusted to be highly accurate. In this case, R is larger than the actual variance to account for a mean offset likely due to different loading conditions in the vehicle between the data sets. Without this modification, the algorithm could converge too quickly to another portion of the map where the pitch value is similar to the biased pitch measurement, resulting in an erroneous solution.

The vehicle used in this study could not be equipped to measure odometry, however GPS measurements were available at all time samples. The odometry measurement was instead reconstructed from GPS by calculating the distance between data samples and adding a variance. The value of the odometry measurement variance, R_O, was calculated using the results of an effective tire radius study [18], where a tire with a specified radius of 321.65 mm was measured to have a nominal effective radius of 310.4 mm. Under different loading and tread conditions the effective radius was shown to vary by as much as 0.8%, so a dead-reckoning odometry measurement could vary by 0.8%. Thus, to be conservative, the variance in the odometry measurement was chosen to be $R_O = (0.01\cdot dX)^2$ m^2.

The convergence results on the track are shown in Figures 4 and 5. Figure 4 is an overhead view of the test track with the position estimates as dots, the mean estimate as a circle, and the actual "true" position measured from differential GPS shown as a box. It can be seen that as the vehicle travels, the position estimates converge to the measured vehicle location. Figure 5 demonstrates the convergence on the pitch response map where the dots represent the pitch at each position estimate along the map and the mean position estimate as the circle. Convergence is clearly seen within about 150 meters of forward roadway travel.

At every 10 meters the vehicle's position is predicted by the mean of the particle's position estimates and the error in the prediction is calculated. Because a driver is not capable of driving over the exact same position around the track at every pass, a path error is introduced called the lane-keeping error. An estimate of error between the predicted position to the actual vehicle position as measured by DGPS would include the lane-keeping error. In order to remove the lane-keeping error the measured vehicle position is projected to the nearest position on the map. The corrected error is calculated as the distance from the predicted position to the corrected measured position as

$$E^k = |\bar{X} - x_c^k|, \tag{5}$$

where $\bar{X}$ is the mean location of the position estimates and x_c is the vehicle's DGPS measured location projected to the map. The lane-keeping error is calculated as

$$E_l^k = |x_c^k - x_m^k|, \tag{6}$$

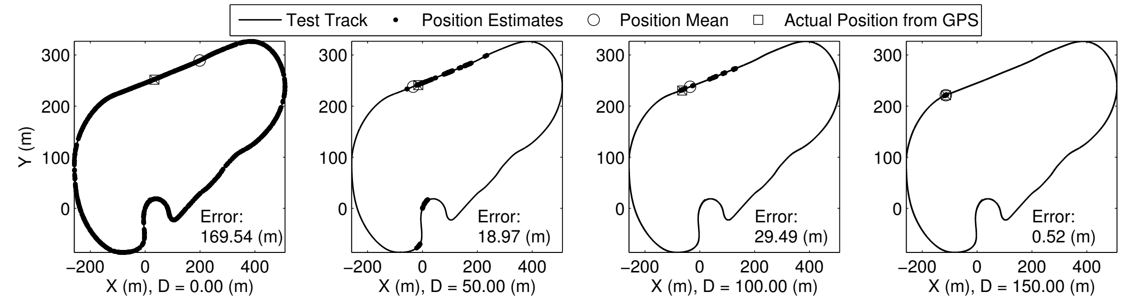

Fig. 4. Position estimates after the vehicle traveled a distance D.

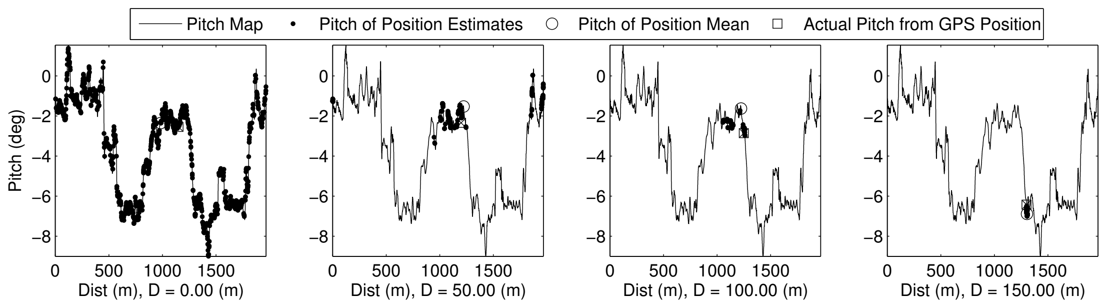

Fig. 5. Pitch estimates after the vehicle traveled a distance D.

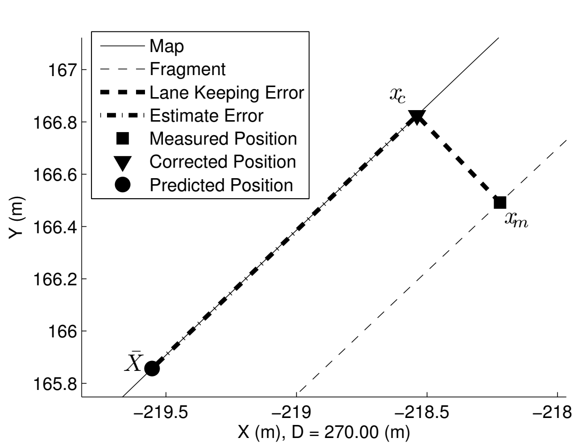

Fig. 6. Lane keeping error and its correction.

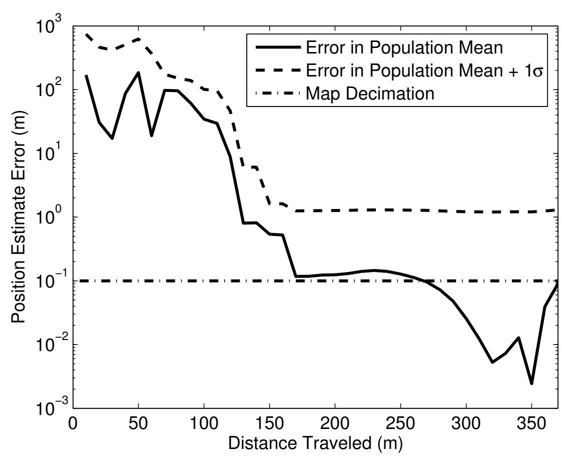

Fig. 7. Position estimate error as a function of the distance traveled.

where x_m is the vehicle's actual DGPS measured location. An example of the projection and error measurement is shown in Figure 6.

The standard deviation of the population of particle position estimates is also calculated at every 10 meters. These results are shown in Figure 7 and demonstrate the algorithm converged to an accuracy of about 1 meter after moving over 150 meters. Also shown in this figure and the following is a line representing the mapping interval of the terrain map, 10 cm, a value which places a lower-limit on the achievable estimator accuracy. Compared to the off-line correlation method mentioned earlier[8], the on-line particle filter algorithm estimate is similar in accuracy.

One disadvantage of particle filters is the computational burden imposed by using a large number of particles to achieve an accurate yet robust prediction. To determine the relationship between particle population size and estimator accuracy, the same algorithm was tested using various numbers of particles N distributed randomly across the one mile map. The converged estimate error was inferred by examining the standard deviation of the population averaged over the final 100 meters. This was repeated 10 times, each time with a different random initial population, and averaging the results. The plot of population size versus estimator

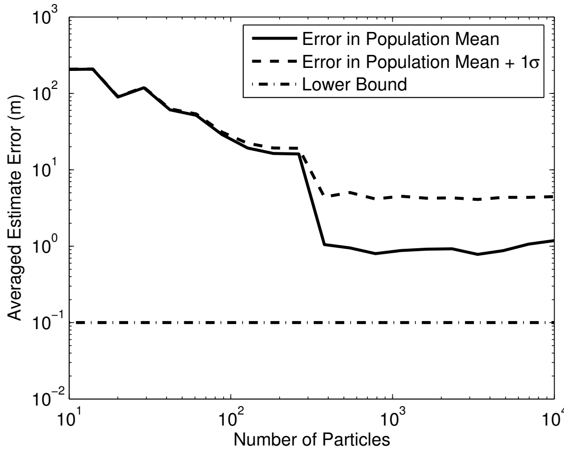

Fig. 8. Position estimate error as a function of the number of position estimates.

error is graphed in Figure 8. It can be seen that, as the numbers of particles increases to 1000 per mile, the accuracy improves greatly. It can also be seen that the algorithm resulted in about the same accuracy with $N = 1,000$ as when $1,000 < N \leq 10,000$. This suggests that the position of the vehicle can be estimated most efficiently at about 1,000 particles per mile and further particles are unnecessary. This can also be used to estimate the computational cost of using a particle filter to estimate vehicle position over large area maps for a given measure of initial uncertainty. These preliminary estimates suggest that, for each mile of locational uncertainty within a mapped roadway, 1000 particles would be needed.

V. CONCLUSIONS

This work shows that a vehicle's longitudinal position can be estimated given a terrain map and pitch measurements. To achieve on-line estimation, a particle filter algorithm is presented. The convergence of the estimate is seen to occur within approximately 150 meters of driving, with converged longitudinal positioning error of consistently achieving accuracies of about 1 meter or better as compared to a DGPS system. By examining population size, a particle density of 1,000 particles per mile of unknown roadway has been shown to be sufficient for position estimation and can be used as a basis for larger maps covering greater distances.

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] H.-C. Chu and R.-H. Jan, "A GPS-less self-positioning method for sensor networks," *2005. Proceedings. 11th International Conference on Parallel and Distributed Systems*, vol. 2, pp. 629–633, Jul. 20–22, 2005.

[5] 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.

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

[7] 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.

[8] R. D. Martini, "GPS/INS sensing coordination for vehicle state identification and road grade positioning," MS, Pennsylvania State University, August 2006.

[9] H. Bae, J. Ryu, and J. Gerdes, "Road grade and vehicle parameter estimation for longitudinal control using GPS," in *Proceedings of the IEEE Conference on Intelligent Transportation Systems*, 2001.

[10] R. S. Sharp, "Wheelbase filtering and automobile suspension tuning for minimizing motions in pitch," *2002. IMechE 2002. Proceedings 2002 IMechE, Part D: J. Automobile Engineering*, vol. 216, no. 12, pp. 933–944, 2002.

[11] C. Kwok, D. Fox, and M. Meila, "Adaptive real-time particle filters for robot localization," *2003. Proceedings. ICRA '03. IEEE International Conference on Robotics and Automation*, vol. 2, pp. 2836–2841, Sep. 14–19, 2003.

[12] ——, "Real-time particle filters," *Proceedings of the IEEE*, vol. 92, no. 3, pp. 469–484, Mar. 2004.

[13] F. Caron, M. Davy, E. Duflos, and P. Vanheeghe, "Particle filtering for multisensor data fusion with switching observation models: Application to land vehicle positioning," *Signal Processing, IEEE Transactions on [see also Acoustics, Speech, and Signal Processing, IEEE Transactions on]*, vol. 55, pp. 2703–2719, Jun. 2007.

[14] F. Chausse, J. Laneurit, and R. Chapuis, "Vehicle localization on a digital map using particles filtering," in *Intelligent Vehicles Symposium, 2005. Proceedings. IEEE*, Jun. 6–8, 2005, pp. 243–248.

[15] D. Fox, W. Burgard, H. Kruppa, and S. Thrun, "A probabilistic approach to collaborative multi-robot localization," *Autonomous Robots*, vol. 8, no. 3, pp. 325–344, 2000. [Online]. Available: citeseer.ist.psu.edu/fox00probabilistic.html

[16] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking," *Signal Processing, IEEE Transactions on [see also Acoustics, Speech, and Signal Processing, IEEE Transactions on]*, vol. 50, no. 2, pp. 174–188, Feb. 2002.

[17] C. Boucher, A. Lahrech, and J. C. Noyer, "Non-linear filtering for land vehicle navigation with GPS outage," *Systems, Man and Cybernetics, 2004 IEEE International Conference on*, vol. 2, pp. 1321–1325, Oct. 2004.

[18] C. R. Carlson and J. C. Gerdes, "Consistent nonlinear estimation of longitudinal tire stiffness and effective radius," *IEEE Transactions on Control Systems Technology*, vol. 13, no. 6, pp. 1010–1020, Nov. 2005.