Identifying and Filtering Noise Caused by Sensor Set Switching in UWB Indoor Position Tracking

Abstract: Trilateration calculations are affected by errors in distance measurements from the set of fixed points to the object of interest. When these errors are systemic, each distinct set of fixed points can be said to exhibit a unique set noise. For ultra-wideband (UWB) indoor position tracking, the set of fixed points is a set of sensors measuring the distance to a tracked tag. In this paper we develop a noise model for this sensor set noise, along with a particle filter that uses our set noise model. We test our methods on a real UWB system. Our methods showed an approximately 15% improvement in accuracy over the raw measurements.

Keywords: ultra-wideband (UWB), local positioning system (LPS), particle filter.

1 Introduction

Trilateration-based tracking relies upon measuring the distances from a fixed set of points ("sensors") to an object of interest ("tag"). In this work we consider the noise at the level of a set of sensors used in a single trilateration calculation. This noise changes when the set changes. In a Global Navigation Satellite System (GNSS), the set of sensors changes slowly because of the scale of the tracking system (see Figure 1). In an unobstructed area, the set of visible satellites changes approximately every 15 minutes (Hofmann-Wellenhof et al. 1997). However, in indoor ultra-wideband (UWB) indoor position tracking, sensor sets change with every new measurement (typically 100 ms). The sets change while moving around a single room, and sometimes even while standing still, depending upon the received signal strengths as shown in Figure 2.

The problem caused by set switching is illustrated in Figure 3. The sequence shows three consecutive trilateration calculations that use different sensor sets, each resulting in a different tracked location, even though the object of interest has not moved. At time t, the position of the tag is computed from sensors A, B and C. At time t+1, a new sensor set consisting of A, B, C and D is used to calculate the position of the tag. It can be observed that a change in the sensor set has caused a shift in the calculated position of the tag, due to the changing collective set of noises in the distance measurements. At time t+2, a new set of sensors consisting of A, B and E causes another shift in the calculated position. Hence, switching between sensor sets at each time instant adds a different noise to the measurements corresponding to the noise model of each sensor set. This causes a "jump" in the calculated position of the tag, even when the tag is not moving. A video of such a behavior occurring at our facility using a real UWB position tracking system can be seen at http://youtu.be/B-oCDTBQLd4.

Figure 1 Trilateration using a GNSS (earth-sized).

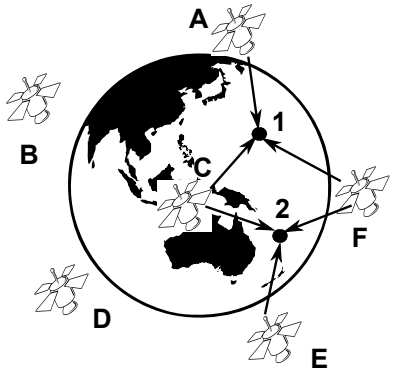

Figure 2 UWB indoor trilateration (building-sized).

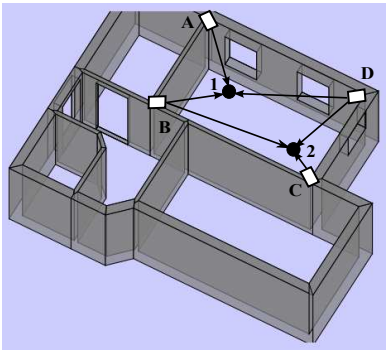

Figure 3 Changing noise due to sensor set switching.

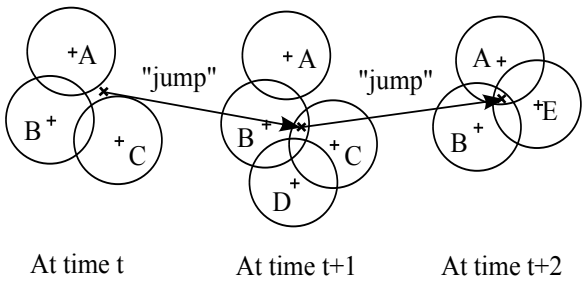

Previous works have examined several different sources of noise in UWB position tracking, but none have considered set switching. Zetik et al. (2004) examined mitigating the effects of multipath, synchronization, antenna effects and peak detection. Low et al. (2005) considered peak search and match filtering in indoor line-of-sight (LOS) signals to improve accuracy. Lie et al. (2005) combined an envelope detector with a leading-edge pulse detection method using a tunnel diode for UWB ranging in a multipath environment. Meier et al. (2007) suppressed multipath effects in highly reflective laboratory environment by detecting the LOS and NLOS signals by filtering the measurement noise using a Kalman filter. Mahfouz et al. (2008) also conducted experiments in a small test area and concentrated on reducing the effects of multipath, synchronization, antenna effects, peak detection and sensor placement. In order to determine the impact of each noise source, they conducted experiments to measure the position of the transmitter over a small (1 - 2 m) unobstructed area. Jing et al. (2010) studied the impact of propagation of signals through various building materials. Shen et al. (2010) proposed a new method to identify and mitigate the effects of non-line-of-sight (NLOS) signal propagation by comparing the mean square errors of the range estimates with the estimated LOS ranges. They conducted simulations to compare the performance of the proposed method with other methods in this area. Zhang et al. (2010) proposed an architecture for combining traditional energy-based and carrier-based detection schemes to minimize carrier phase noise and timing error effects for static and dynamic tracking. Further research by Yang et al. (2010) proposed a modified correlation algorithm to improve the accuracy and reduce the computational burden of systems. Caron et al. (2007) proposed different particle filters which can handle synchronous and asynchronous measurements received from different sensors in a multisensor system. These particle filters can switch between different observation models and also handle cases where sensors fail or their functioning changes. Denis et al. (2005) used modified extended Kalman filter and modified regularized particle filter to track biases caused by transitions from LOS to NLOS and from NLOS to NLOS conditions. It should be noted that some of these works were only theoretical and not applied to real world systems.

In this paper we identify noise due to the switching of sensor sets. This noise is present in any trilateration-based tracking system, but its effect is much more pronounced in an indoor positioning system. In preliminary work our group examined this issue in simulations (Ganjali 2009); this paper is the first to study it in a real system. After providing a mathematical model for this noise, we describe a particle filter for reducing its effect. We then demonstrate the operation of this particle filter on data from a real system, showing an approximately 15% improvement in accuracy over the raw measurements.

2 Methods

2.1 Facility

Our facility is located in the basement of Riggs Hall at Clemson University. The test area covers approximately 8 m × 8 m, covering the majority of a laboratory and part of an adjacent hallway. Figure 4 shows a picture of part of this area, where it can be observed that the laboratory and the adjacent hallway are separated by a stone wall which is approximately 20 cm thick. The walls are approximately 5 m high with false ceilings at a height of 3 m. The false ceilings are made up of thermocol and placed on metal railings. In addition, there are two metal mailboxes and a vending machine in the hallway, and two cupboards in the laboratory. Figure 5 shows the locations of furniture, walls and sensors in the test area.

Figure 4 The facility and location of one of the eight sensors.

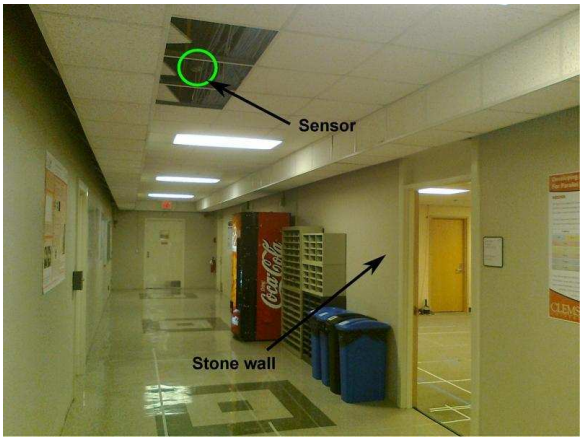

2.2 UWB positioning system

We used a commercially available UWB based local positioning system developed by Ubisense Inc. (Cambridge, U.K.). We installed eight Series 7000 sensors (Ubisense 2011*b*) in the facility at fixed locations. These sensors detect UWB pulses from Ubisense tags (Ubisense 2011*a*), which are tracked moving throughout the test area. Sensors are powered over network cabling using a Power-over-Ethernet switch. The Ubisense system uses a combination of angle of arrival and time difference of arrival, followed by multilateration or hyperbolic positioning, to calculate the position of a tag (Ubisense 2011*c*). The system chooses five sensors which have the highest demodulation power. We refer to this collection of five sensors as a sensor set.

Figure 5 shows the positions of the eight sensors distributed across our facility. The company recommends an install where the sensors are placed in a rectangular pattern surrounding the area of interest, with minimum NLOS conditions. However, the promise of UWB indoor position tracking is that it can be accomplished without direct LOS between the tracked object and fixed sensor

Figure 5 Layout of the facility (0–7 indicate UWB sensor positions).

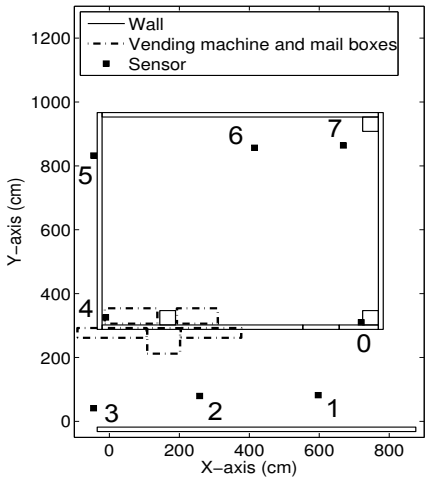

Figure 6 1000 measurements collected at a single location (450, 590, 92) cm.

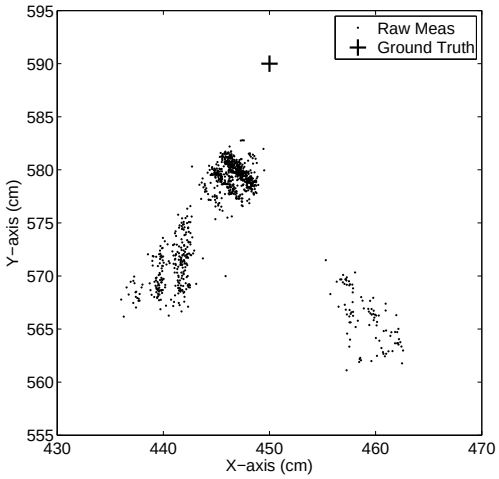

points. Our install purposefully introduces some NLOS conditions from the facility in order to explore this challenge. However, it must be noted that NLOS is not the only noise source that contributes to errors in distance measurements for trilateration, and that even with a completely LOS install, we have observed significant sensor set switching noise.

2.3 Noise model

2.3.1 Sensor set

We assume that a tracking measurement is calculated from a subset of the available sensors and each sensor set has a noise model associated with it. We model the noise associated with each sensor set independently. We model the total set of I sensors as $\{1, 2, ..., i_{max}\}$. Let a sensor set s represent any subset of size ≥ 5 sensors drawn from I, denoting a specific sensor subset. The total possible subsets can grow large as the number of sensors I grows. For example, if $\parallel I \parallel = 8$, then there are a possible total of $s_{max} = \binom{8}{5} = 56$ sensor sets. However, we assume that a relatively small number of sensor sets dominates the possibilities used for tracking measurements. Figure 6 shows a plot of 1,000 measurements made by our Ubisense system, all at a single ground truth location shown by a '+'. The dots represent the actual measurements received from the system. Figure 7 shows the frequency distribution of the sensor sets for these measurements. The most common 4 sensor sets account for 795 measurements, or 79.5% of the data. These 4 sensor sets are '76540', '75420', '65410' and '65420', where the numbers indicate the sensors (see figure 5) used to provide a measurement. The most commonly occurring sensor set corresponds to the sensors that are most LOS, and therefore generally the most powerful signals. However, it accounts for less than 45% of the total data. The second most commonly occurring sensor set contains a sensor which lies in the hallway and contributes approximately 20% of the measurements. Similarly, other sensor sets contain at least one sensor which lies in the hallway, providing better angular coverage but more NLOS conditions.

Figure 7 Distribution of measurements.

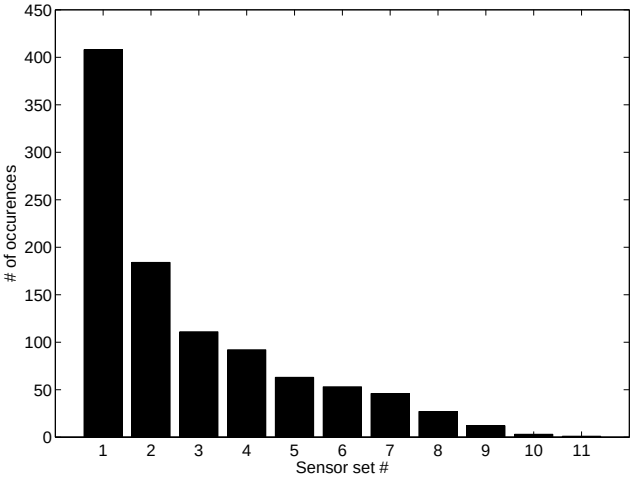

Figure 8 shows the measurements for the top four sensor sets from the data collected in Figure 6. The measurements from each sensor set appear clustered at different offsets from the ground truth position. To capture this behavior, we use a Gaussian noise model with non-zero mean:

$$\mathcal{N}(\mu_x^s, \mu_y^s, \sigma_x^s, \sigma_y^s) \tag{1}$$

The noise model for each sensor set is given by $(\mu_x^s, \mu_y^s, \sigma_x^s, \sigma_y^s)$ where (μ_x^s, μ_y^s) corresponds to the average shift of the measurements from the sensor set relative to the ground truth location, and (σ_x^s, σ_y^s) corresponds to the standard deviation of the measurements from their mean. The length of the axes of the ellipses in figure 8 correspond to three standard deviations. The use of other noise models is possible and is discussed in the conclusion.

Figure 8 Noise models of four sensor sets at (X,Y,Z) = (450,590,92) cm.

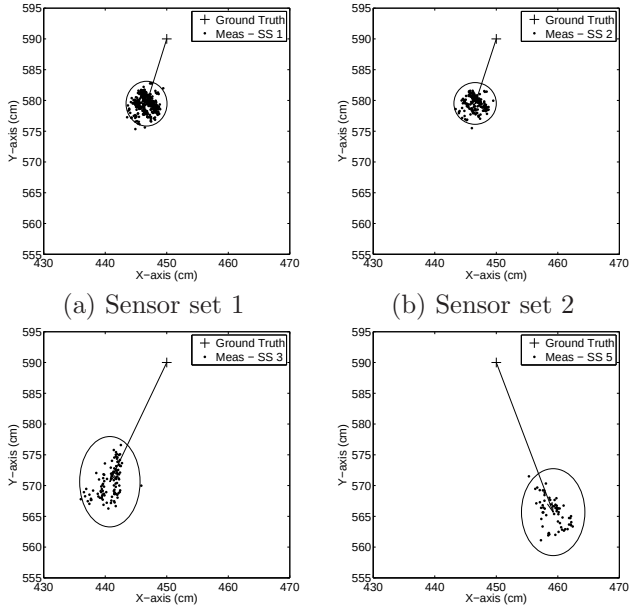

(a) Sensor set 1 (b) Sensor set 2

(c) Sensor set 3 (d) Sensor set 4

2.3.2 Calibration

In order to calculate our noise model, we conduct a calibration step. A tag is placed at a known location, and 1000 measurements are collected. This process is repeated at 6 different locations distributed throughout the facility. The noise parameters $(\mu_x^s, \mu_y^s, \sigma_x^s, \sigma_y^s)$ for each sensor set are calculated at each location, and then weighted-averaged by the number of measurements for each sensor set across the 6 locations. At a single location, if a sensor set has less than 30 measurements then no noise model is calculated at that location. After weighted-averaging, some sensor sets may have no model. We therefore also calculate a facility-wide noise model that is used by default for measurements taken from a non-modeled sensor set. The facility-wide noise model is taken as the average of all measurements taken during the calibration step.

Figure 9 shows how calibration data was collected. A tag was placed on a wooden sawhorse 92 cm in height. This height was chosen since it corresponds to the typical height of the waist of an upright person. The base of the setup is marked at intervals of 10 cm so that it can be aligned with markings on the floor of the facility. This enables accurate positioning of the sawhorse with respect to the ground truth coordinate system. We used laser levels and tape measures to ensure that the ground truth locations are accurate to within 1 cm relative to the calibrated Ubisense coordinate system.

2.4 Basic particle filter

Before describing how we use our noise model, we describe a basic particle filter algorithm. In our experiments we compare the performance of our set noise particle filter to this basic particle filter.

Figure 9 Setup used to collect calibration data.

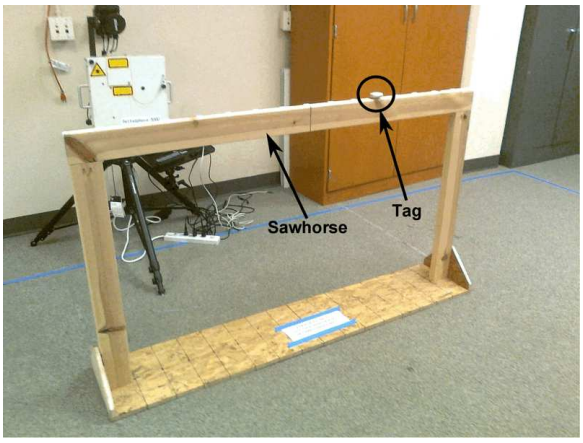

We assume a 2D linear constant velocity model in our experiments. Let the state of the system $\mathbf{X}$ be defined as

$$\mathbf{X} = \begin{bmatrix} x_t \\ \dot{x}_t \\ y_t \\ \dot{y}_t \end{bmatrix} \tag{2}$$

where, x_t, y_t are the positions along the x and y axes at time t and $\dot{x}_t, \dot{y}_t$ are the velocities.

The system transition equations $\boldsymbol{f}$ are

$$\boldsymbol{f} = \begin{bmatrix} x_{t+1} = x_t + T\dot{x}_t \\ \dot{x}_{t+1} = \dot{x}_t + \mathcal{N}(0, \sigma_d) \\ y_{t+1} = y_t + T\dot{y}_t \\ \dot{y}_{t+1} = \dot{y}_t + \mathcal{N}(0, \sigma_d) \end{bmatrix} \tag{3}$$

where $\mathcal{N}(0, \sigma_d)$ is a continuous zero-mean Gaussian random variable. The dynamic noise $\mathbf{U_t}$ denotes the dynamic noise during a state transition:

$$\mathbf{U}_t = \begin{bmatrix} 0 \\ \mathcal{N}(0, \sigma_d) \\ 0 \\ \mathcal{N}(0, \sigma_d) \end{bmatrix} \tag{4}$$

The dynamic noise models a potential change in velocity during each time step.

At each time t, the set of observed values $\mathbf{Z}$ is

$$\mathbf{Z} = \begin{bmatrix} \tilde{x}_t \\ \tilde{y}_t \end{bmatrix} \tag{5}$$

where $\tilde{x}_t, \tilde{y}_t$ is the measured position of the tag. The observation equations $\boldsymbol{g}$ are

$$\boldsymbol{g} = \begin{bmatrix} \tilde{x}_t = x_t + \mathcal{N}(0, \sigma_{n_x}) \\ \tilde{y}_t = y_t + \mathcal{N}(0, \sigma_{n_y}) \end{bmatrix} \tag{6}$$

where it is assumed that a 2D zero-mean Gaussian random noise with standard deviation $\sigma_{n_x}, \sigma_{n_y}$ has been added to the actual position to produce the measurement.

Particle filtering is a sequential Monte Carlo methodology where the posterior density function is

recursively approximated using a set of random samples and weights, from which estimates are computed (Arulampalam et al. 2002, Djuric et al. 2003). The number of samples K depends on how accurately we want to model the posterior density; in this work we use $K = 1000$.

In our experiments, we initialize all particles to have equal weight and the same initial state:

$$\chi = \left\{\mathbf{X}^k, w^k\right\}_{k=1}^{K} = \left\{ \begin{bmatrix} x_0 \\ 0 \\ y_0 \\ 0 \end{bmatrix}, \frac{1}{K} \right\} \tag{7}$$

where x_0 and y_0 are the known starting position of the tag along the x and y axes with zero initial velocities.

In an iteration of the particle filter, each particle is transitioned through the system transition equations $\boldsymbol{f}$ with a randomly computed dynamic noise $\mathbf{U}_t^k$:

$$\left\{\mathbf{X}_t^k\right\}_{k=1}^{K} = \boldsymbol{f}\left\{\mathbf{X}_{t-1}^k, \mathbf{U}_t^k\right\}_{k=1}^{K} \tag{8}$$

An observation $\mathbf{Z}_t$ is taken. We use sequential importance sampling using the prior importance function Djuric et al. (2003), so that the weights for the particles are updated according to

$$w_t^k = w_{t-1}^k p(\mathbf{Z}_t|\mathbf{X}_t^k) \tag{9}$$

Based upon equation 6, we assume $p(\mathbf{Z}_t|\mathbf{X}_t^k)$ is a 2D Gaussian that can be written as

$$p(\mathbf{Z}_t|\mathbf{X}_t^k) = \exp - \left(\frac{(x_t^k - \tilde{x}_t)^2}{2(\sigma_{n_x})^2} + \frac{(y_t^k - \tilde{y}_t)^2}{2(\sigma_{n_y})^2} \right) \tag{10}$$

Equation 10 calculates the likelihood of obtaining the actual observed measurement relative to the most probable measurement of the particle, according to the measurement noise distribution.

Collectively, the discrete values of the particles represent a continuous probability distribution function (pdf) of the state at a particular time. After updating the weights, they must be normalized:

$$\left\{ w_t^k = \frac{w_t^k}{\sum_{i=1}^{K} w_t^i} \right\}_{k=1}^{K} \tag{11}$$

We then calculate the expected value of the set of particles as output:

$$E[\chi] = \sum_{k=1}^{K} \mathbf{X}^k w^k \tag{12}$$

Particle filtering is well-known to suffer from a degradation over time where particle weights tend towards zero. In order to counter this effect, we follow the general strategy of resampling when a sizable percentage of particles reaches negligible weight. Following methods outlined in Rekleitis (2004), we calculate the coefficient of variation (CV) as

$$CV = \frac{VAR(w^k)}{(E[w^k])^2} \qquad = \frac{1}{K}\sum_{k=1}^{K}(Kw^k - 1)^2 \tag{13}$$

The effective sample size is calculated as $ESS = \frac{K}{1+CV}$. This factor indicates if there are sufficient particles having appreciable weight and whether resampling is necessary. If $ESS < 0.5$, we resample using the select with replacement algorithm (Rekleitis 2004).

2.5 Set noise particle filter

Using our sensor set noise model, the basic particle filter algorithm is adjusted as follows. At each time t, the set of observed values $\mathbf{Z}$ is

$$\mathbf{Z} = \begin{bmatrix} \tilde{s}_t \\ \tilde{x}_t \\ \tilde{y}_t \end{bmatrix} \tag{14}$$

where $\tilde{s}_t$ is the sensor set used to measure $\tilde{x}_t, \tilde{y}_t$. The observation equations $\boldsymbol{g}$ are

$$\boldsymbol{g} = \begin{bmatrix} \tilde{s}_t \leftarrow \{1, 2, ..., s_{max}\} \\ \tilde{x}_t = x_t + \mathcal{N}(\mu_x^{\tilde{s}_t}, \sigma_x^{\tilde{s}_t}) \\ \tilde{y}_t = y_t + \mathcal{N}(\mu_y^{\tilde{s}_t}, \sigma_y^{\tilde{s}_t}) \end{bmatrix} \tag{15}$$

where it is assumed that a random non-zero mean Gaussian noise associated with sensor set $\tilde{s}_t$ has been added to the actual position to produce the measurement.

The weight update step given in equation 10 is replaced with

$$p(\mathbf{Z}_t|\mathbf{X}_t^k) = \exp - \left(\frac{((x_t^k - \mu_x^{\tilde{s}_t}) - \tilde{x}_t)^2}{2(\sigma_x^{\tilde{s}_t})^2} + \frac{((y_t^k - \mu_y^{\tilde{s}_t}) - \tilde{y}_t)^2}{2(\sigma_y^{\tilde{s}_t})^2} \right) \tag{16}$$

where $(x_t^k - \mu_x^{\tilde{s}_t})$ and $(y_t^k - \mu_y^{\tilde{s}_t})$ gives the most probable measurement of each particle. Equation 16 calculates the likelihood of obtaining the actual observed measurement relative to the most probable state of the particle, according to the measurement noise distribution associated with the sensor set used to take the measurement.

All the other steps are the same as described for the basic particle filter.

2.6 Data collection

Figure 10 shows the apparatus used to record experimental data for testing. A tag was placed on a tripod resting on a trolley. The tripod was adjusted so that it would match up to the same height (92 cm) used to collect calibration data. The trolley was pulled manually by a person standing at maximum arm length in order to minimize any disturbances in the measurements from the presence of a person. The trolley was moved along a track laid on the ground at different speeds. For each recording, the apparatus was pulled back and forth seven times along a 250 cm straight line. The total distance covered in each recording is 1750 cm.

Figure 10 Setup used to collect recordings.

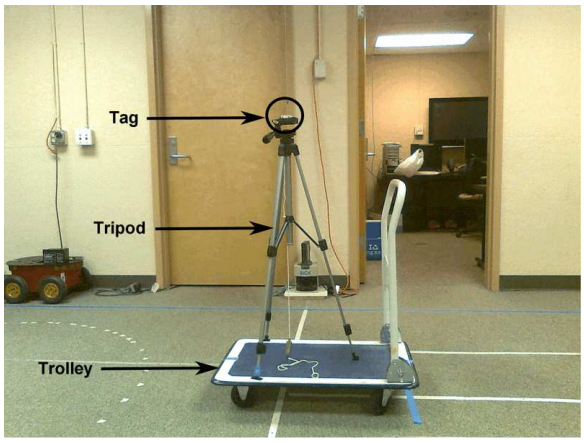

Figure 11 Test tracks in the facility.

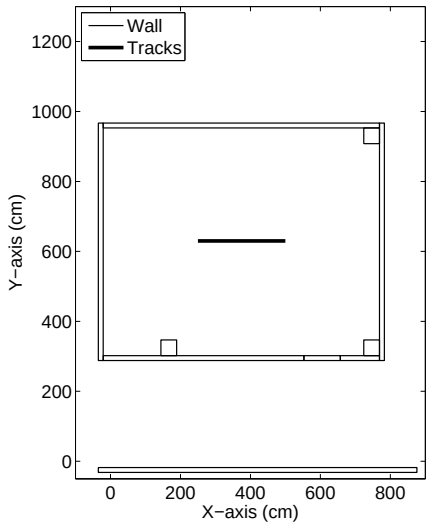

Figure 11 shows the location of the track in the test area. For each recording along the track, the Ubisense system provides raw measurements of the tag along the X and Y axes, and the sensors used to calculate each measurement. By default, the Ubisense system uses a simple averaging filter to remove outliers by comparing each measurement to the average and variance in position of previous measurements. We turned off this filter to avoid any potential confusion in understanding the noise due to sensor set switching, and collected raw measurements. In practice we have observed that the noise due to sensor set switching does not appear similar to outlier noise.

Five recordings were collected along the track at different speeds. The speed was varied from ≈11 cm/s (extremely slow motion) to ≈120 cm/s (walk speed) (Knoblauch et al. 1996, Waters & Mulroy 1999). Table 1 lists the recording number and the approximate speed of the recording. The speeds were chosen to test the viability of our method for a range of motion dynamics resembling a slow moving robot to the walking of a person.

Table 1 Range of motions

Recording #	Total Measurements	Speed (cm/s)	Raw Error (cm)
1	1521	11	20
2	467	35	22
3	250	65	23
4	164	100	25
5	135	120	23

2.7 Ground Truth

We use a least squares approach to calculate the ground truth data. The tag is initially placed in the start position and measurements are collected for 15 seconds at this position before moving the tag. After the tag reaches the end position, we wait for another 15 seconds before ending the recording. This gives us "flat" regions near the start and end positions along with linear regions in between which indicate the movement of the tag. Now, a subset of measurements in each region is used to provide a least squares fit to the set of measurements in that region. The measurement closest to the intersection of the lines determines the start, end or change in dynamics (change in direction of the tag). Figure 12 shows a partial output of this approach. In this figure, we can observe the flat region indicating the start region and two linear regions indicating the first two passes of motion long the track. (A full trial has seven passes of motion, with a single flat area at the beginning and end.)

Since we assume that the tag is moving with a constant velocity, we can associate each measurement with a ground truth location. The velocity is calculated by dividing the total ground truth distance covered by the total time taken when the tag is in motion. Now, multiplying the velocity by the time at which the measurement was received gives us the corresponding ground truth position of the tag at that time instant. This can be written as

$$\breve{x}_t = \frac{D_x}{T} \times t \tag{17}$$

$$\breve{y}_t = \frac{D_y}{T} \times t \tag{18}$$

where $\breve{x}_t, \breve{y}_t$ are the ground truth data at time t, D_x, D_y are the total ground truth distances along x and y axes respectively, and T is the total time taken to complete a recording .

2.8 Error metric

In order to evaluate the performance of the filtered output we calculate the average Euclidean distance between the filtered data and the corresponding ground truth data over the total number of measurements. This distance is known as the position error (P.E.) and can be defined as

$$P.E. = \frac{1}{N} \sum_{i=1}^{N} \sqrt{(x_i - \breve{x}_i)^2 + (y_i - \breve{y}_i)^2} \tag{19}$$

Figure 13 Illustration of the effect of σ_d on filter output.

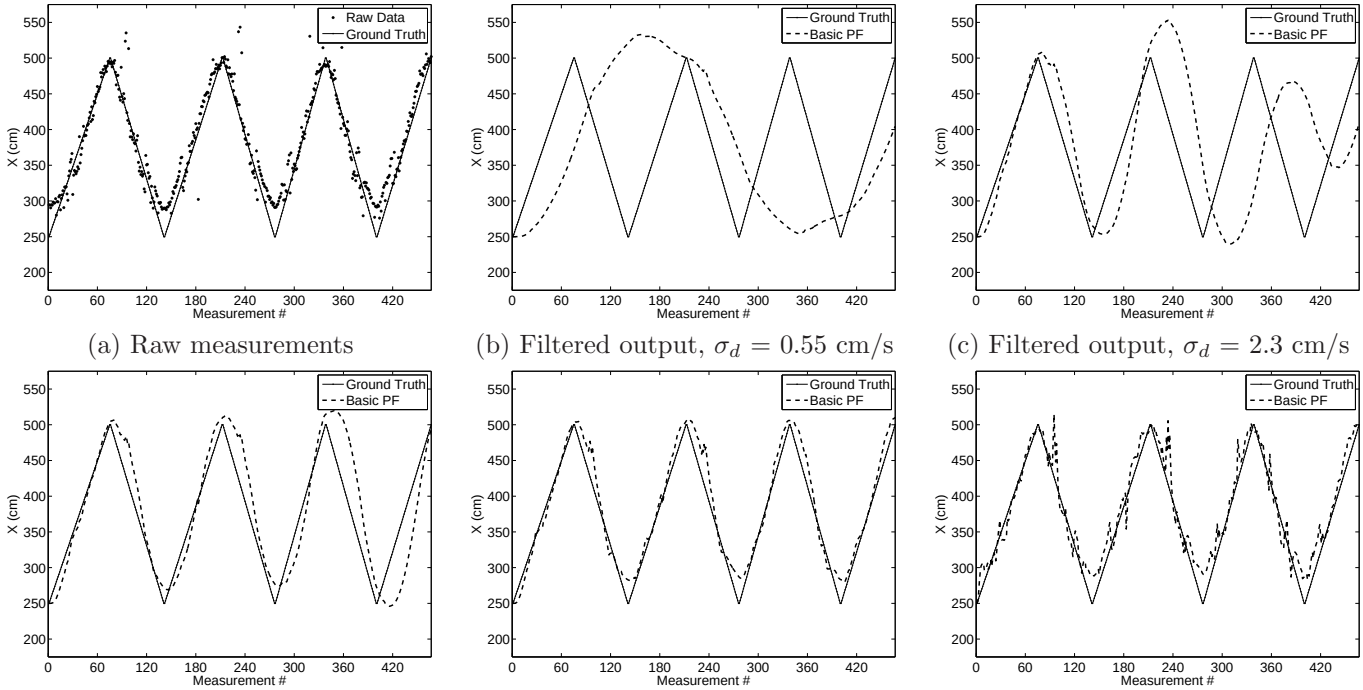

(a) Raw measurements (b) Filtered output, $\sigma_d = 0.55$ cm/s (c) Filtered output, $\sigma_d = 2.3$ cm/s

(d) Filtered output, $\sigma_d = 2.8$ cm/s (e) Filtered output, $\sigma_d = 3.7$ cm/s (f) Filtered output, $\sigma_d = 55.0$ cm/s

Figure 12 Least squares approach to generate ground truth data (partial output).

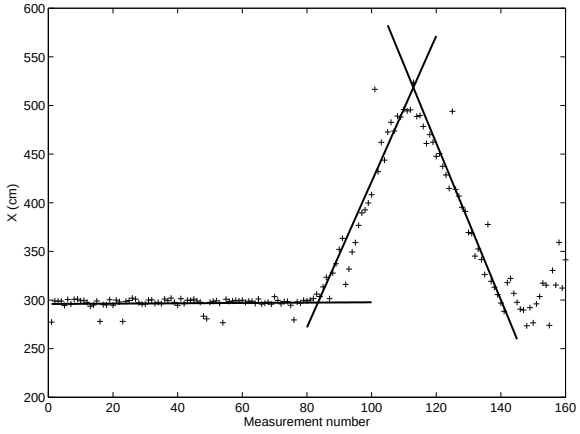

where, x_i, y_i are the filtered data, $\breve{x}_i, \breve{y}_i$ are the corresponding ground truth data for measurement i and N is the total number of measurements.

3 Experimental Results

The output for the particle filter is to some degree controlled by the value chosen for σ_d (see equation 4), the dynamic noise in the motion model. This value represents the amount of expected change in velocity at each time step. The lower this value, the more the filter weights the output towards the system equations, in essence providing more smoothing. The higher this value, the more the filter weights the output towards the measurements, allowing a quicker reaction to actual dynamics at the cost of less smoothing. Figure 13 demonstrates this effect. Part (a) shows the raw measurements for a recording; parts (b)-(f) show the basic particle filter output for increasing values of σ_d. In part (b), the actual dynamics in the recording (the step changes in velocity) are considerably larger than the value chosen for σ_d, so that the filter output is not able to reliably track the motion. In effect, the filter is smooothing too much. In parts (c)-(d), the larger values for σ_d cause the filter to more reliably track the actual motion, but there is a noticeable lag, particularly at the points where the actual dynamics change. In part (e), σ_d most closely matches the actual dynamics so that the optimal filter output is obtained. Part (f) shows the output for an even larger value of σ_d, where the filter is giving too much weight to individual measurements.

Figure 15 shows the error curves comparing the raw measurements, basic particle filter, and set noise particle filter, for one recording. The error is shown over a range of $\sigma_d = 0.01$ to 60 cm/s. It is important to evaluate performance across a range of σ_d because in practice the actual dynamics of tracked targets are unknown (for example, people can walk at varying speeds). The error curves are the average errors of the recordings over 100 trials (repeated runs of each filter at each value of σ_d); this is necessary because the particle filter is a Monte Carlo approach and a single trial of limited length does not necessarily provide a typical representative output. From the figure, it can be seen that the set noise model particle filter performs better than basic particle filter over the entire range of dynamic noise. The minimum error is 15 cm, and occurs at approximately $\sigma_d = 6.0$

Figure 14 Recording 2 at $\sigma_d = 6.0$ cm/s (partial output)

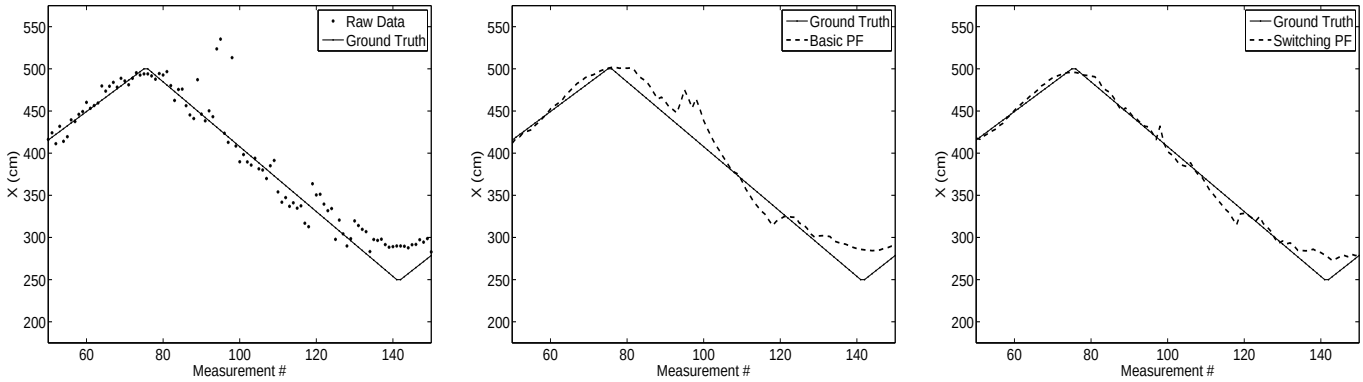

(a) Raw measurements (b) X - axis for basic particle filter (c) X - axis for set noise particle filter

Figure 15 Error curve for one recording.

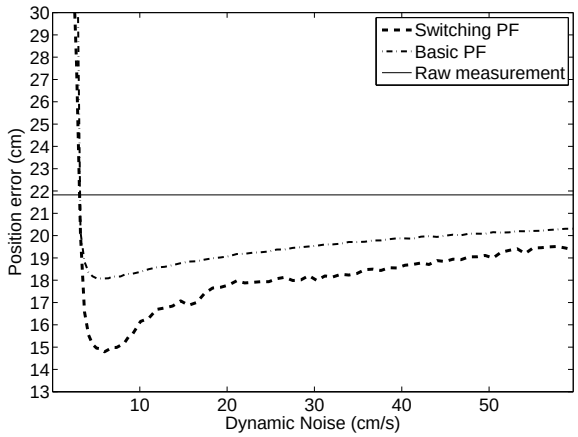

Figure 16 Average error curve for all recordings.

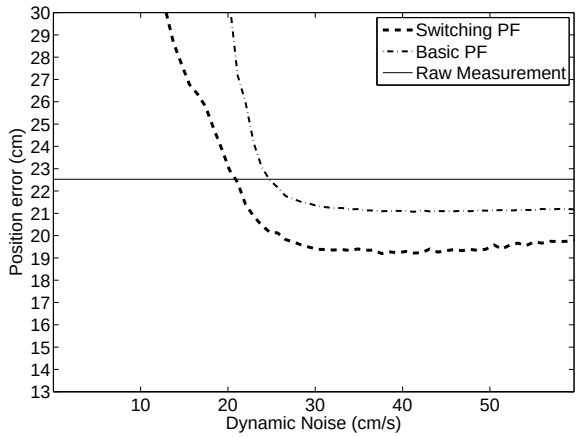

cm/s, where σ_d best matches the actual dynamics of the motion for this recording.

Figure 14 shows the basic particle filter and set noise particle filter output for this recording at $\sigma_d = 6.0$ cm/s. For clarity, only a subset of the data is shown, and only the X-coordinates are shown (the motion is along a straight line of constant Y). Because the optimal value of σ_d was chosen for this figure, both filters provide a fairly good output that is better than the raw measurements. However, it can also be seen that the set noise particle filter output is more accurate, particularly in the range of measurements from 80 to 120.

Figure 16 shows the average error curves for all five recordings. From this figure, it can be observed that the set noise particle filter performs better than the basic particle filter across the entire range of dynamic noise. There is no global minimum like there was in Figure 15 because the actual dynamics of all the recordings vary (see Table 1). For the range of motions tested, any value of σ_d between 30 to 50 cm/s is appropriate. The average accuracy of the raw measurements is approximately 23 cm. In the range of $\sigma_d = 30$ to 50 cm/s, the set noise particle filter improved the accuracy of the raw measurements by approximately 4 cm on average, about double that of basic particle filter. Thus, our set noise particle filter shows an approximately 15% improvement over a basic particle filter.

4 Conclusion and Future Work

The contribution of this paper is that we have identified a new noise source due to the switching of fixed point sets for trilateration. While this noise is theoretically present in all trilateration-based systems, it is not readily apparent in large-scale systems like a GNSS, but it can cause noticeable jump-like behavior in indoor UWB position tracking. We have developed a mathematical model and particle filter that accounts for this noise. We tested our methods on a real UWB indoor position tracking system. Our set noise particle filter showed an approximately 15% improvement in accuracy over the raw measurements. We use a particle filtering approach since it can handle non-zero mean noise models, unlike Kalman filtering approaches. In the experiments, we used Gaussian noise models for the sensor sets. However, our approach could be applied using other noise models by substitution for $\mathcal{N}$ in equation 15 and the reformulation of equation 16 to account for the alternative noise distribution.

Our experiments have been conducted in a real world setting, where we have achieved modest improvement in the tracking accuracy over a range of dynamics. However, we have observed that noise due to NLOS, multipath and timing errors tend to be more significant than noise caused by sensor set switching. Hence, research studying the impact of these noise sources in isolation have been able to achieve sub-decimeter (Low et al. 2005, MacGougan et al. 2010) and sometimes even sub-centimeter accuracies (Mahfouz et al. 2008, Zetik et al. 2004), while we have been able to observe only a modest improvement. In practice, we believe that our set noise filter should be combined with filters specifically designed for other noise sources. This is a subject for future work.

References

Arulampalam, M., Maskell, S., Gordon, N. & Clapp, T. (2002), 'A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking', *IEEE Transactions on Signal Processing* **50**(2), 174 –188.

Caron, F., Davy, M., Duflos, E. & Vanheege, P. (2007), 'Particle filtering for Multisensor Data Fusion With Switching Observation Models: Application to Land Vehicle Positioning', *IEEE Transactions on Signal Processing* **55**(6), 2703–2719.

Denis, B., Ouvry, L., Uguen, B. & Tchoffo-Talom, F. (2005), Advanced bayesian filtering techniques for uwb tracking systems in indoor environments, *in* 'IEEE International Conference on Ultra-Wideband', p. 6 pp.

Djuric, P., Kotecha, J., Zhang, J., Huang, Y., Ghirmai, T., Bugallo, M. & Miguez, J. (2003), 'Particle filtering', *IEEE Signal Processing Magazine* **20**(5), 19 – 38.

Ganjali, D. (2009), Filtering Noise Caused by Sensor Selection for an Ultra-Wideband position tracking system, Master's thesis, Clemson University.

Hofmann-Wellenhof, B., Lichtenegger, H. & Collins, J. (1997), *Global Positioning System: Theory and Practice*, 4 edn, Springer.

Jing, M., Nai-tong, Z. & Qin-yu, Z. (2010), 'IR-UWB Waveform Distortion Analysis in NLOS Localization System', *Information Technology Journal* **9**, 139–145.

Knoblauch, R., Pietrucha, M. & Nitzburg, M. (1996), 'Field Studies of Pedestrian Walking Speed and Start-Up Time', *Transportation Research Record* **1538**(1), 27–38.

Lie, J., See, C. M. & Ng, B. P. (2005), 'UWB Ranging with High Robustness Against Dominant Jammer and Multipath', *Microwave and Wireless Components Letters, IEEE* **15**(12), 907 – 909.

Low, Z., Cheong, J., Law, C., Ng, W. & Lee, Y. (2005), 'Pulse detection algorithm for line-of-sight (LOS) UWB ranging applications', *Antennas and Wireless Propagation Letters, IEEE* **4**, 63 – 67.

MacGougan, G., O'Keefe, K. & Klukas, R. (2010), 'Tightly-coupled GPS/UWB Integration', *The Journal of Navigation* **63**(01), 1–22.

Mahfouz, M., Zhang, C., Merkl, B., Kuhn, M. & Fathy, A. (2008), 'Investigation of High-Accuracy Indoor 3-D Positioning Using UWB Technology', *IEEE Transactions on Microwave Theory and Techniques* **56**(6), 1316 –1330.

Meier, C., Terzis, A. & Lindenmeier, S. (2007), A robust 3D high precision radio location system, *in* 'IEEE/MTT-S International Microwave Symposium', pp. 397 –400.

Rekleitis, I. M. (2004), A particle filter tutorial for mobile robot localization, Technical Report TR-CIM-04-02, Technical report, Centre for Intelligent Machines, McGill University.

Shen, G., Zetik, R., Hirsch, O. & Thomä, R. (2010), 'Range-Based Localization for UWB Sensor Networks in Realistic Environments', *EURASIP Journal on Wireless Communications and Networking* .

Ubisense (2011*a*), 'Series 7000 Compact Tag', `http://www.ubisense.net/en/resources/factsheets/series-7000-compact-tag.html`.

Ubisense (2011*b*), 'Series 7000 Sensor', `http://www.ubisense.net/en/resources/factsheets/series-7000-ip-sensors.html`.

Ubisense (2011*c*), 'Ubisense Precise Location', `http://www.ubisense.net/en/resources/factsheets/ubisense-precise-location.html`.

Waters, R. L. & Mulroy, S. (1999), 'The energy expenditure of normal and pathologic gait', *Gait & Posture* **9**(3), 207 – 231.

Yang, D., Fathy, A., Li, H., Mahfouz, M. & Peterson, G. (2010), Millimeter accuracy UWB positioning system using sequential sub-sampler and time difference estimation algorithm, *in* 'IEEE Radio and Wireless Symposium (RWS)', pp. 539 – 542.

Zetik, R., Sachs, J. & Thoma, R. (2004), UWB localization - active and passive approach [ultra wideband radar], *in* 'Proceedings of the 21st IEEE Instrumentation and Measurement Technology Conference', Vol. 2, pp. 1005 – 1009.

Zhang, C., Kuhn, M., Merkl, B., Fathy, A. & Mahfouz, M. (2010), 'Real-Time Noncoherent UWB Positioning Radar With Millimeter Range Accuracy: Theory and Experiment', *IEEE Transactions on Microwave Theory and Techniques* **58**(1), 9–20.