Monte Carlo Localization in Dense Multipath Environments Using UWB Ranging†

Damien B. Jourdan, John J. Deyst, Jr., Moe Z. Win, Nicholas Roy
Massachusetts Institute of Technology
Laboratory for Information and Decision Systems
Computer Science and Artificial Intelligence Laboratory
32 Vassar St., Cambridge, MA 02139
{*jourdan*|*deyst*|*moewin*|*nickroy*}@*mit.edu*

Abstract— **For most outdoor applications, systems such as GPS provide users with accurate position estimates. However, reliable range-based localization using radio signals in indoor or urban environments can be a problem due to multipath fading and Line-of-Sight (LOS) blockage. The measurement bias introduced by these delays causes significant localization error, even when using additional sensors such as an Inertial Measurement Unit (IMU) to perform outlier rejection.**

We describe an algorithm for accurate indoor localization of a sensor in a network of known beacons. The sensor measures the range to the beacons using an Ultra-Wideband (UWB) signal and uses statistical inference to infer and correct for the bias due to LOS blockage in the range measurements. We show that a particle filter can be used to estimate the joint distribution over both pose and beacon biases. We use the particle filter estimation technique specifically to capture the non-linearity of transitions in the beacon bias as the sensor moves. Results using real-world and simulated data are presented.

I. Introduction

Since the Global Positioning System (GPS) became widely accessible [1], localization in the absolute frame (or geolocation) has found application in many different fields. In areas with good Line-of-Sight (LOS) to GPS satellites, this technique provides a good estimate (within a few meters) of the user's location on the earth. However, indoor geolocation has always been a more difficult problem for several reasons.

First, the GPS signal is not strong enough to penetrate through most materials. As soon as an object occludes the GPS satellite, the signal is corrupted. This constrains the usefulness of GPS to open environments, and limits its performance in forests or in dense urban environments as retaining a lock on the GPS signals becomes more difficult. GPS typically becomes completely useless inside buildings. However the need for accurate geolocation is not constrained to open environments, both in civil and military applications. In the commercial world for example, the tracking of inventory in warehouses or cargo ships is an emerging need. In military applications the problem of "blue force tracking", i.e. knowing where friendly forces are, is of vital importance, especially in urban scenarios.

Our scenario is that of an agent (such as a person or a vehicle) entering a building and *accurately* tracking its absolute position over time. The position estimate should have a precision of one meter or less (i.e. on the order of some of the building feature dimensions, such as hallway width), a precision current indoor localization systems lack. The solution presented in this paper relies on establishing a local GPS-like network, where fixed beacons emit an Ultra-Wideband (UWB) signal for ranging purposes (in that sense they act just like GPS satellite). We assume that the location of these beacons is known, for example because they are placed outside and can rely on GPS. We also assume that the agent carries an Inertial Measurement Unit (IMU) that provides attitude rates and instant accelerations. The agent then determines the time-of-arrival of the signals from which it infers ranges to these fixed beacons. These ranges, coupled with the IMU information, are then used to update the agent position.

If the ranges to the beacons were accurate, then three beacons would be sufficient to determine the agent position with accuracy. However this is typically not the case in indoor environments because of multipath fading and LOS blockage.

First, if a signal with a small bandwidth is used, it will suffer from multipath fading. In the case of narrowband signals, the distance between two points is inferred from the phase difference between received and transmitted signals. In areas with dense multipath, the received signal can be the sum of a number of carriers arriving along different paths. The phase measured will therefore be different than that of the sole LOS carrier, and the range measurement will not reflect reality [2].

Second, in some areas of the building there may be no LOS between transmitter and receiver. The received signal can then be of two kinds. It may be from a direct path, where the signal traveled along a straight line but had to go through materials other than air: because the propagation of electromagnetic signals is slower in some materials than in the air, the signal arrives with a greater delay. Alternatively, the received signal may have come from reflections only. In both cases the consequence of non-LOS (NLOS) propagation is the same: the range estimate is larger than the true one.

The first difficulty–the fading of the signal due to multipath–is satisfactorily resolved by choosing an appropriate radio signal. Since the ranging accuracy increases with the bandwidth [2], using a signal with a large bandwidth should resolve this problem. UWB signals have been shown to be immune to

† This research was partially supported by Draper Laboratory. Their support is gratefully acknowledged.

multipath fading [3]. For this reason UWB was selected as an appropriate signal to perform indoor ranging in our proposed concept.

The NLOS propagation difficulties, however, remain: in fact, they cannot be resolved at the hardware level. As mentioned above, NLOS propagations add a positive bias to the true range between agent and beacon, so that the measured range is larger than the true value. This error has been identified as a fundamental limiting factor in UWB ranging performance [4]. If accurate UWB channel modeling was available, then it may be possible to predict these biases throughout the environment. However in many applications such modeling will not be accurate, and often not available at all (e.g. military scenarios). In GPS receivers there is a redundancy in measurements so that corrupted measurements can be discarded [5]. However in an indoor environment we will not usually have this luxury since most if not all range measurements are likely to be biased (see Section III): the biases will have to be estimated. An indoor localization method oblivious to them is unlikely to perform well.

We show in this paper that these biases can be estimated jointly with the position estimate of the sensor, allowing for the range measurements to be corrected and the sensor accurately localized. We use a particle filter estimation technique specifically to capture the non-linearity of bias transitions. In particular, conventional Extended Kalman Filter (EKF) approaches [6] to the problem fail. We demonstrate experimentally that the biases can dominate the signal to the extent that outlier rejection causes the EKF to lose track of its position.

In this paper we will focus on indoor localization in an environment with a set of beacons with *a priori* known positions; we will not deal with the problem of Simultaneous Localization And Mapping (SLAM) of the beacon position: the goal is to have a system that provides the same capabilities as GPS, but for an indoor environment. We do, however, assume that no map of the physical layout building is available to infer the signal bias.

The paper is organized as follows. In Section II we describe Monte Carlo localization. In Section III we present results from UWB ranging tests and use them to model the biases in the localization algorithm. Section IV contains the results from experiments utilizing real-world data from Section III. We conclude this paper in Section VI.

II. Localization

Our goal is to estimate the position of an agent moving indoor, given a set of UWB beacon ranges and IMU measurements. In this paper we restrict ourselves without loss of generality to two dimensions. We also assume for simplicity that the translational velocity of the agent is known at all times. Our state vector $\mathbf{x}(t)$ at time t for the agent contains the following variables:

- $x(t)$, $y(t)$ are the coordinates of the agent
- $\theta(t)$ and $\dot{\theta}(t)$ are its heading angle and heading angle rate

Since we are in a 2D environment with known velocity, the IMU considered in this paper needs only be a rate gyro. The measurement received at time t are then:

- $z_\theta(t)$, the heading angle rate of the agent from the IMU
- $\mathbf{r}(t) = \{r_1(t) \ldots, r_n(t)\}$, the ranges from the n beacons.

Given the IMU and range measurements of the beacons, we use the Bayes' filter equation [7] to maintain a probability distribution over the current state $\mathbf{x}(t)$:

$$\begin{aligned} p(\mathbf{x}(t)|z_\theta(t), \mathbf{r}(t), z_\theta(t{-}1), \mathbf{r}(t{-}1), \ldots) \quad = \\ \alpha \cdot p(z_\theta(t)|\mathbf{x}(t)) \cdot p(\mathbf{r}(t)|\mathbf{x}(t)) \times \qquad (1) \\ \int p(\mathbf{x}(t)|\mathbf{x}(t{-}1)) \cdot p(\mathbf{x}(t{-}1)|z_\theta(t{-}1), \mathbf{r}(t{-}1), \ldots) d\mathbf{x}(t-1) \end{aligned}$$

where α is a normalization term. We have assumed that the current state $\mathbf{x}(t)$ only depends on the previous state $\mathbf{x}(t-1)$ (by the Markov assumption), and that the measurements $z_\theta(t)$ and $\mathbf{r}(t)$ are independent given $\mathbf{x}(t)$.

A conventional EKF is often used to maintain the distribution $p(\mathbf{x}(t))$ by linearizing the prediction and measurement functions and modeling the noise terms as Gaussian [6]. These two constraints allow the distribution over $\mathbf{x}(t)$ to be approximated as a Gaussian. The advantage to such a representation is that the distribution can be represented using only a small number of parameters (a mean vector and covariance), and updated very efficiently. If, however, the prediction and measurement models are not easily linearized and the noise terms are not Gaussian, then the EKF typically does an increasingly poor job of approximating the true distribution over $\mathbf{x}$, often leading to filter divergence.

An alternate technique for representing $p(\mathbf{x}(t))$ is to maintain a set of sample states drawn from the distribution [8]. Good techniques exist to sample from distributions even when the distribution itself cannot be represented, and statistics such as the mean, variance and higher order moments of the distribution can be computed directly from the samples instead of from the distribution parameters.

Monte Carlo Localization [9] is a form of robot localization using Importance Sampling [10], in which samples from a target distribution $p(\mathbf{x})$ are desired but cannot be drawn directly. Instead, samples are drawn from some known proposal distribution $q(\mathbf{x}(t))$ that does permit direct sampling. Each sample is assigned an importance weight $p(\mathbf{x}(t))/q(\mathbf{x}(t))$, and the set of weighted samples can be used in place of the distribution $p(\mathbf{x}(t))$. In sampling problems where the target distribution changes over time, the sample weights can be updated directly to reflect the new distribution, although finite numerical precision can cause the sample weights to converge eventually to 0. To avoid this problem, in Importance Sampling Resampling [11], the weighted samples are periodically resampled according to their weights to generate a new set of uniformly weighted samples.

In the localization problem, our target distribution is $p(\mathbf{x}(t)|z(t), r(t), z(t-1), r(t-1), \ldots)$. Under the assumption that we do not have a parametric representation of this distribution, we maintain a set of particles where the i^{th} particle $\mathbf{x}^{[i]}(t)$ is written:

$$\mathbf{x}^{[i]}(t) = [x^{[i]}(t), y^{[i]}(t), \theta^{[i]}(t), \dot{\theta}^{[i]}(t)] \qquad (2)$$

Since we can sample from the prediction model (by sampling simulated motions), our proposal distribution can be given by:

$$q(\mathbf{x}(t)) = \int p(\mathbf{x}(t)|\mathbf{x}(t{-}1))p(\mathbf{x}(t{-}1)|z_\theta(t{-}1),\mathbf{r}(t{-}1),\ldots)d\mathbf{x}(t-1). \tag{3}$$

We can then use the measurement likelihood as the importance weight, since

$$\begin{aligned} & \frac{p(\mathbf{x}(t)|z_\theta(t),\mathbf{r}(t),\ldots)}{\int p(\mathbf{x}(t)|\mathbf{x}(t{-}1))p(\mathbf{x}(t{-}1)|z_\theta(t{-}1),\mathbf{r}(t{-}1),\ldots)d\mathbf{x}(t-1)} \\ = \quad & \alpha \cdot p(z_\theta(t)|\mathbf{x}(t)) \cdot p(\mathbf{r}(t)|\mathbf{x}(t)) \end{aligned} \tag{4}$$

If the measured ranges are unbiased, the measurement models $p(\mathbf{r}(t)|\mathbf{x}(t))$ and $p(z_\theta(t)|\mathbf{x}(t))$ will be sufficient to unambiguously determine the agent position given measurement of three beacons.

III. Beacon Bias Modeling in the Particle Filter

As we have argued however, the ranges are positively biased. In order to accurately estimate the position of the agent, it is necessary to estimate both the robot position and the set of beacon biases $\mathbf{b}(t) = \{b_1(t), \ldots, b_n(t)\}$, where n is the number of beacons. The gyro, in addition to giving noisy measurements, also has a bias that evolves over time. We therefore included the gyro bias $g(t)$ as a state variable to be estimated.

The i^{th} particle $\mathbf{x}^{[i]}(t)$ is then written:

$$\mathbf{x}^{[i]}(t) = \left[x^{[i]}(t), y^{[i]}(t), \theta^{[i]}(t), \dot{\theta}^{[i]}(t), g^{[i]}(t), b_1^{[i]}(t), \ldots\right] \tag{5}$$

where $b_j^{[i]}(t)$ is the bias estimate of the *j*th beacon for particle *i*. In order to estimate the biases $\mathbf{b}(t)$ with the particle filter, we need a proposal distribution (cf. equation 3) and a likelihood model (cf. equation 4). Our likelihood model does not change as a result of estimating biases, but we need to modify the proposal distribution to model how these biases change over time. We can learn this model by looking at actual UWB measurements.

A. UWB Range Measurements

Using data collected by Win and Scholtz [12] we are able to build a probabilistic model for the beacon bias transitions. These measurements were collected on a floor of an office building. A bandwidth in excess of 1GHz was used, and the UWB transmitter was placed in a specific room, while measurements were taken in different rooms. For an in-depth description of the experiment, we refer the reader to [12]. For the purpose of this paper, we focus on two sets of data, one with measurements taken in one room, and the other with measurements taken at regular intervals along a corridor.

B. Measurements From a Grid

For the first set of data, 49 measurements were taken in a 7x7 square grid with a 6 inch spacing between measurement points. These were collected in a different room from the transmitter (so that the signal had to propagate through several walls). From the received signal we determined the time-of-arrival of the signal by manual inspection (for algorithms, see e.g. [13]), and inferred a range measurement. These range

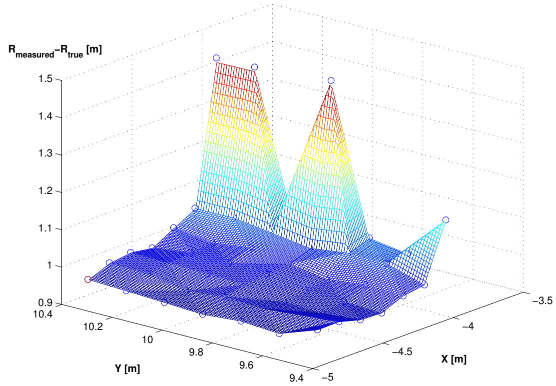

Fig. 1. Difference in metres (on the z axis) between the measured and true ranges for 49 points using a UWB signal.

estimates were then compared to the true distances between transmitter and receiver (this true distance was obtained from the building floor plan, with a precision of about 0.1m). Figure 1 shows the differences between the range estimates using the measured signal and the true range for the 49 points. It can be seen that for most points there is a constant difference of about 1m. This can be attributed to either LOS blockage or propagation delays (as mentioned before, these two phenomena produce the same effect). We also observe that for 4 points located toward one edge of the grid, this difference is higher. This shows that toward this edge either the propagation delay increases, or a multipath signal becomes the first arrival.

We draw two conclusions from this set of data. First, the range estimates are positively biased, as expected. Second, the value of the bias remains constant locally, and then suddenly changes in a discrete amount as we move through the building. For example, in Figure 1, notice the local plateau at 1m for $x < -4m$ and the sudden increase to 1.5m around $x = -4m$. We use these observations to build a probabilistic model of bias transitions in IV-A.

C. Measurements Along the Building's Corridor

In this second experiment, 33 measurements were made at regular intervals as the receiver was moved around the floor along the building corridor. Again, for each point a range

Fig. 2. Difference in metres between the measured and true ranges for the 33 corridor locations.

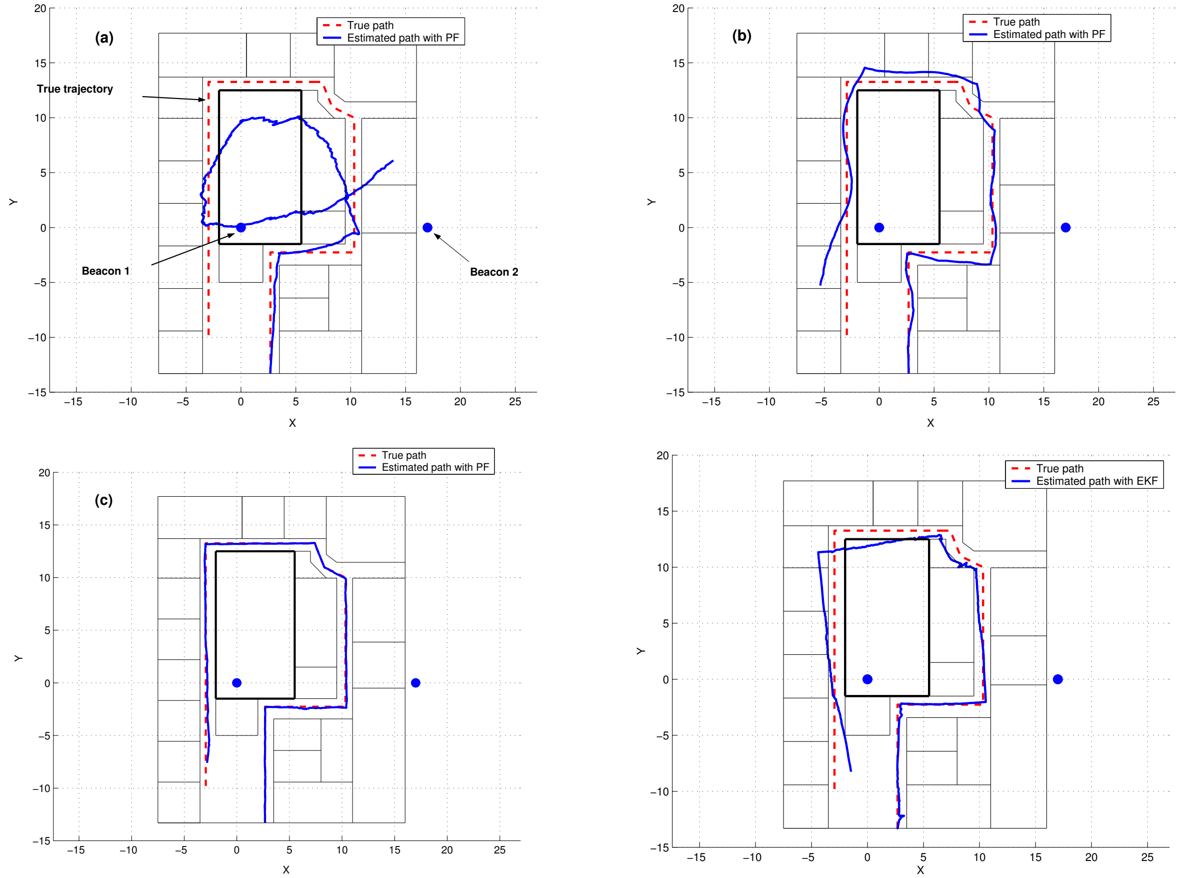

Fig. 3. Path estimate for (a) Case 1 (IMU only), (b) Case 2 (IMU and beacons, but no bias estimation), (c) Case 3 (full pose and bias estimation) (d) EKF with outlier rejection.

estimate was extracted from the received signal and compared to the true range. This difference is plotted on Figure 2 for all 33 points. It is seen that the positive bias changes as the receiver moves along the corridor. In this case, the value of the bias varied between 0.15m and 1m, and we will use these values in our experiments in section IV. The distance between two measurement points was about 2m, and we make an approximation that the bias value remains constant between two points. This assumption is consistent with our observation above (III-B).

D. The Beacon Model

We can now construct a probabilistic model of beacon biases based on this experimental data. We first define the rate of beacon bias change r_{change}, which is the expected number of times the bias changes per second and per beacon. This rate typically depends on the environment, but we found that the particle filter performed well even for inaccurate values of r_{change}. In our experimental results, we set $r_{change} = 1$ Hz. At each time step dt, the probability that the bias of the j^{th} beacon changes for a given particle is $r_{change}dt$. When this is the case, the particle's j^{th} bias value is assigned a uniformly distributed random number between 0 and $b_j^{[i]} - \epsilon$ or between $b_j^{[i]} + \epsilon$ and β, where β is the maximum value the bias can take, and ϵ is a positive number smaller than β. The role of ϵ is to ensure that the bias change is sufficiently large, in order to model the fact that when the bias changes, it is likely to change to a value significantly different from the previous one (e.g. in Figure 1, it jumps from 1m to 1.5m). Our bias motion model is therefore a uniform distribution notched about the current bias, where the notch has a width of 2ϵ. In practice, we set ϵ to be equal to 3 standard deviations of the beacon measurement noise.

We may have been tempted to simply model the bias transition as a Gaussian centered at the current bias value. But this does not agree with the measurements presented earlier, and in fact simulations showed that the resulting estimation performs poorly.

The probabilistic model of the agent's dynamics is as follows:

$$
\begin{aligned}
g^{[i]}(t) &= g^{[i]}(t{-}1) + N(0, \sigma_g^2) && (6)\\
\dot{\theta}^{[i]}(t) &= \dot{\theta}^{[i]}(t{-}1) + N(0, \sigma_\theta^2) && (7)\\
\theta^{[i]}(t) &= \theta^{[i]}(t{-}1) + \dot{\theta}^{[i]}(t{-}1)dt && (8)\\
x^{[i]}(t) &= x^{[i]}(t{-}1) + V\cos(\theta^{[i]}(t))dt && (9)\\
y^{[i]}(t) &= y^{[i]}(t{-}1) + V\sin(\theta^{[i]}(t))dt && (10)
\end{aligned}
$$

where V is the (known) translational velocity, and σ_g and σ_θ are the standard deviations of the noise for the gyro and the gyro bias, respectively. $N(0, \sigma^2)$ is a normally distributed random number with mean 0 and variance σ^2. In our experimental

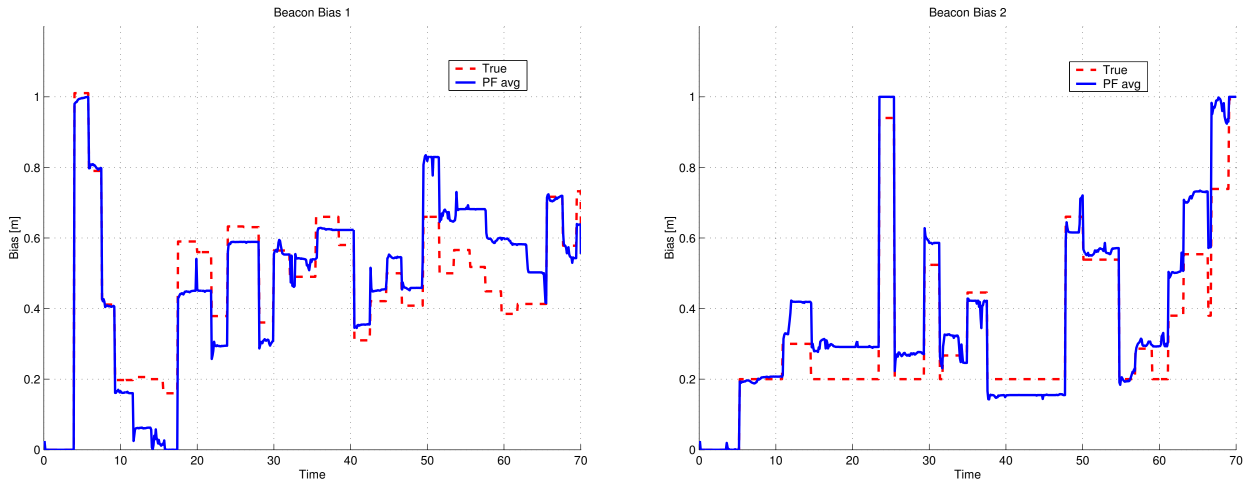

Fig. 4. Estimates of the biases of beacon 1 (left) and beacon 2 (right) are shown by the solid line. Note that beacon 1 is an actual beacon, and the bias is being induced by some real physical process. Beacon 2 is simulated, and its bias is artificial induced.

results, $\sigma_g = 0.001$ and $\sigma_\theta = 3$.

Note that our measurement model assumes not only that the gyro measurements are independent of the beacon measurements, but also that the beacon measurements are independent of each other. Therefore, for a given particle $\mathbf{x}^{[i]}(t)$, the likelihood is a product of $n+1$ factors. The first factor is the likelihood of the current measurement $z_\theta(t)$, which we model as a normal distribution

$$p(z_\theta(t)|\mathbf{x}(t)) = N(\dot{\theta}, \sigma_{\dot{\theta}}^2), \tag{11}$$

where $\sigma_{\dot{\theta}} = 0.1$.

The n remaining factors are the likelihoods of the current n range measurements given each particle's position $\mathbf{x}(t)$ and its n beacon biases $\{b_1^{[i]}(t), b_2^{[i]}(t), \ldots, b_n^{[i]}(t)\}$,

$$p(r_i(t)|\mathbf{x}(t), b_i(t)) = N(||\mathbf{x}(t) - \mathbf{y}_i(t)|| - b_i(t), \sigma_{rf}^2), \tag{12}$$

where $\mathbf{y}_i(t)$ is the location of the i^{th} beacon, and the σ_{rf} is the error in the range sensor taken from the sensor specification. This was specified as $\sigma_{rf} = .025m$ for our experiments. This value in fact was a slight underestimate; using the range measurements from the equal-bias points on the plateau ($x < -4m$) in the room experiment (Figure 1) showed the variance to be $\sigma_{rf} = 0.03m$.

IV. Experimental Results

In this section we describe our simulation and then compare the results for different cases. We only had access to actual UWB signal data for a single transmitter, so the second signal is simulated. Although in theory three beacons are necessary to unambiguously localize an agent in 2D, two beacons are sufficient in our case because we use an IMU and the agent's initial position is known, so that the ambiguity is removed.

A. Experimental Setup

An agent travels at a constant speed of 1m/s along a corridor in an office building. It carries a rate gyro providing its instantaneous heading angle rate, and receives range measurements from two beacons at a rate of 10Hz. The range measurements from those two beacons are positively biased by $\mathbf{b}(t)$. For the bias of one beacon, we use the actual range measurements from a physical transmitter (as in section III). For the second beacon, we simulated range measurements and a bias profile. The map of the environment is shown in Figure 3. We assume that the agent knows its initial state since prior to entering the building it can use GPS to determine its exact location. We show the results of the estimation using the particle filter for the following cases:

- Case 1: the agent uses only its IMU to estimate $\mathbf{x}(t)$
- Case 2: the agent uses its IMU and beacon range to estimate $\mathbf{x}(t)$, but does not estimate $\mathbf{b}(t)$
- Case 3: the agent uses its IMU and beacon range, estimating the joint distribution over $\mathbf{x}(t)$ and $\mathbf{b}(t)$.

We also compare our results against an EKF with outlier rejection.

Case 1 Results: The estimated trajectory of the agent (Figure 3a) does not track the true trajectory, as the gyro noise quickly dominates.

Case 2 Results: The results shown in Figure 3(b) are better than in Case 1, but the path estimate oscillates about the true trajectory, yielding a position error of more than 1m on average. This is due to the fact that the ranges have unmodeled bias. The position estimate oscillates as the particle filter tries to best adjust to these biased measurements.

Case 3 Results: In this case the beacon biases are estimated and the results of Figure 3(c) show a very close tracking of the true path. The reason for such good performance comes from the fact that the beacon biases are being estimated: the particle likelihoods incorporate the biases, so the measurements can still be used with confidence. Figure 4 shows the bias estimates for beacon 1 and beacon 2.

We also show an EKF with outlier rejection for comparison. Since only two beacons are present, there is not enough redundancy to ensure sufficient measurement updates, so the results are poor, as shown in Figure 3(d).

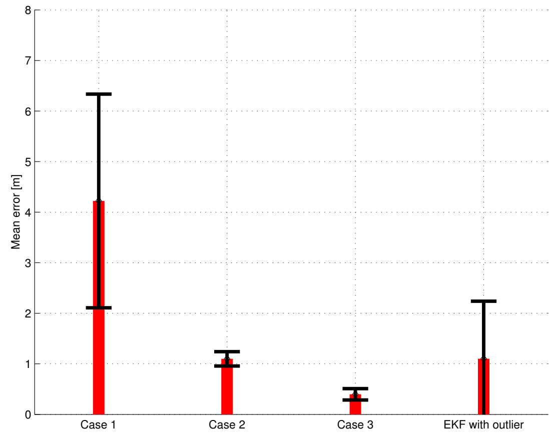

Fig. 5. Mean error between the true position and its estimate after 100 simulation runs for Cases 1, 2, 3 and the EKF. The standard deviation is shown.

B. *Systematic Comparison*

Cases 1, 2 and 3 and the EKF were performed over 100 simulation runs, shown in Figure 5. The mean error of Case 1 is not surprisingly much worse than the other two mean errors, and Case 3 provides localization with twice the accuracy of Case 2. Estimating the beacon biases doubles the average accuracy in this example. It should be noted that in this example the beacon biases were limited to a maximum of 1m, so that the errors are fairly limited. If the biases are larger than 1m (in Figure 1 some equal 1.5m), then we can expect Case 2 to perform ever worse than Case 3.

V. Related work

Ours is not the first approach to using Monte Carlo techniques for inference in sensor networks, however, we believe that ours is the first range-based localization to demonstrate robustness to hidden biases. In contrast, Ladd *et al.* [14] build an explicit model of the spatial distribution of biases which they then use to build an HMM and solve the global localization problem. They are able to use 802.11 signals to localize a laptop based only on beacon measurements, however, they do require the substantial initial training phase.

Smith *et al.* [15] avoid many of the bias issues by explicitly using two different range sensors in the Cricket system. Their results indicate that the EKF approach with outlier rejection provides accurate localization in the face of bias given a sufficient number of beacons (several per room). However, they do encounter periodic EKF failures; their assumption is that having the moving agent transmit additional signals to the beacons is sufficient to recover from the EKF failure.

Biswas *et al.* [16] take a similar approach to ours in factoring the likelihood model, however, they are attempting to solve a fundamentally different problem in assessing the presence of enemy agents in the sensor network.

VI. Conclusions

A method based on the coupling of an IMU and UWB ranging beacons has been proposed to improve indoor geolocation. It has been shown to overcome the two main difficulties for indoor localization. Using UWB signals resolves problems due to multipath fading. The remaining issue (NLOS propagation) adds a positive bias to the range measurement, which degrades the localization accuracy and robustness if nothing is done. However we showed that a particle filter can be used to simultaneously estimate the state of the agent and the beacon measurement biases. Experimental results incorporating real and simulated UWB measurements demonstrated the efficacy of the particle filter approach, which enabled us to localize the agent within a few tenths of a meter. Although these results are based on a limited set of data, we believe that this example shows the validity of our concept to provide a realistic solution to the challenge of accurate indoor geolocation. We hope to further validate our system in the near future as we perform more experiments.

References

[1] J. J. Spilker, Jr., "GPS signal structure and performance characteristics," *Journal of the Institute of Navigation*, vol. 25, no. 2, pp. 121–146, Summer 1978.

[2] K. Pahlavan, X. Li, and J.-P. Makela, "Indoor geolocation science and technology," *IEEE Communications Magazine*, pp. 112–118, February 2002.

[3] M. Z. Win and R. A. Scholtz, "On the performance of ultra-wide bandwidth signals in dense multipath environment," *IEEE Communications Letters*, vol. 2, no. 2, pp. 51–53, February 1998.

[4] J.-Y. Lee and R. A. Scholtz, "Ranging in a dense multipath environment using an uwb radio link," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 9, pp. 1677–1683, December 2002.

[5] E. D. Kaplan, Ed., *Understanding GPS: Principles and Applications*. Boston, MA: Artech House, 1996, ch. Performance of Standalone GPS.

[6] B. D. Anderson and J. B. Moore, *Optimal Filtering*. New Jersey: Prentice-Hall, 1979.

[7] D. Fox, W. Burgard, and S. Thrun, "Markov localization for mobile robots in dynamic environments," *Journal of Artificial Intelligence Research*, vol. 11, pp. 391–427, 1999.

[8] A. Doucet, N. de Freitas, and N. Gordon, Eds., *An introduction to sequential Monte-Carlo methods*. New York, NY: Springer-Verlag, 2001, ch. Sequential Monte Carlo Methods in Practice.

[9] D. Fox, W. Burgard, F. Dellaert, and S. Thrun, "Monte Carlo localization: efficient position estimation for mobile robots," in *Proceedings of the National Conference on Artificial Intelligence*, Orlando, Florida, 1999.

[10] J. Geweke, "Bayesian inference in economic models using Monte Carlo integration," *Econometrica*, vol. 45, no. 7, pp. 788–799, 1989.

[11] N. J. Gordon, D. J. Salmon, and A. F. M. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEE Proceedings-F*, vol. 140, no. 2, pp. 107–113, 1993.

[12] M. Z. Win and R. A. Scholtz, "Characterization of ultra-wide bandwidth wireless indoor channels: a communication-theoretic view," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 9, pp. 1613–1627, December 2002.

[13] Z. N. Low, J. H. Cheong, C. L. Law, W. T. Ng, and Y. J. Lee, "Pulse detection algorithm for Line-of-Sight (LOS) UWB ranging applications," *IEEE Antennas and Wireless Propagation Letters*, vol. 4, 2005.

[14] A. M. Ladd, K. E. Bekris, A. Rudys, L. E. Kavraki, and D. S. Wallach, "On the feasibility of using wireless ethernet for indoor localization," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 3, June 2004.

[15] A. Smith, H. Balakrishnan, M. Goraczko, and N. Priyantha, "Tracking moving devices with the Cricket location system," in *Proceedings of the 2nd International Conference on Mobile Systems, Applications, and Services*, Boston, MA, June 2004.

[16] R. Biswas, L. Guibas, and S. Thrun, "A probabilistic approach to inference with limited information in sensor networks," in *Proceedings of the 3rd International Symposium on Information Processing in Sensor Networks*, 2004.