802.11 Positioning in the Home

James Salter, Binghao Li
Daniel Woo, Andrew G. Dempster, Chris Rizos
{jsalter|danielw@cse.unsw.edu.au}, {binghao.li|a.dempster|c.rizos|@unsw.edu.au}
University of New South Wales, Sydney, NSW 2052 Australia

***Abstract*—802.11 positioning systems are established as a low-cost solution to positioning within context aware computing. Prior research has mostly focused on either indoor positioning within commercial environments, where access points are prevalent, or outdoors positioning over large areas, using discovered networks to augment GPS. In this paper we test 802.11 positioning in a medium sized domestic house with only a small number of access points, examining its suitability in a domestic context aware computing system.**

We use a signal strength fingerprint map approach, in which empirical signal strength data is gathered over the area prior to use. Estimation then involves comparing the online input to the map, and selecting the best position, and direction, based on similarity. We compare different strategies for gathering signal strength, and implementations of the Nearest Neighbor and Bayesian methods. Our results demonstrate that with good placement, only two access points are sufficient to estimate position with an error of less than 4 meters 90% of the time.

I. Introduction

802.11 Wireless LAN infrastructure has been used in several contexts for positioning, in addition to its usual data transport function. Prior research has established the performance of 802.11 positioning in office and laboratory environments [1]–[9]. It provides a cheap, low-cost solution in indoors environments where satellite positioning is unavailable.

In this paper, we examine the performance of 802.11 positioning in the context of a domestic house. In several respects, the domestic context differs from office and laboratory environments: cost, and hence the number of access points (APs), are both severely restricted, the areas involved are proportionally smaller, and the surveying phase must require a minimum of effort.

Many smart home applications rely on room occupancy sensing, such as automated air conditioning and lighting. Currently these services are mostly driven by infra-red sensors. We aimed to determine whether 802.11 positioning could provide sufficient accuracy to perform the same function. Occupants could be tracked by wireless signals transmitted by either portable consumer devices such as mobile phones, or low power 802.11 tags. Assuming that any smart home will already have a domestic wireless network, this could reduce costs by avoiding the need for sensor infrastructure.

We simulate a low number of APs by testing various arrangements of only one and two APs, achieving distance errors less than 5 and 4 meters respectively in 90% of cases, and we demonstrate that the calibration phase of our system takes only 10 minutes for a 100m^2 house.

II. Methods of positioning

802.11 positioning systems infer position from signal strength information incidental to the operation of 802.11 wireless LAN. Received radio signal strength varies according to the distance from the transmitter and how the signal path is reflected, absorbed, and scattered by objects in the environment. This propagation can be simulated by specialized software, but a more accurate image can be constructed by empirically surveying the area [1], [8]. During the surveying process, a *signal strength map* is constructed that characterizes the environment by associating signal strength with position. Each entry in a signal strength map is known as a *fingerprint*, named such because the signal strength it contains identifies the associated position as uniquely as possible. The exact format of these fingerprints varies according to the method used, but generally it involves some kind of representation of the signal strength to be expected at that position, such as the mean.

During positioning, the incoming input of signal strength is compared to the signal strength map, and the similarity between each fingerprint and the input is calculated. The output is a list of the fingerprints in the map ranked such that the first is the fingerprint calculated as closest to the current position, and the last is the furthest away. Several different methods have been used to perform the similarity test; in this paper we use the Nearest Neighbor and Bayesian methods.

A. Point and Sector sampling

The surveying process is complicated by the fact that the human body affects signal strength readings as it both attenuates and reflects the RF signal. Hence, at different directions for a particular position, different readings can be observed as the signal path is attenuated, or not, by the body of the person holding the device. The way the device is held also affects readings: a device in the user's pocket is reached by different signal path than a device held in their hand.

In surveying, two prior approaches have been used: *point surveying*, in which signal strength is gathered at a number of discrete locations over the environment facing north, south, east and west [1], [4], [8] and *sector surveying*, where data is associated with room sized sectors as the surveyor sweeps the area while walking around it [3], [5], [7]. Each method has its advantages. Point sampling generates a limited image of the space because it records data at a small set of positions within the relatively larger area, but because data is recorded directionally, it allows estimation of the mobile agent's direction as well as their heading [10]. However, point sampling can generate fingerprints distorted by pockets of multi-path interference peculiar to their point location, but unrepresentative of the immediate surrounds; this results in an inaccurate signal strength map. In contrast, sector sampling results in a more complete image of the space as the surveyor sweeps the device through the entire sector, but the data can contain no direction information. In the experimental section, we compare the accuracy available using each method.

B. Nearest Neighbor method

The Nearest Neighbor algorithm, proposed by [1], calculates the p-norm distance between signal strength vectors to determine similarity. The problem is modeled using a state space of fingerprint locations $L = \{l_1, \ldots, l_n\}$, and an observation space $O = \{o_1, \ldots, o_m\}$. Each state l_i corresponds to the location at which fingerprint i was recorded. Each observation is a signal strength vector $o = \{\lambda_1, \ldots, \lambda_n\}$, where λ_j is a signal strength reading from AP j.

Fingerprints are generated by standing at a location in the coverage area and recording signal strength data for some period of time. Using this data, the mean signal strength for each AP is calculated; each mean forms an element in a signal strength vector associated with that location. Hence, a fingerprint can be expressed as $f_i^{NN} = (l_i, \{\mu_{i,1}, \ldots, \mu_{i,n}\})$, where l_i is the location of the fingerprint, and $\mu_{i,j}$ the mean signal strength for APs b_j at location l_i. Here the location l can refer to either a Cartesian coordinate, with or without a direction (in the case of point sampling), or refer to a less well defined region such as a room or part of a hallway, as is common with sector sampling.

Once this map has been generated, online positioning is performed by calculating the distance between the currently observed signal strength o and each fingerprint. The fingerprints are then ranked according to this distance, with the top-ranked fingerprint containing the most likely position and the bottom-ranked fingerprint the least likely.

The p-norm distance is used to calculate the distance between the fingerprint and observation vectors.

$$D(f_i, o) = \left(\sum_j |\mu_{i,j} - \lambda_j|^p \right)^{1/p} \tag{1}$$

[1] used the Euclidean distance, $p = 2$ to calculate the distance between vectors, although any $p \geq 1$ can be used.

C. Bayesian methods

Bayesian algorithms [3], [9], [11] use Bayes' rule to calculate the mobile agent's location. The final output is a discrete probability distribution, $\overrightarrow{P}$, over L, representing the likely location of the agent. $\overrightarrow{P}(l_i)$ expresses the probability that the mobile agent is at the location of fingerprint i. In other words, the output is a one to one mapping between each fingerprint location and the probability of the agent being at that location.

$\overrightarrow{P}$ is calculated upon receiving each new observation o using Bayes' Rule as follows:

$$\overrightarrow{P}(l_i) = p(l_i|o) = \frac{p(o|l_i)p(l_i)}{p(o)} \tag{2}$$

$p(l_i|o)$ is a conditional probability representing the likelihood that the mobile agent is at location l_i, given observation o. This is known as the *posterior* probability. $p(o|l_i)$ is the *likelihood function*: the probability of making observation o at location l_i. $p(l_i)$ is the *prior probability* and represents the probability that the device can be at location l_i. $p(o)$ is a normalizer that can be calculated within this context as $\sum_{i=1}^{n} p(o|l_i)p(l_i)$. It can also be used as a measure of the *confidence* of the distribution.

The likelihood function can be calculated using the signal strength map. Prior to positioning, we record the training data similarly to the nearest neighbor method. However, instead of associating each location l_i with a vector of average signal strengths for each AP, we associate it with a vector of empirical probability distributions. Fingerprints can be expressed as $f_i^{Bayesian} = (l_i, \{\Lambda_{i,1}, \ldots, \Lambda_{i,n}\})$. $\Lambda_{i,j}$ is the empirical distribution at location l_i for signal strength from AP b_j. Each empirical distribution is generated by taking the recording of signal strength values for an AP, calculating the frequency of each signal strength value, and normalizing the frequencies. Hence, $\Lambda_{i,j}(\lambda) = p(\lambda_j|l_i)$: the probability of observing signal strength λ from AP b_j, at the location l_i.

The likelihood function is calculated by multiplying the conditional probability of observing each reading in the observation vector $o = \{\lambda_1, \ldots, \lambda_n\}$:

$$p(o|l_i) = \prod_j p(\lambda_j|l_i) = \prod_j \Lambda_{i,j}(\lambda_j) \tag{3}$$

This calculation assumes independence between the readings from different APs for a fixed location. This is not strictly true: proximate APs broadcasting on overlapping channels can interfere with each other, but in a properly configured network such interference is minimal. The alternative to this method is to maintain multi-dimensional distributions, but this greatly increases the complexity of the calculations and the size of the signal strength map, and is unlikely to deliver better results assuming the dependancy violation is insignificant.

The prior can be used to bias the result towards one location or another; this can be useful for considering historical information or the presence of impassable objects such as walls. We used a uniform prior distribution, which introduces no bias towards any particular location, for simplicity.

Within this framework, implementations can differ in their derivation of $\Lambda_{i,j}$. It is possible to use the empirical distributions directly in the Bayesian framework, but the probability distribution is only a rough estimation of the underlying probability based on a limited amount of data. Further grooming of the distributions can lead to better results. We have used two methods in our work: a histogram method similar to [11], and summarization of the signal strength as a normal distribution similarly to [7].

1) Histogram: The *histogram method* functions similarly to the method described above, but generates a histogram $H_{i,j}$ in place of the empirical distribution $\Lambda_{i,j}$. Rather than calculating the frequency of each signal strength value λ, the frequencies are grouped into q contiguous bins $\{h_1, \ldots, h_q\}$ of fixed and equal width w. Each bin covers a continuous, non-overlapping range of signal strength values. The histogram constitutes a discretized version of the original distribution $\Lambda_{i,j}$, and is used directly in its place when calculating the likelihood function.

2) Gaussian: In the Gaussian method, $\Lambda_{i,j}$ is summarized as a Gaussian probability distribution by storing only its mean $\mu_{i,j}$ and standard deviation $\sigma_{i,j}$. The probability of observing signal strength λ_j at location l_i is given by

$$p(\lambda_j | l_i) = \frac{1}{\sigma_{i,j}\sqrt{2\pi}} \exp\left(-\frac{(\lambda_j - \mu_{i,j})^2}{2\sigma_{i,j}^2}\right) \quad (4)$$

Although the data is not necessarily Gaussian, or even unimodal, fitting the data to a Gaussian distribution reduces the influence of outliers and smoothes any gaps in the distribution. The Gaussian method can also result in a more compact fingerprint map than the Histogram method: only the mean and standard deviation need be stored for each AP rather than an entire distribution. Fingerprints generated using this method can be expressed as $f_i^{Gaussian} = (l_i, \{(\mu_{i,1}, \sigma_{i,1}), \ldots, (\mu_{i,n}, \sigma_{i,n})\})$.

D. Output Averaging

Taking the output ranking of fingerprints and their associated distances, we can spatially average the locations of the k closest fingerprints. This technique is directly applicable to point sampling, where each location is a Cartesian coordinate; for sector sampling we use the geometric centre of the sector. Either the arithmetic or weighted mean can be used. For the Nearest Neighbor and Distribution Distance methods, the weights are the inverses of the distances, and for the Bayesian methods the probabilities can be used directly.

III. Experimental Setup

Our experimental data was recorded using a Sony VAIO U8G running Fedora Core 4 Linux, using the Kismet packet sniffer. The U8G is a palmtop computer ideal for this scenario; small enough to carry around with one hand but with the power and programmability of a laptop computer. Kismet was slightly modified so that we could extract per-packet signal strength from it. The Kismet approach generalizes to a range of hardware and the per-packet sampling allows a sampling rate of at least 10Hz, the default rate at which beacon packets are broadcast by most 802.11 APs. We used a simple Perl script to connect to the Kismet server and gather the data for a list of visually pre-selected fingerprint coordinates in a local frame.

The setup is shown in figure 1. It contained six 802.11b APs of various models, 19 fingerprint locations, and 39 test point locations. Two types of fingerprint maps were recorded: Point samples, facing north, south, east and west (relative to the house's orientation) for 10 seconds at each fingerprint location, and corresponding sector samples, for 20 seconds each.

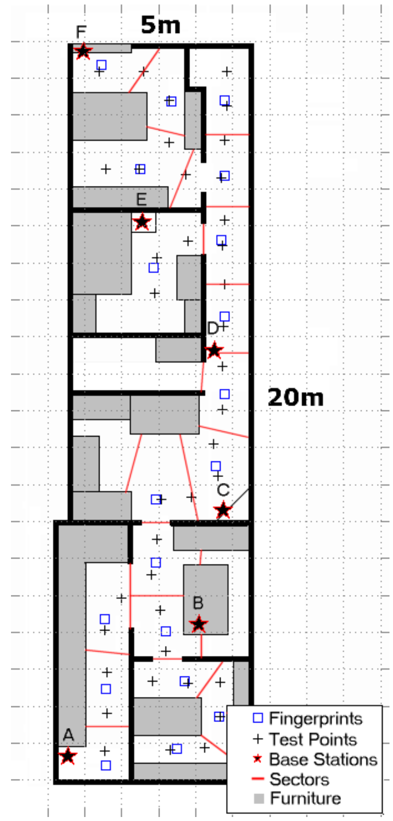

Fig. 1. House test bed, showing fingerprint, test point, and AP locations.

Each sector covered a 2m^2 area centered on a corresponding fingerprint location, such that coverage was contiguous. Each sector was recorded by continuously sweeping the device and changing direction, attempting to gather the most even coverage possible. Sector sampling took around 10 minutes, and point sampling took around 20 minutes. We recorded test points at a separate set of locations spaced one meter apart facing north, south, east and west, 5 seconds per direction, a day after recording the fingerprint data.

For the sake of our accuracy analysis, each sector was assigned a coordinate identical to the point fingerprint in it's centre. Note that it is not strictly necessary to have a raster map and a coordinate frame for the area in order to perform positioning. The fingerprints can be instead associated with labels, such as room names, and output can be in terms of these instead.

In this paper we focus on the performance available with one and two APs rather than the complete set of six, to simulate a realistic setup that might be found in a domestic house. The advantage of using six APs to gather data is that only one survey is required, but many lower density configurations can be simulated by removing the superfluous entries from the signal strength map and test data.

IV. Experimental Results

We tested the data using a wide range of parameters for the Nearest Neighbor and Histogram methods (the Gaussian method has no parameters). For the Nearest Neighbor we tested values of p between 1 and 10, finding the best values were 2, 3 and 4. For the Histogram method, we tested w between 1 and 40, finding the best results were achieved between 5 and 15. The results of the Nearest Neighbor and Histogram methods shown here are for the outright best performing respective parameters $p = 2$ and $w = 5$.

When using output averaging, the point and sector maps performed best with the non-weighted mean using $k = 5$ and $k = 2$ respectively. This is due to the differing spatial coverage of the two types. In the point map, four fingerprints represented a single point, one recording for each direction. In the sector map, there was only one fingerprint for each sector. For the point map, this meant that the first k fingerprints were likely to share locations; for the sector map the fingerprints represented unique locations. Hence, the best k represented a wider spread of locations for the sector map. The results shown were calculated using these k values.

Accuracy is described in terms of the *distance error* - the distance between the output of the algorithm and the location at which the test data was recorded. For each method, we calculated the 50th percentile (median) error and the 90th percentile. The 50th and 90th percentile distance errors describe the error which 50 percent and 90 percent of the results are less than, respectively. Analysis of the 90th errors, in addition to the 50th, provides some indication of the stability of the method in pathological cases.

Figures 2 and 3 show the median and 90th percentile distance errors when using each AP setup. The positions of APs A-F are shown in figure 1. The results are fairly similar for all methods, and for both types of surveying method. The Nearest Neighbor usually performs marginally better than the Bayesian methods. The sector surveying method usually results in slightly lower median distance errors, but has similar or slightly worse performance than point surveying in the 90th percentile.

For single AP setups, the results are extremely sensitive to placement. The best results were achieved for A and F - when the AP was at either the north or south ends of the house. The worst results were achieved using a single AP in the centre of the house. In contrast, for double AP setups the results seem to be relatively indifferent to placement. For example, A,F represents placement of APs at either end of the house, while A,B represents placement at the south end only, yet the former shows only marginally better accuracy.

The sensitivity towards placement here can be explained by considering the distribution of signal strength. Generally, APs radiate power omni-directionally, such that similar signal strengths will be found in concentric bands about each AP, shaped by the obstacles in the environment. In our current context, this means that when a single AP is placed in the vertical centre of the house, there will be roughly symmetrical signal strengths to the north and south of the AP. In other words, the one dimensional signal strength data is ambiguous, as the same signal strength is observed at positions north and south of the AP. In the worst case, fingerprints at the extreme north and south of the house are indistinguishable, leading to the large errors we can observe for APs C and D. Under the same principle, placing APs at the north or south ends of the house results in less ambiguous fingerprints within the area of the house, and better positioning performance. Placement is apparently less of a problem when there are two APs, probably because the two dimensional data already contains an adequate amount of unique combinations.

TABLE I
POSITIONING ACCURACY EXPRESSED IN TERMS OF ROOM ESTIMATION.

BS	Result	Gaussian	Histogram	NN
A	Correct	38.5	38.6	39.9
	Adjacent	45.4	44.4	46.0
F	Correct	46.0	40.5	45.7
	Adjacent	48.1	53.6	47.6
A,E	Correct	60.4	59.2	60.6
	Adjacent	35.1	33.5	34.1
B,E	Correct	55.9	51.9	55.3
	Adjacent	37.9	42.1	38.2
All	Correct	70.7	69.8	73.4
	Adjacent	26.7	29.0	25.4

Note that the area used here is longer than it is wide, and benefits from this effect more than a square area would. The same approach might still be useful for a square area; if APs are placed at the corners, or edges, the number of ambiguous fingerprints should be reduced.

Table I shows the accuracy expressed in terms of room estimation, for the two best performing single and double AP setups. We calculated the percentage of tests in which the positional estimate was in the same room as the test point, or in an adjacent room. The system was very efficient in selecting either the correct or adjacent room, but was not very effective at determining the outright correct room. Using all six APs only resulted in a 10% increase in correctness, which suggests diminishing returns with a greater number of APs.

V. Conclusions

We have demonstrated that relatively good positioning accuracy is achievable using only one or two APs, when appropriately placed. Our results were similar for the different positioning algorithms and surveying methods used.

The optimal placement for one AP was at the north or south edge of the house, as this minimizes the ambiguity of the resultant one dimensional signal strength data; this kind of placement might penalize data transfer speed or connectivity at the far side of the area. In a situation where this would be problematic, two APs would be more appropriate. Double AP setups were less sensitive to placement, but performed slightly better when the APs were at opposite ends of the house.

In this paper, we have not analyzed the impact of human traffic on positioning accuracy, which is likely to be significant. In busy households, it is likely that a single AP would be insufficient for positioning, as any unmapped obstacle to the single signal path would result in an erroneous position.

The system's room estimation could be refined further. For example, we could weight each room according to the weights associated with its constituent fingerprints, rather than just taking the containing room of the final estimate as we do here.

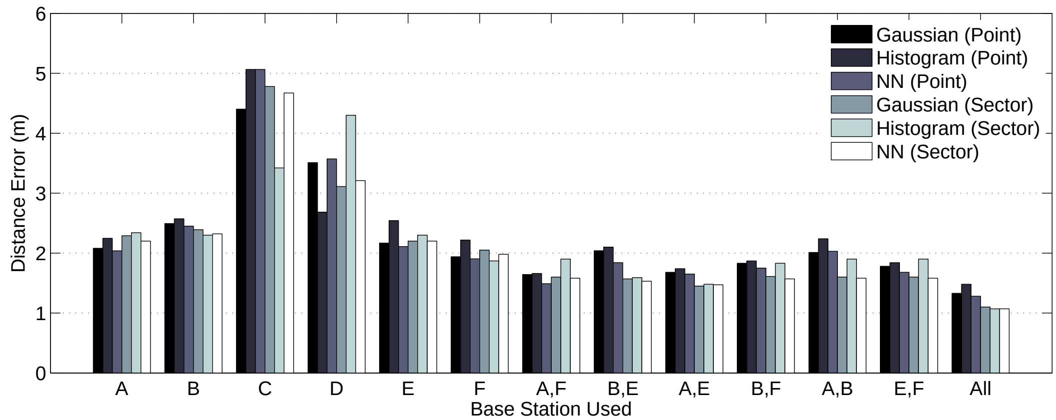

Fig. 2. Median distance errors obtained using the nearest neighbor, histogram, and Gaussian methods, using point and sector sampling.

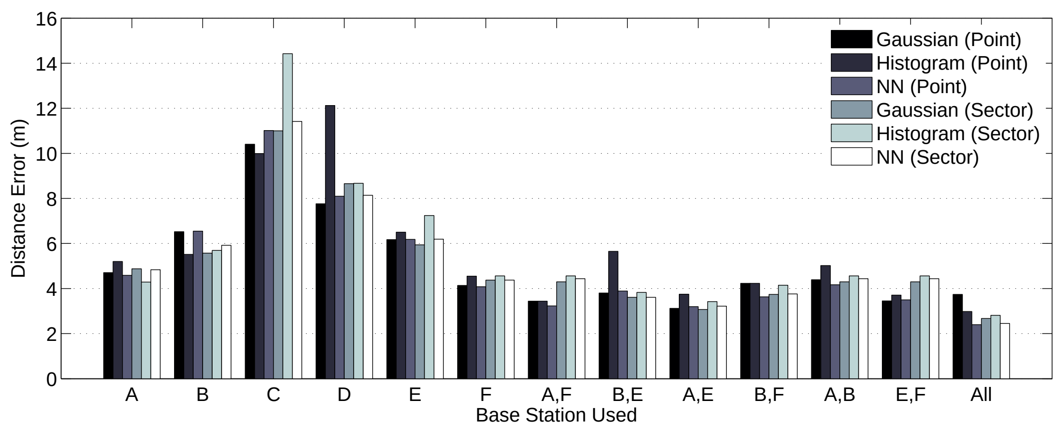

Fig. 3. 90th percentile errors obtained using the nearest neighbor, histogram, and Gaussian methods, using point and sector sampling.

Biasing the prior in the Bayesian methods could also improve accuracy by preventing output in the middle of a wall or a piece of furniture.

Given the low hardware requirements, the short surveying time and the accuracy achievable, we believe single or double 802.11 AP setups are useful as a low-cost positioning solution for smart home applications where the maximum accuracy required is to the actual or adjacent room. Increasing the number of APs can provide a more robust estimate in situations where the actual room is required.

Acknowledgment

This research took place in conjunction with the Audio Nomad project, lead by Dr Daniel Woo, Dr Nigel Helyer and Prof Chris Rizos, supported by and Australian Research Council Linkage Project (LP0348394) in conjunction with the Australia Council for the Arts as part of the Synapse Initiative.

References

[1] P. Bahl and V. N. Padmanabhan, "RADAR: An In-Building RF-based User Location and Tracking System," in *Proceedings of IEEE Conference on Computer Communications (INFOCOM)*, 2000.

[2] J. Small, A. Smailagic, and D. P. Siewiorek, "Determining user location for context aware computing through the use of a wireless LAN infrastructure," 2000. [Online]. Available: http://www-2.cs.cmu.edu/aura/docdir/small00.pdf

[3] P. Castro, P. Chiu, T. Kremenek, and R. Muntz, "A Probabilistic Room Location Service for Wireless Networked Environments," in *3rd International Conference on Ubiquitous Computing (UBICOMP)*, 2001.

[4] A. M. Ladd, G. Marceau, K. E. Bekris, L. E. Kavraki, A. Rudys, and D. S. Wallach, "Robotics-Based Location Sensing using Wireless Ethernet," Department of Computer Science, Rice University, Tech. Rep. TR02-393, 2002.

[5] M. A. Youssef, A. Agrawala, and A. U. Shankar, "WLAN Location Determination via Clustering and Probability Distributions," in *Proceedings of the first IEEE International Conference on Pervasive Computing and Communications*, 2003.

[6] D. Niculescu and B. Nath, "VOR Base Stations for Indoor 802.11 Positioning," in *Proceedings of Tenth Annual International Conference on Mobile Computing and Networking, (MOBICOM)*, 2004.

[7] A. Haeberlen, E. Flannery, A. M. Ladd, A. Rudys, D. S. Wallach, and L. E. Kavraki, "Practical robust localization over large-scale 802.11 wireless networks," in *Proceedings of the 10th annual international conference on Mobile computing and networking*, 2004.

[8] B. Li, A. Dempster, C. Rizos, and J. Barnes, "Hybrid Method for Localization Using WLAN," in *Proceedings of SSC 2005 Spatial Intelligence, Innovation and Praxis*, 2005.

[9] B. Li, J. Salter, A. G. Dempster, and C. Rizos, "Indoor Positioning Techniques Based on Wireless LAN," in *AusWireless '06, Sydney, Australia*, March 2006.

[10] B. Li, J. Kam, J. Lui, and A. Dempster, "Use of directional information in wireless lan based indoor positioning," in *Symp. on GPS/GNSS (IGNSS2007), Sydney, Australia,*, Dec. 2007.

[11] T. Roos, P. Myllymaki, H. Tirri, P. Misikangas, and J. Sievanen, "A Probabilistic Approach to WLAN User Location Estimation," *International Journal of Wireless Information Networks*, vol. 9, no. 3, 2002.