14th European Signal Processing Conference (EUSIPCO 2006), Florence, Italy, September 4-8, 2006, copyright by EURASIP

POSITION AND VELOCITY RECOVERY FROM INDEPENDENT ULTRASONIC BEACONS

Michael McCarthy, Henk L Muller, Andrew Calway, and R Eddie Wilson

Deparmtent of Computer Science, University of Bristol
Merchant Venturers Building, Woodland Road, Bristol, BS8 1UB, UK
phone: +44 117 954 5104, fax: +44 117 954 5108, email: mccarthy@cs.bris.ac.uk
web: http://www.cs.bris.ac.uk/

ABSTRACT

In this paper we present a study on how to estimate the position of a mobile receiver using ultrasonic beacons fixed in the environment. Unlike traditional approaches, the ultrasonic beacons are independent, and positioning is performed by measuring the Doppler shift within their observed periods. We show that this approach allows us to deduce both position and velocity, but an analysis of the space indicates that we can recover the direction of velocity very well, the magnitude of velocity less well, and that location estimation is the least accurate. Based on the characteristics of the solution space, we suggest a method for improving positioning accuracy.

1. INTRODUCTION

Indoor location systems have been developed for a variety of purposes. Early location systems were built to allow autonomous robots navigate through buildings (for example [3]). More recently, location systems have been developed by the ubiquitous computing community to track people and objects, and by the mobile and wearable community to allow mobile devices to position themselves. In both domains of research, the systems provide support for location based services. One such service takes the form of a guide in a museum or art gallery. For example, a location based application running on a device worn by a tourist has the potential to enhance the experience of a visit. Such an application would remove the need for visual references in the environment that are usually in the form of markers, descriptions or maps. A properly designed positioning system can therefore improve the appearance of displays and give the curator more freedom in the design of the experience.

There are a large number of decisions affecting the design of an indoor location system. The system designer must select the sensing technology used (video, RF, ultrasound), determine whether or not it is feasible or desirable to modify the environment, and decide on the orientation of the architecture: infrastructure-centric (tracking) or user-centric (positioning). Camera based systems can identify and track features in a scene to estimate the location of a mobile camera [2, 6]. They are appealing to some applications because they do not require modification of the environment. Other systems use a combination of RF and ultrasound to estimate distances between receivers and transmitters. These systems can be configured as infrastructure-centric [8] where the transmitter is tracked, or user-centric [7] where the receiver tracks itself. Most of these systems use techniques where distance or range based measurements are the basis for position estimation.

Funding for this work is received from the U.K. Engineering and Physical Sciences Research Council as part of the Equator IRC, GR-N-15986.

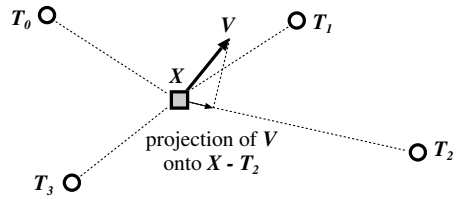

Figure 1: A receiver at position $\mathbf{X}$ with velocity $\mathbf{V}$ moving relative to beacons positioned at $\mathbf{T_i}$

The particular application motivating our research requires the design of a light-weight, easy-to-install location system that can be retrofitted to existing installations, such as museums or art galleries. We have opted to use narrowband ultrasonic beacons fixed at known locations, and carry an ultrasonic receiver with a processing unit to process the signals emitted by the beacons. Compared with other solutions, our beacons are completely independent with no wiring between them, they are low power such that they can be run from small solar cells, and they are made from readily available components that make them cheap to produce.

In this paper we discuss the signal processing needed to recover the location and velocity of our receivers. In particular, we discuss the type of data that we receive and the limitations of the algorithms that we use. We explore the solution space, and show that direction of travel seems to be the easiest to estimate, followed by speed then location. We have used both a Kalman Filter and Particle Filter solution and suggest why the latter seems to be more amenable to this specific problem.

2. LOCATION DETECTION USING USING DOPPLER PERIODICITY

The approach that we have taken emphasises hardware simplicity. We use off-the-shelf 40 kHz narrowband ultrasonic transducers on the beacons and receivers. The beacons are fixed on the ceilings and walls of a room, while the receiver is mounted on a mobile device.

In operation, each beacon "chirps" with a certain periodicity, for example 501 or 507 ms. The receiver unit classifies chirps arriving at regular intervals as being transmitted from a particular beacon and monitors the differences from the expected periodicities in order to establish whether the receiver

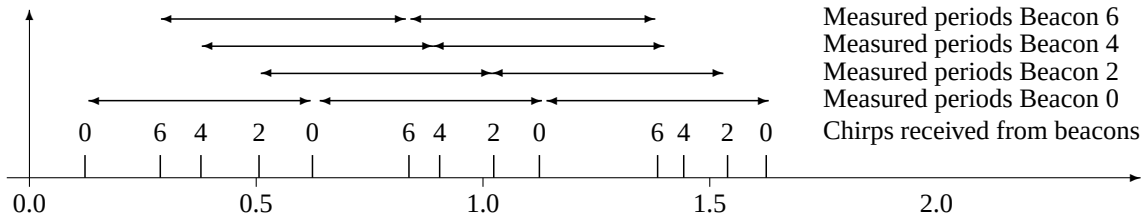

Figure 2: Chirp trace received from four beacons, and the resulting measurement periods.

has moved towards a beacon, or away from it.

2.1 Doppler shift

Each beacon is programmed to send a 250 μs ultrasonic pulse every $\sim$500 ms. The precise periodicity is dependent on the beacon – the period behaves as a unique signature that allows the receiver to identify the source of each chirp. Given that our mobile receiver will not travel faster than 1-2 ms^{-1}, the shift in periodicity will be limited to a fraction of a percent, allowing us to identify the signature of the chirp train [5]. The receiver is equipped with a pick-up, two op-amps, and a PIC micro-controller. When a signal is received, the PIC records the time and passes this on to an attached processing unit.

The method that we use for estimating the position and velocity of the receiver involves measuring the variation in periodicity of each of the beacons. For example, if the receiver moves towards a beacon, the observed periodicity of that beacon will be reduced within the time frame of the period. This happens because chirps arriving later do not take as long to reach the receiver. The reverse is true when the receiver moves away from a beacon: the period will increase depending on the velocity of the receiver. As shown in Figure 1, a moving receiver will move towards a number of beacons while simultaneously moving away from others. The period shift ΔP observed for each of the beacons depends on the location of a beacon ($\mathbf{T}$), and the position ($\mathbf{X}$) and velocity ($\mathbf{V}$) of the receiver:

$$\frac{\Delta P}{P+\Delta P} v_s = \frac{\mathbf{X}-\mathbf{T}}{|\mathbf{X}-\mathbf{T}|} \cdot \mathbf{V} \qquad (1)$$

The right part of Equation 1 is the projection of the velocity vector $\mathbf{V}$ onto the unit vector $\mathbf{X}-\mathbf{T}$. This is the magnitude of the receiver velocity in the direction of the beacon, which is also proportional to the observed periodicity shift. The relative distance that the receiver moves in the time between successive observations from a single beacon is $\Delta P v_s$, where v_s is the speed of sound. Dividing by the elapsed time $P+\Delta P$ gives the magnitude of velocity in the direction of the beacon expressed in terms of ΔP, as is shown on the left side of the equation. This equation relates six unknowns ($\mathbf{X}$ and $\mathbf{V}$), to a single measurement ΔP. Once we receive measurements from n beacons, we will obtain n equations, with six unknowns, which can be solved if there are no dependencies between equations, and if $n \geq 6$. In practice we use eight transmitters ($n = 8$), and we place the transmitters in an irregular pattern, avoiding dependencies.

2.2 Measurement restrictions

There are two issues that make it difficult to solve the sets of equations. First, the right-hand side of Equation 1 is expressed in terms of the instantaneous position and velocity of the receiver. Therefore, given an $\mathbf{X}$ and $\mathbf{V}$, the left-hand side of the equation should provide a measure of instantaneous scalar velocity in the direction of the beacon. However, the left-hand side of Equation 1 measures the Doppler shift over a period P. As a result, we measure an "average" position and velocity over the period P.

The second problem is that each chirp arrives asynchronously as the receiver is moving. Hence, the n equations will, strictly speaking, have $6n$ unknowns; we have to make an assumption that $\mathbf{X}$ and in particular $\mathbf{V}$ do not change dramatically between measurements, which is the case if we sample the system often enough. Each of our eight beacons chirps at a periodicity of around 520 ms, giving us an average update rate of 65 ms, or 16 Hz.

The two issues interact, as the measurement of each beacon overlaps with the measurement of the other beacons. This is depicted in Figure 2, which shows the arrival of chirps from four beacons over a 1.5 second period. The measured periods of each of the four beacons is shown. Note that the measurements shift relative to each other since the beacons transmit asynchronously with unique periods. If we only use Equation 1 to estimate location and velocity, we will end up with approximations that are averaged over a, typically, $2P$ period of about one second.

3. SAMPLING THE SPACE

In order to explore the nature of the Doppler equation (Equation 1) we have performed a number of Monte Carlo simulations. In these experiments, we randomly sample position and velocity pairs and compare their ΔP values with a reference pair. The aim of the experiments is to determine how the Doppler equation constrains values for position and velocity. We assume that we can take instantaneous measurements for ΔP, and that we can identify the source of each of the measurements with 100% accuracy.

3.1 Experimental set up

The experiments use two different eight-beacon configurations that cover a virtual room. The first configuration situates the beacons on the corners of a $4 \times 4 \times 4$ metre cube at locations $(-2,-2,-2),(-2,-2,2),...(2,2,2)$. The second configuration simulates the set-up that we have in our lab, with the eight beacons placed on the walls and ceiling.

The reference pair is placed near the centre of the space, at location $(-0.5, 0.5, 0.4)$ m and we assume that it is moving with a velocity of $(0.5, 0.7, 0.1)$ ms^{-1}. This allows us to

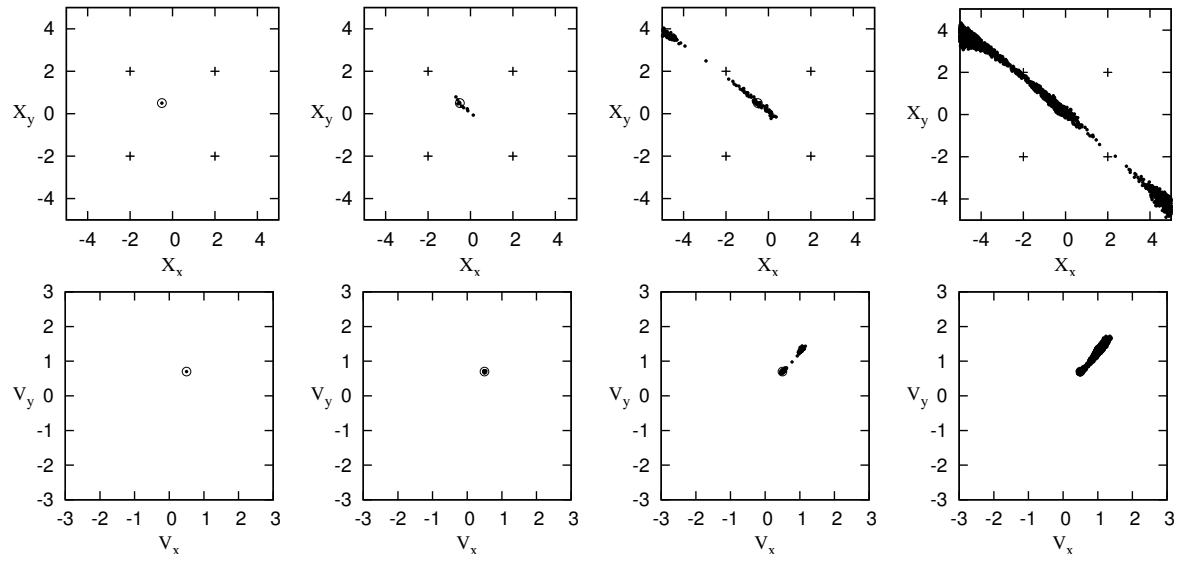

Figure 3: Position and velocity samples with δ less than (a) 2 cm, (b) 4 cm, (c) 6 cm, and (d) 8 cm using the cubic beacon configuration

compute a reference ΔP_i^{ref} that corresponds to each beacon i. For example, ΔP_0^{ref} for the beacon at $(-2,-2,-2)$ is 1.1 ms and ΔP_7^{ref} for the beacon at $(2,2,2)$ is -1.2 ms. The collection $(\Delta P_0^{\text{ref}}, \Delta P_1^{\text{ref}}, ...\Delta P_7^{\text{ref}})$ forms a vector of reference measurements, denoted $\mathbf{\Delta P}^{\text{ref}}$.

For each experiment, we sample a large number of random pairs $(\mathbf{X},\mathbf{V})_j$ in the six dimensional space spanning location and velocity. The pairs are uniformly distributed in location (with bounds $[-5..5]$) and velocity (with bounds $[-3..3]$). For each pair $(\mathbf{X},\mathbf{V})_j$ we calculate the ΔP_i^j for each beacon i and compare them with the reference measurements by taking the magnitude of their difference:

$$\delta^j = |\mathbf{\Delta P}^{\text{ref}} - \mathbf{\Delta P^j}|$$

Ideally, position-velocity pairs that are further away from the reference pair in the 6D space should have higher δ values, while pairs that are nearby have lower δ values. However, it turns out that there are many pockets of local minima that are distributed over the 6D space. They are widely spread in the location space, while the velocity space is slightly more constrained.

The importance of δ lies in the fact that our location hardware has limited accuracy and precision. As such, measurements for ΔP contain errors. In an ideal situation, small measurement errors will cause small perturbations in position and velocity estimates. However, if there are pockets of solutions that have smaller δ values than our measurement error, then our estimators could easily move into those pockets. In the two succeeding sections we present results from our experiments on the cubic and real-world configurations of the beacons. We use a varying threshold for δ to observe how deep the local minima are relative to the measurement errors.

3.2 Results on cubic layout of beacons

Figure 3 depicts the results of the cubic beacon experiment. Each plot in the figure shows the X-Y placement of position (in the top row) and velocity (in the bottom row). The Z-axis has been left out for simplicity. The dot surrounded by a circle is the reference point, while all other dots are the randomly generated position-velocity pairs. For the plots detailing position, the stars depict the position of the beacons in the X-Y plane. For this experiment, only four beacons can be seen since the other four have the same X and Y coordinates with different Z coordinates.

The figure has been separated into columns (a) through (d), which correspond to different thresholds on δ. Column (a) includes all pairs that have δ values less than 2 cm, and Column (d) includes all pairs that have δ values less than 8 cm. We use distance values (cm) to describe δ (by dividing by the speed of sound) to make it easier to compare with our observed sensor noise (which we commonly describe in terms of distance). Nevertheless, it can be seen that, as the threshold for δ increases, there is more of the 6D space that is covered by local minima.

Viewing the location space, one can see that the local minima seem to occupy, roughly, a plane (which extends in the Z direction – not visible in the plots). It appears that the plane is perpendicular to the velocity of the receiver. In the velocity space, the local minima appear as an ellipsoid of velocities that have the same direction as the reference velocity, but with higher magnitude (speed).

From the simulation we can also get a rough idea on the fraction of the 6D space that contains local minima that are lower than the measurement error. The number of pairs with a δ of 15 cm or less (which is a very conservative estimate for the maximum sensor noise we observe; the standard deviation of the sensor noise measured in our lab is around 1 to 2 cm), is less than 0.0005%. For a 64 m^3 room, that corresponds to a space approximately 6.7 cm cubed.

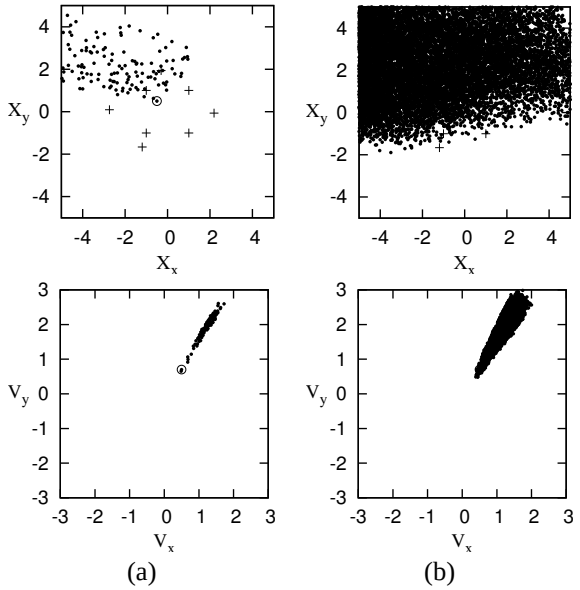

Figure 4: Position and velocity samples with δ less than (a) 4 cm, (b) 8 cm using the lab beacon configuration

3.3 Results on laboratory layout of beacons

The results from the experiment using the beacon configuration employed in our lab is given in Figure 4; we provide two thresholds for δ, 4 and 8 cm. In this configuration, there is only a small variation in the Z-coordinates for the location of the beacons. As a result, we observe a much more unpleasant distribution of local minima. Specifically, the local minima are widely distributed over the X and Y axes. We believe that the higher dilution of precision of the beacons in the Z-axis, means that many ghost solutions are created under the floor, and are widely spread over X and Y.

As with the previous configuration, the ghost solutions in velocity are approximately constrained to an ellipsoid extending in the same direction of the reference velocity. However, compared with the previous experiment, the error in position seems to be determined by the relative positions of the reference point and the beacons. The fraction of solutions that fall under the 15 cm threshold is less than 0.0186%, or 23 cm cubed in terms of a 64 m^3 space.

3.4 Results along a line in space

In order to further explore the presence of the local minima in the cubic beacon experiment, we have made three one-dimensional cuts through the 6D space and plotted the error against the distance traversed along these cuts. This is shown in Figure 5. Each cut follows the same path through the 3D location space, intersecting the reference point. The lines follow the diagonal band of points displayed in Figure 3(d). The horizontal line in Figure 5 is the corresponding 6 cm δ threshold. Each of the three lines depicting the path through the 6D space have a fixed but different velocity. We observe that one of the lines reaches zero-error. This is the line with velocity set to the reference velocity and, as expected, it hits zero when it goes through the reference point. We notice that the other lines have local minima at other points on the diagonal path traversed. In the next section, we detail our experience of using the Doppler model with two types of estimator, a Kalman Filter and a Particle Filter.

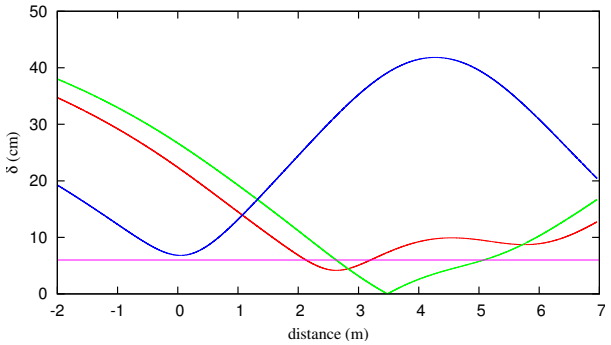

Figure 5: δ measured along three lines in the 6D space

4. APPROXIMATION METHODS

We have explored two methods to retrieve the location and velocity using Equation 1: one uses a Kalman Filter and the other uses a Particle Filter. Details of these two filters are in two companion papers [4, 5].

4.1 Kalman Filter

For this algorithm, the receiver uses a Kalman Filter to model position and velocity in three dimensions. Each measurement is incorporated using Equation 1 as it arrives; this method is known as "single constraint at a time" filtering [9]. In order to deal with collisions and ambiguities, the algorithm uses a multi-hypothesis [1] approach. For example, when a chirp arrives that could possibly come from multiple sources, the algorithm creates *m* hypotheses, where *m* is the number of possible sources (including noise or reflections). Each hypothesis contains a Kalman Filter that uses a unique guess as to the source of the ambiguous measurement. The idea behind this approach is that, as more chirps arrive, incorrect hypotheses can be eliminated.

4.2 Particle filter

The particle filter models the 6D space for location and velocity, where each particle models a point in this space. The particle state is progressed by changing the particle location in line with the particle velocity, and by adding a random variation to both. Variation in the X and Y components of the velocity has a standard deviation of 2 ms^{-2}, and the variation on the Z component is assumed to be 0.2 ms^{-2}; we assume that our receivers are mounted on humans and, hence, sudden variations in X and Y are possible, but only minor variations in Z are expected.

4.3 Results

To date we have managed to get the Kalman Filter approach to work with simulated data and are in the process of configuring it to work with real data. Initial observations have suggested that it may be difficult to determine which hypotheses are stronger than others and we have yet to find a set of metrics to provide this evidence consistently. It may also be the case that ambiguous measurements do not actually make enough of a difference to warrant the branching of separate

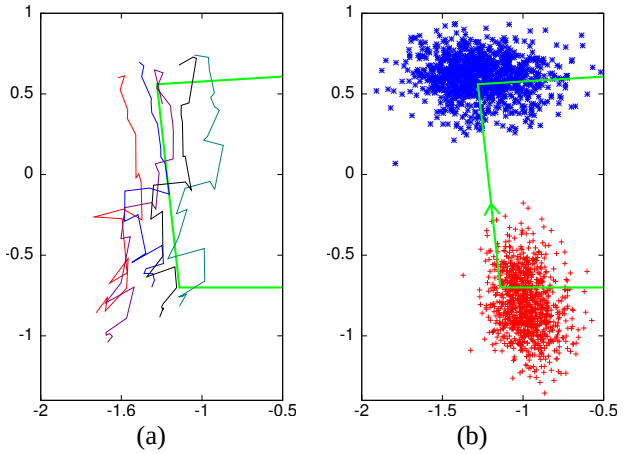

Figure 6: Position estimation with particle filter: (a) the tracks, made between two corners; (b) the positions in two of the corners of the test track.

hypotheses. For example, an ambiguous measurement could be ignored, used as a chirp from the most likely beacon, or interpreted as coming from both. If this is true, it may explain the lack of evidence to support "stronger" hypotheses.

The Particle Filter works on real data, and results are detailed in Figure 6. Figure 6(a) shows fragments of paths that the particle filter recovered on five separate runs with the same recorded measurements. The plot highlights the variation in position that the filter produces, independent of the input data. We observe that the tracks have slightly odd shapes (due to the particles being updated with a single measurement at a time), but we also notice that the direction of the tracks mostly agree.

Figure 6(b) shows the positions that the particle filter estimates on 1000 separate runs over the set of recorded measurements. The two clouds of points are the location for two separate time snapshots in the test track, three seconds apart. Each point in the cloud is the location produced by one run of the particle filter. The 50% CEP (circular error probability) is around 25cm, but the error seems to be less pronounced in the direction of travel. For example, the receiver was moved clockwise along the green line. One can see that on arriving at the bottom left-hand corner, the error in the X-direction is about half the error in the Y-direction. Similarly, arriving at the top left-hand corner produces errors in the Y-direction that are about half the error in the X-direction. These results support the observations in Section 3, where position error is largest perpendicular to the direction of travel.

5. DISCUSSION

The presence of local minima within the bounds of our measurement error suggests that the Doppler equation is relatively insensitive to the position of the receiver, especially for a beacon arrangement with a high dilution of precision. Despite this, we have shown that it is possible to recover position to a coarse degree of accuracy using a particle filter. We speculate that, because of the non-linearity of the measurement function and the associated distribution of local minima, a Kalman filter will perform less well than the particle filter.

Unlike ranging solutions that are common to ultrasonic positioning, errors in measurements used with the Doppler equation will have pronounced effects on the positioning accuracy. Ranging systems with an optimal layout of beacons, for example a cubic layout, tend to have localised position errors in response to measurement errors. From our discussion in Section 3, we observe that this is not the case for the Doppler solution.

However, it should be feasible to introduce a ranging component to our model. For example, we can use the Doppler equations to provide an estimated velocity and location and subsequently build a model of the beacon transmission clocks (a form pseudo-ranging). Using that model, we may be able to recover our location with a higher accuracy.

REFERENCES

[1] Yaakov Bar-Shalom and Thomas E. Fortmann. *Tracking and data association*. Academic Press Professional, Inc., San Diego, CA, USA, January 1988.

[2] A.J. Davison. Real-time simultaneous localisation and mapping with a single camera. In *Proc. International Conference on Computer Vision, Nice*, October 2003.

[3] C. Durieu and H. Clergeot. Navigation of a mobile robot with ultrasonic beacons. In Louis O. Hertzberger and Frans C. A. Groen, editors, *IAS*, pages 203–208. North-Holland, 1986.

[4] Michael McCarthy and Henk L. Muller. Positioning with independent ultrasonic beacons. Technical Report CSTR-05-005, Department of Computer Science, University of Bristol, September 2005.

[5] Henk Muller, Michael McCarthy, and Cliff Randell. Particle Filters for Position Sensing with Asynchronous Ultrasonic Beacons, 2006.

[6] Mark Pupilli and Andrew Calway. Real-time camera tracking using a particle filter. In *Proceedings of the British Machine Vision Conference*, pages 519–528. BMVA Press, September 2005.

[7] Adam Smith, Hari Balakrishnan, Michel Goraczko, and Nissanka Priyantha. Tracking Moving Devices with the Cricket Location System. In *Proceedings of the 2nd international conference on Mobile systems, applications, and services*, pages 190–202. ACM Press, June 2004.

[8] A. Ward, A. Jones, and A. Hopper. A New Location Technique for the Active Office. In *IEEE Personnel Communications, volume 4 no.5*, pages 42–47, October 1997.

[9] Greg Welch and Gary Bishop. SCAAT: Incremental Tracking with Incomplete Information. In *SIGGRAPH 97 Conference Proceedings, Annual Conference Series*, August 1997.