

RF-Based Initialisation for Inertial Pedestrian Tracking

Oliver Woodman and Robert Harle

University of Cambridge Computer Laboratory
{ojw28,rkh32}@cam.ac.uk
<http://www.cl.cam.ac.uk>

Abstract. Location information is an important source of context for ubiquitous computing systems. We have previously developed a wearable location system that combines a foot-mounted inertial unit, a detailed building model and a particle filter to locate and track humans in indoor environments. In this paper we present an algorithm in which a map of radio beacon signal strengths is used to solve two of the major problems with the original system: scalability to large environments and uncertainty due to environmental symmetry.

We show that the algorithm allows the deployment of the system in arbitrarily large buildings, and that uncertainty due to environmental symmetry is reduced. This reduction allows a user to be located after taking an average of 38 steps in a 8725 m² three-storey building, compared with 76 steps in the original system. Finally, we show that radio maps such as those required by the algorithm can be generated quickly and automatically using the wearable location system itself. We demonstrate this by building a radio map for the 8725 m² building in under two and a half hours.

Keywords: Radio, localisation, inertial tracking, particle filters.

1 Introduction

Some of the first context-aware computing systems made use of location information as the primary source of context. Today, GPS provides localisation outdoors, but precise indoor tracking of people remains an open research problem. We have seen indoor location systems based on infra-red [1], ultrasound [2], WiFi signal strength [3], ultra-wideband¹ (UWB), vision [4], and many others [5]. Nearly all location systems based on these technologies require the physical installation of fixed infrastructure in the environment. Furthermore, there is often a correlation between the amount of infrastructure and the positioning accuracy achieved, as shown in Figure 1. The amount of infrastructure is often prohibitively expensive to deploy and maintain.

¹ <http://www.ubisense.net>

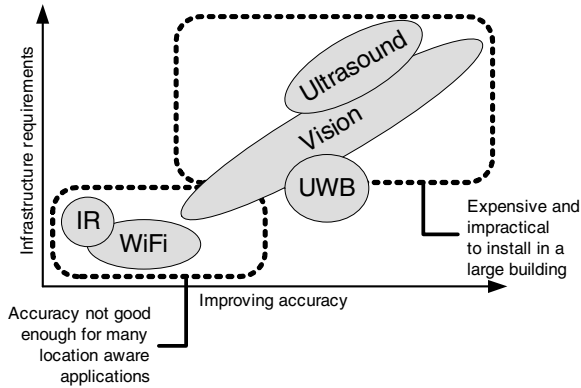


Fig. 1. The correlation between accuracy and infrastructure requirements in existing indoor location systems

In a related paper [6] we developed a wearable location system, in which a foot-mounted inertial measurement unit (IMU), a detailed building model and a particle filter were combined to locate and track a user with sub-metre accuracy. This system suffered from two major problems. Firstly, the computational power required by the system was relatively high and scaled as $\mathcal{O}(\text{Alog}_2(A))$, where A was the floor area of the building in which it was deployed. This made it impractical to deploy the system in large buildings. Secondly, ambiguity caused by symmetry in the environment could delay or even prevent the system from determining the user's true location.

It has been shown that radio beacons can be used to seed our location system with an approximate initial position for the user [6], helping to alleviate the problems described above. In this paper we present an improved algorithm in which a map of radio beacon signal strengths (a radio map) is used throughout the localisation process. We show that when using this algorithm, the user is localised after taking an average of 38 steps in our building, compared with 55 when radio beacons are only used to seed the initial position, and 76 when they are not used at all. We also show that the algorithm makes it feasible to deploy the system in arbitrarily large buildings.

In addition to their use in this paper, radio maps are also used by many existing RF-based location systems. Until now the construction of such maps has been done manually, which is very time consuming. The second contribution of this paper is to show how our wearable location system can be used to construct such maps quickly and automatically.

The structure of this paper is as follows. Section 2 introduces our wearable location system. Existing work in the field of RF-based location systems is outlined in Section 3. Section 4 describes the automatic construction of radio maps. Section 5 presents an algorithm for using such maps during the localisation process. Finally, we evaluate the effectiveness of the algorithm in Section 6.




Fig. 2. Wearable location system overview

2 Wearable Location System

An overview of our existing wearable location system is shown in Figure 2. The user wears a small and light inertial measurement unit (IMU), which is mounted on the foot and connected to a hip-mounted ultra-mobile PC (UMPC). The IMU measures acceleration and angular velocity at a frequency of 100 Hz. These measurements are processed by a pedestrian-dead-reckoning (PDR) filter, which generates a noisy step event of the form $(l, \delta z, \delta \theta)$ for each step taken by the pedestrian, where l is the computed step length, δz is the change in height of the foot and $\delta \theta$ is the change in heading relative to the previous step. The sequence of step events generated by the PDR filter describes the approximate path followed by the pedestrian, relative to an unknown starting position and orientation.

The localisation particle filter combines the step events with knowledge of constraints that exist naturally in indoor environments (specifically walls and floors) to determine the user’s absolute position. This process is known as localisation, and is described in Section 2.1. The user’s position is then tracked as described in Section 2.2. The building constraints are defined by a 2.5D map, which consists of a collection of rooms. Each room is defined by one or more planar floor polygons, which correspond to surfaces within the room on which a user’s foot may be grounded. Each edge of a floor polygon is either an impassable wall or a connection to the edge of another polygon. It is possible to represent even complex rooms using this format, such as the lecture theatre shown in Figure 3.

2.1 Localisation

Initially it is assumed that the user can be located anywhere in the building. Many particles (each of which represents a possible location and orientation of the user) are generated uniformly over all floor polygons in the map. Every particle is assigned an identical weight, which is the probability that the particle corresponds to the user’s actual position. Each time a new step event is received, a new set of particles is generated to represent the user’s updated position. The new set is generated by randomly sampling old particles in proportion to their weights (this is known as re-sampling). The position and orientation of each particle in the new set is then updated according to the received step event,




Fig. 3. (a) A lecture theatre. (b) Its 2.5D representation. The edges separating adjacent rows of seating are treated as walls (solid blue lines), whereas the edges between stairs in the aisle are connections (dashed green lines). Grid size = 1 m^2 .

perturbed by some noise to model uncertainty. If the perturbed step causes a particle to pass through a wall then that particle is assigned a weight of zero, which prevents it from being re-sampled during the next update. This is equivalent to removing the particle. Hence with each step, the user's possible location is narrowed down as shown in Figure 4. This continues until only a single cluster of particles around the user's true position remains. There are two major problems faced during the process of localisation, which are outlined below.

Scalability. The localisation particle filter requires $\mathcal{O}(n \log_2(n))$ time to update the set of particles when a new step event is generated, where n is the number of particles [6]. The number of particles is varied at each step and is proportional to the floor area within which the user may be located. Since the user may initially be located anywhere in the building, it is clear that for a large enough building it will not be possible to perform localisation in real time. Our lab is an example of such a building, with a floor area of 8725 m^2 . We found that around 4,530,000 particles were required to achieve reliable localisation. Our current implementation (written in Java and run on a 2.6 GHz Linux machine) is only able to update 800,000 particles during the average step duration.

Environmental symmetry. Symmetry of the environment can delay or prevent convergence to a single cluster of particles. If the layout of a building exhibits translational or rotational symmetry then the relative path described by the step events is often consistent with multiple different locations. In such cases




Fig. 4. An example of localisation in a three-storey building. (a) The route taken by the pedestrian (dashed red line) and a manually-aligned overlay of the steps generated by the inertial navigation component (solid black line); (b) The initial distribution of particles; (c-f) The particle distribution at four points during localisation. Particles which have passed through walls in the previous step are coloured black. An arrow indicates the actual position of the pedestrian in each figure. Grid size = 10 m². Diagrams are exploded 10x in the z-axis.

model and the known locations of radio beacons. In the latter, a radio map of the environment is constructed during an offline phase, in which beacon strengths are recorded at positions throughout the environment. The radio map is then used to locate the user during an online phase via pattern matching.

Radio propagation in an indoor environment is dependent on many factors, such as the wall and floor materials, which are difficult to accurately model. Complex propagation models have been devised which attempt to take such factors into account [11]; however model-based RF-location systems have been unable to match the accuracies achieved by map-based systems [3]. Radio propagation is also dependent on dynamic factors such as whether doors are open or closed, and whether other people are present in the environment. Hence it cannot be assumed that the distribution of signal strengths at a particular location is constant.

For human location systems (in which beacon strengths are measured by a device attached to the user), the user's body has a significant impact on the measured signal strengths. Obstruction of the line-of-sight to a beacon by the body causes signal attenuation, which can result in a reduction of up to 9 dBm in the received signal strength for WLAN access points [12]. The overall effect of the user's presence is that mean signal strengths are reduced and variances are increased. Hence if a radio map is to be used for human location, then the map itself should be constructed when a user is present.

A number of different algorithms exist for computing the user's position from a beacon scan made by the user at his current position and a pre-constructed radio map. In RADAR [3] the Euclidian distances between the measured signal strengths and those in each entry of the radio map are calculated. Taking the location for which the smallest distance was calculated to be the position of the user results in an error of 2.94 m 50% of the time, and 9 m 95% of the time. Wang et al. [13] used a similar algorithm to obtain a mean error of 6.44 m, which was improved to 4.30 m by using relative movement information obtained with an accelerometer.

The systems described above are deterministic, in that they compute a best guess of the user's position. An alternative is to compute a probability distribution $p(\mathbf{s}|\mathbf{z})$, which describes the probability that the user is located at position \mathbf{s} given the scan \mathbf{z} . The best guess location is then the position \mathbf{s} which maximises $p(\mathbf{s}|\mathbf{z})$. The Horus system [8] used a probabilistic approach, with a reported accuracy of 1.4 m 95% of the time. It is unclear how much control was exerted on the environment within which these results were obtained (e.g. whether other humans were present in the environment), and how vulnerable the system was to changes in the environment such as the opening and closing of doors.

Until now most radio maps used in RF-based location systems have been constructed manually. To do this a human stands at different locations in the building, manually marking his position on a map and then performing a beacon scan at each location. This method is very time consuming and is also prone to error, since the user must estimate their true position each time a scan is performed. Furthermore, the map would have to be manually updated if new

radio beacons were added, or if existing beacons were moved to new locations. As a result many deployments have been small [8], or restricted to the corridors of larger buildings [3]. Castro et al. [10] describe map-building as “tedious”. In one of the few large-scale deployments of a RF-based location system, 28 man-hours were required to construct a radio map covering a 12,000 m² building [14]. Robots have been used to acquire radio maps for use in robot location systems [15]; however such a map would not be well suited for use in a human location system since it would not take the presence of the user’s body into account.

4 Constructing Radio Maps

Radio maps can be constructed quickly and easily by a user who is being tracked using our wearable location system. In this paper we focus on mapping WiFi access points, however the same approach could be used to map other radio beacons. The WiFi hardware embedded within the hip-mounted UMPC is used to scan repeatedly for visible access points and their received signal strength indications (RSSIs). When the PDR-filter signals that a new step has been taken, the localisation filter updates the set of particles and clusters them to obtain a new position and orientation. The most recent WiFi scan to have been completed since the last position update (if one exists) is then incorporated into the radio map at this location. Note that radio map construction does not start until after the user has been localised. If environmental symmetry prevents localisation, then the user may manually select the cluster of particles which corresponds to his true position in order to complete the localisation process.

The radio map data structure (and the amount of data stored within it) can be changed according to the desired application. In this paper we construct a radio map for use during the localisation step of our location system. To do this we divide the 2.5D map into cells. A cell is a region in space within which radio measurements are grouped, with the cell size defining the granularity of the map. Each cell contains a histogram of measured RSSI values for each access point that has, at some point, been visible from the cell. A new measurement is added to a cell by simply updating the histogram for each access point that was sighted in the scan. If an access point is sighted for the first time in a cell then a new histogram is created. We considered two factors when deciding how to divide a 2.5D map into cells:

1. The radio map is most useful when the variability of radio measurements within cells is minimised, and variability between cells is maximised. Since walls attenuate radio signals more than air, we should not allow cells to span multiple rooms.
2. To get a complete radio map it is necessary to take measurements in every cell. Some rooms such as the lecture theatre shown in Figure 3 are made up of many small floor polygons. Therefore it should be possible for cells to span multiple floor polygons within a room.




Fig. 6. The cells generated for the lecture theatre shown in Figure 3 with a cell granularity of 6 m. Each cell is shaded with a different colour.




Fig. 7. Tracks generated by a user over 2 hours and 28 minutes, during which a radio-map was constructed.




Fig. 8. The radio map generated for a single access point. Cell colour indicates the mean RSSI value. Cell granularity = 5 m.

Table 1. The map coverage and mean scans per cell at different cell granularities

Cell granularity (m)	Coverage (%)	Mean scans per cell
1	59	1.3
2	85	4.3
3	96	9.14
4	99	12.9
5	99	18.9

Hence we divide each room into cells as follows:

1. Compute the 3-dimensional bounding box of each room.
2. Divide the bounding box into the minimum number of identical cuboids such that no edge is longer than the specified cell granularity.
3. A cell is defined as the fragments of floor polygon which are contained within a single cuboid.

The Weiler-Atherton clipping algorithm [16] is used to compute the fragments of each floor polygon that are contained in a cell, as shown in Figure 6.

We used our location system to construct a map of WiFi access points in a large (8725 m^2) three storey building. A user was tracked for 2 hours and 28 minutes (split into three sessions), travelling a total distance of 8.7 km as shown in Figure 7. In total 33 access points were observed, one of which is shown in

Figure 8. We define coverage as the percentage of the total floor area covered by cells in which at least one measurement was recorded (excluding restricted areas to which we could not gain access). The coverage at different cell granularities is shown in Table 1. The map with a cell granularity of 5 m is used in the next section to improve the localisation performance of the system.

5 RF-Assisted Localisation

There are two stages at which WiFi measurements can be incorporated into the localisation particle filter. The first is when the initial distribution of particles, known as the prior, is generated. The second is during the localisation process itself, in which measurements can be used to weight particles according to the probability of observing the measurement given a particle's state (i.e. the position and orientation of the particle). Formally, the weight of a particle at time t is calculated as

$$w_t = w_{t-1} \cdot p(\mathbf{z}_t | \mathbf{s}_t) \quad (1)$$

where w_{t-1} is the previous weight of the particle and $p(\mathbf{z}_t | \mathbf{s}_t)$ is the probability of observing the measurement \mathbf{z}_t given the state of the particle \mathbf{s}_t . A particle which is consistent with the measurement will be assigned a high weight, making it more likely to be re-sampled during the next update relative to a particle that is not consistent with the measurement. The result is that over multiple updates more and more of the particles that closely match the observed measurements are generated, and particles which do not match the measurements die out.

5.1 Prior Generation

In previous work, a simple uniform prior was generated based on visible access points in order to solve the scalability problem [6]. Here we present an improved algorithm which takes into account the observed RSSI values of each visible access point. During initialisation the WiFi hardware is used to scan for visible access points and obtain corresponding RSSI measurements. Let the vector of observed RSSI values be $\mathbf{z}_0 = (\text{ap}_1, \text{ap}_2, \dots, \text{ap}_n)$, where ap_n is the RSSI observed for access point n .

To constrain the initial location of the user we first compute the Euclidian distance between \mathbf{z}_0 and the vector \mathbf{c}_i of mean RSSI values for each cell (indexed by i) in the radio map. A prior region is then generated, consisting of all cells for which the distance is under a threshold $\epsilon = 9.5$ dBm, which was determined experimentally to ensure a low probability of failure as described in Section 5.3. Formally the prior region is given by

$$R_{\text{prior}} = \bigcup_i \text{Cell}_i \mid \text{ndist}(\mathbf{z}_0, \mathbf{c}_i) < \epsilon \quad (2)$$

where Cell_i is the set of polygon fragments in the i^{th} cell and $\text{ndist}(\mathbf{z}_0, \mathbf{c}_i)$ is the normalised Euclidian distance between \mathbf{z}_0 and \mathbf{c}_i as described below. The initial distribution of particles is generated uniformly across the prior region.

Since we are deploying our tracking system in a large building, careful consideration must be paid to computing the distance when different access points are visible at different locations in the building. This is a problem rarely addressed in existing literature, since most experimental deployments of RF-based location systems are small and implicitly assume that all access points are visible at all locations in the building. Consider the following scan and cell vectors, where a question mark indicates that an access point was not observed:

$$\mathbf{z}_0 = (-82, -73, ?) \quad (3)$$

$$\mathbf{c}_1 = (-84, -75, ?) \quad (4)$$

$$\mathbf{c}_2 = (-84, -75, -75) \quad (5)$$

As a first attempt we might choose to compute the distance between the observation and a cell using only the access points for which measurements are present in both vectors. This approach computes equal distances to both cells, yet intuitively the distance to cell 1 should be smaller, since an access point that is strongly visible in cell 2 was not observed by the user. Our solution is to compute the Euclidian distance over all access points, filling in missing readings in both the cell and the measurement vectors with the weakest RSSI value that can be measured by the WiFi hardware, in this case -96. Each distance measurement is then normalised by dividing by the number of access points that are visible from the cell according to the radio map. This makes it possible to compare distance measurements between areas where different numbers of access points are visible. Note that this normalisation penalises a missing value in the cell vector more than a missing value in the measurement vector:

$$\mathbf{z}_{0a} = (-82, -73, ?) \quad (6)$$

$$\mathbf{c}_{1a} = (-84, -75, -75) \quad (7)$$

$$\text{ndist}(\mathbf{z}_{0a}, \mathbf{c}_{1a}) = 7.1 \text{ dBm} \quad (8)$$

$$\mathbf{z}_{0b} = (-84, -75, -75) \quad (9)$$

$$\mathbf{c}_{1b} = (-82, -73, ?) \quad (10)$$

$$\text{ndist}(\mathbf{z}_{0b}, \mathbf{c}_{1b}) = 10.6 \text{ dBm} \quad (11)$$

This is a desirable property, since each cell in the radio map typically contains many measurements. If an access point were visible from a cell, then you would expect that it would have been sighted in at least one of them. In contrast the measurement vector \mathbf{z}_0 is from a single scan, and an access point that it is possible to sight from the cell may have been obscured by the user's body or some other object in the environment.

Figure 9 shows the result of using R_{prior} in the same example as shown in Figure 5. Note that the prior region covers only a small proportion of the total floor area. Since the number of particles is proportional to the floor area that they cover, the computational cost of localisation has been reduced. By using the prior the number of distinct clusters is reduced from eight to three, since only three of the clusters shown in Figure 5 originated from inside the prior region. Hence




Fig. 9. The remaining clusters for the same example as in Figure 5 when R_{prior} (indicated by the shaded region) is used. The cluster at the bottom right corresponds to the user’s true position.

the uncertainty due to symmetry in the environment has also been reduced, but the user’s location has not been uniquely identified. One problem which a prior region does not solve is distinct clusters that arise due to rotational symmetry and originate from the same location. For example the top left and bottom right clusters shown in Figure 9.

5.2 Localisation

To remove clusters which arise due to only rotational symmetry, we generate containment regions for each step during the localiasation process. For a step at time t we generate the region

$$R_t = \bigcup_i \text{Cell}_i \mid \text{ndist}(\mathbf{z}_t, \mathbf{c}_i) < \epsilon \tag{12}$$

where \mathbf{z}_t is the last WiFi scan to be completed prior to the completion of the step. We use this as a containment measurement with distribution

$$p(\mathbf{z}_t | \mathbf{s}_t) = \begin{cases} 1 & \text{if } \mathbf{s}_t \text{ is in } R_t \\ 0 & \text{otherwise} \end{cases} \tag{13}$$

which sets the weight of all particles outside the region to 0. These measurements remove clusters of particles which arise due to rotational symmetry in the environment. Figure 10 shows the result of using both the prior and containment measurements. Only a single cluster remains, which corresponds to the true location of the user.




Fig. 10. The remaining clusters for the same example as in Figure 5 when both R_{prior} and containment measurements are used. The shaded area indicates the most recently applied containment measurement.

5.3 Robustness

In the solution presented above we have opted to use containment regions, where the probability $p(\mathbf{z}_t|\mathbf{s}_t)$ is uniform across the entire region. An alternative (and perhaps more conventional) approach would be to compute a probability $p(\mathbf{z}_t|\mathbf{c}_i)$ for each cell and then apply a measurement with the distribution:

$$p(\mathbf{z}_t|\mathbf{s}_t) = p(\mathbf{z}_t|\mathbf{c}_i) \quad \text{where } \mathbf{s}_t \text{ is in Cell}_i \quad (14)$$

The distributions $p(\mathbf{z}_t|\mathbf{c}_i)$ could be computed directly from the histograms in the radio map and the measurement vector \mathbf{z}_t . We chose to avoid the use of probabilistic measurements because they assume that the distribution of RSSI measurements recorded for a given access point at a given location is constant. It has been noted elsewhere that this assumption is not valid [12], since the distribution can be affected by a change in the environment such as the opening or closing of a door. Such a change can introduce a systematic error into the resulting probability distributions, which can in turn draw the particle cloud away from the user's true position.

In our approach we avoid introducing systematic errors into the resulting probability distributions by making the distributions uniform across containment regions. The only problem that can arise is for the cluster corresponding to the user's true position to fall outside the containment region. To ensure that this is unlikely to occur, we processed just under one hour of trace data (distinct from that used to construct the radio map) gathered by the system. After each step during tracking, we calculated the normalised distance $\text{ndist}(\mathbf{z}_t, \mathbf{c}_i)$ for each cell overlapped by the single cluster corresponding to the user's position. The cumulative distribution of the calculated distances is shown in Figure 11. This




Fig. 11. The cumulative distribution of $\text{ndist}(z_t, c_i)$, for cells overlapped by the correct cluster of particles

distribution shows that for a threshold $\epsilon = 9.5$ dBm, the probability of the cluster falling outside the containment area is 0.00028 for a single step. We have not found any localisation to require more than 60 steps at this threshold. If 60 steps are required, the overall probability of failure is 1.7%. When a failure does occur, the localisation process is simply restarted.

6 Localisation Performance

In this section we quantify the localisation improvements provided by containment measurements generated using a radio map. To do this we randomly select 50 starting points from the first 30 minutes of a 35 minute continuous trace of the user moving throughout the building. For each starting position, we run the localisation process without using the radio map, using only R_{prior} , and finally using both R_{prior} and containment measurements.

The results obtained for a single starting position are shown in Figure 12. In this example 4,478,049 particles were used in the prior and 82 steps were required to localise the user without the radio map. When R_{prior} was used to constrain the initial location of the user, the prior consisted of 190,850 particles and 68 steps were taken before localisation. When both R_{prior} and containment measurements were used, the prior contained 228,001 particles and localisation required only 31 steps. The number of particles in the prior for the two cases where R_{prior} was used differ only due to the random process which is used to generate the priors [6]. Note that the number of particles is well under 800,000 when the radio map is used, which is the number of particles that can be updated in real time by our current implementation.




Fig. 12. The data gathered for a single localisation without using the radio map (red dashed), using R_{prior} only (blue dot-dashed), and using both R_{prior} and containment measurements (black solid). (a) The number of particles generated at each step. (b) The number of clusters at each step (clustering is only performed when the number of particles is below 10,000).

Table 2. Summary of localisation performance

	Radio map use	Mean	Std-dev	Max	Min
Particles in prior	None	4,449,156	44,195	4,482,151	4,371,901
	Prior only	535,493	446,751	1,435,550	211,850
	Throughout	539,644	462,869	1,493,701	228,001
Steps until localisation	None	76	14.1	96	63
	Prior only	55	24.7	96	24
	Throughout	38	12.2	53	23

Table 2 summarises the results. Using R_{prior} reduces the average number of particles to 535,493. There is however a large variation between the number of particles generated in different runs, with one run generating a prior of 1,435,550 particles. This is explained by the presence of large open plan areas in our building within which radio signals propagate much more freely than in areas where there are many walls. Hence larger priors are generated in such areas. In all cases the number of particles fell under the real-time limit of the current implementation within 5 steps, allowing the algorithm to catch up before the user was located. Hence using R_{prior} allows localisation in real time. The number of steps taken by the user before localisation is also reduced by using the constrained prior; however the worst case of 96 steps was not reduced, indicating that R_{prior} alone did not always reduce the problem of environmental symmetry.

Using containment measurements in addition to R_{prior} reduced the average number of steps before localisation to 38. The worst case was reduced to 53 steps from 96. Hence the use of containment measurements throughout the localisation process helps to solve the problems which arise due to environmental symmetry.

7 Conclusions and Future Work

In this paper we have developed an algorithm for using radio beacons during the localisation phase of our wearable location system. A radio map is used to generate containment regions within which the user is almost certainly located. These regions constrain the location of the user, allowing localisation to be performed in real time in arbitrarily large buildings. We have also shown that the algorithm reduces uncertainty caused by symmetry in the environment. In tests performed in a large building, a user was located after taking an average of 38 steps, compared with 55 steps when the beacons were only used to generate the prior, and 76 steps when the beacons were not used at all. Hence we have shown the benefit of using containment regions throughout the localisation process.

Our second contribution has been to demonstrate how the existing system can be used to construct radio maps quickly and automatically. We demonstrated our approach by constructing a radio map of a 8725 m² three-storey building in 2 hours and 28 minutes, with map coverage of 96% at a cell granularity of 3 m. This is of particular interest to developers of existing RF-based location systems, for which radio maps are usually constructed manually in a far more time consuming process.

In the future we plan to investigate how a radio map could be built, shared and updated by multiple users. This would allow the system to respond to changes in the environment including the installation, relocation and removal of radio beacons. We also plan to use our ability to quickly construct radio maps to investigate the temporal and spatial properties of radio beacon signal strengths throughout a large building.

Acknowledgements

The authors would like to thank Andy Hopper and David Cottingham for their insightful comments. This work has been funded by EPSRC.

References

1. Want, R., Hopper, A., Falcao, V., Gibbons, J.: The active badge location system. *ACM Trans. Inf. Syst.* 10, 91–102 (1992)
2. Addelee, M., Curwen, R., Hodges, S., Newman, J., Steggles, P., Ward, A., Hopper, A.: Implementing a sentient computing system. *Computer* 34, 50–56 (2001)
3. Bahl, P., Padmanabhan, V.: RADAR: an in-building RF-based user location and tracking system. In: *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings*, vol. 2, pp. 775–784. IEEE, Los Alamitos (2000)
4. López de Ipina, D., Mendonça, P.R.S., Hopper, A.: Trip: A low-cost vision-based location system for ubiquitous computing. *Personal Ubiquitous Comput* 6, 206–219 (2002)
5. Hightower, J., Borriello, G.: Location systems for ubiquitous computing. *Computer* 34, 57–66 (2001)

6. Woodman, O., Harle, R.: Pedestrian localisation for indoor environments. In: *UbiComp 2008: Proceedings of the 10th international conference on Ubiquitous computing*, pp. 114–123. ACM, New York (2008)
7. Ashbrook, D., Starner, T.: Learning significant locations and predicting user movement with GPS. In: *Proceedings of Sixth International Symposium on Wearable Computers (ISWC 2002)*, pp. 101–108 (2002)
8. Youssef, M., Agrawala, A.: The horus location determination system. *Wirel. Netw.* 14, 357–374 (2008)
9. Smailagic, A., Siewiorek, D.P., Anhalt, J., Kogan, D., Wang, Y.: Location sensing and privacy in a context-aware computing environment. *IEEE Wireless Communications* 9, 10–17 (2001)
10. Castro, P., Chiu, P., Kremenek, T., Muntz, R.R.: A probabilistic room location service for wireless networked environments. In: Abowd, G.D., Brumitt, B., Shafer, S. (eds.) *UbiComp 2001*. LNCS, vol. 2201, pp. 18–34. Springer, Heidelberg (2001)
11. Lott, M., Forkel, I.: A multi-wall-and-floor model for indoor radio propagation. In: *Vehicular Technology Conference, 2001. VTC 2001 Spring*. IEEE VTS 53rd, vol. 1, pp. 464–468 (2001)
12. Kaemarungsi, K.: Distribution of WLAN received signal strength indication for indoor location determination. In: *1st International Symposium on Wireless Pervasive Computing*, pp. 1–6 (2006)
13. Wang, H., Lenz, H., Szabo, A., Bamberger, J., Hanebeck, U.: WLAN-based pedestrian tracking using particle filters and low-cost MEMS sensors. In: *4th Workshop on Positioning, Navigation and Communication, 2007. WPNC 2007*, pp. 1–7 (2007)
14. Haeberlen, A., Flannery, E., Ladd, A.M., Rudys, A., Wallach, D.S., Kavraki, L.E.: Practical robust localization over large-scale 802.11 wireless networks. In: *MobiCom 2004: Proceedings of the 10th annual international conference on Mobile computing and networking*, pp. 70–84. ACM, New York (2004)
15. Ocana, M., Bergasa, L., Sotelo, M., Nuevo, J., Flores, R.: Indoor robot localization system using WiFi signal measure and minimizing calibration effort. In: *ISIE 2005. Proceedings of the IEEE International Symposium on Industrial Electronics*, vol. 4, pp. 1545–1550 (2005)
16. Weiler, K., Atherton, P.: Hidden surface removal using polygon area sorting. *SIGGRAPH Comput. Graph.* 11, 214–222 (1977)