

A Motion Detection Scheme For Wireless LAN Stations

Michael Wallbaum and Stefan Diepolder

{wallbaum, diepolder}@i4.informatik.rwth-aachen.de
Department of Computer Science, RWTH Aachen University
D-52056 Aachen, Germany

ABSTRACT

Wireless LANs not only provide an effective means of communication, but also allow to detect movement and immobility of a mobile node. This information can, for example, be employed by Wireless LAN positioning systems to reduce the lag of location estimates due to sample filtering. In addition, motion detection has many applications in context-aware computing. This paper proposes a simple scheme to reliably detect node movement based on received signal strength samples from one or more access points. The scheme is evaluated through extensive experiments in an office environment.

Keywords: Wireless LAN, context-awareness, motion detection

1 INTRODUCTION

Knowledge about whether a user or device is stationary or in motion can serve a number of purposes. The mobility status is an important part of a user's context and is thus of interest to context-aware systems, e.g. [13]. By itself it can be employed e.g. to change the behaviour of services and applications. Furthermore, it can be used to enhance the information provided by rather coarse positioning systems, such as the well-known Active Badge system [16] or other proximity detection schemes [2] [12] [9]. Also related to user location is the application of movement information to accurate Wireless LAN positioning systems, such as RADAR [1], Horus [18] and Ekahau [11]. These systems essentially perform pattern recognition on the received signal strength typically measured at the mobile node. In general, signal samples fluctuate heavily even in environments with few moving objects. To compensate for these fluctuations Wireless LAN positioning systems often smooth the samples by means of a moving average. Consequently, these systems suffer from a temporal lag, which becomes evident when tracking a moving node. On the other hand, without sample filtering the location estimates tend to be instable, which becomes evident when periodically locating a stationary client. A motion detector can eliminate the necessity for this tradeoff between lag and accuracy. When the client is stationary the location algorithm can employ a filter with a large window size to improve the accuracy. When the client is in motion the window size can be decreased or the filter can be bypassed completely, to minimise the lag. Note that this paper is a spinoff of the work on the Wireless LAN location determination scheme presented in [15]. This scheme mandates a motion detector, since it inherently models user mobility.

Specialised motion tracking and motion detection systems make use of a range of sensing technologies [17]. Examples include mechanical, inertial, acoustic, magnetic and optical sensing. Such dedicated systems typically provide highly accurate information concerning position and even orientation and are thus akin to positioning systems. Unfortunately, dedicated motion detectors require additional hardware, which is often unwieldy, impractical or simply not available. However, the applications described above do not necessarily benefit from accurate and complete information about the mobility status. For the purposes described above it is sufficient to know whether the user is moving or not.

Since Wireless LAN positioning systems were mentioned as a potential application of a motion detector, it should be noted that successive location fixes can in principle be used to infer the mobility status. The scale of motion that can be detected in this manner, depends on the accuracy of the employed positioning system. Simple proximity detection systems, for example, are typically not able to detect small-scale motion within a particular cell. However, even positioning systems with a high spatial and temporal resolution do not necessarily lend themselves to motion detection. For example, Wireless LAN positioning typically requires reference signal patterns and their associated locations. Access to this information may be restricted to certain user groups and devices or it may not be available at all. Consequently, positioning systems in general cannot effectively be used to detect *small-scale* node movement.

This paper presents a simple scheme to detect motion based on the received signal strength measured by a Wireless LAN node. The scheme solely relies on the data provided by the Wireless LAN network interface card and neither requires information about the environment nor the network. The following section provides some background information about the employed signal strength data. Section 3 then discusses how the reference data that is used to evaluate the proposed detector was obtained. In Sect. 4 the motion detection scheme is presented. This section also discusses an attempt to also estimate the direction of movement. Concluding remarks are made in Sect. 5.

2 BACKGROUND

Most commercial IEEE 802.11 [5] Wireless LAN equipment provides access to radio signal strength measurements. Access points periodically emit beacon frames, which carry information about the parameters and capabilities of a cell, known as basic service set (BSS) in 802.11 terminology. Mo-




Figure 1: Floor plan showing where stationary measurements were conducted.

bile nodes use this data when deciding which particular access point to associate with. Among other information, beacon frames report about the network name and the basic service set identifier (BSSID), i.e. the unique identifier of an access point. Implicitly these beacons also carry information about the link quality, which can be derived from the signal strength and the background noise. The node sweeps from channel to channel and records information from any beacon it receives. This process is called *passive scanning* and is performed regularly to determine the access point with the best link quality. In this manner the signal strengths of all visible access points can be determined. 802.11 specifies a second mode of scanning, where the node actively probes for the available BSS. This process is called *active scanning*. The amount of information that can be retrieved in this manner is essentially the same as with passive scanning. Note that a node cannot communicate during the scanning process. Still, scanning is necessary to ensure the best possible link quality.

The idea of the proposed motion detection scheme is to analyse the statistical characteristics of the sample series related to one or more access points. The underlying assumption is that a node's motion influences the sample series in a different manner than the movements of other objects in the environment and the signal fluctuations caused by multipath fading. The next section describes the reference measurements employed to test and evaluate the scheme, that will be presented in Sect. 4.

3 REFERENCE DATA

The following investigations are based on reference measurements conducted in an office wing of the Department of Computer Science of RWTH Aachen University. The wing, depicted in Fig. 1, has a dimension of about 40 m \times 15 m and is covered by three access points (APs) placed on this floor. Two types of 802.11b access points were used: AP_1 and AP_3 were L-11 access points from Lancom Systems and AP_2 was

an Orinoco 1200 from Lucent. All access points (marked by circles) were equipped with omni-directional antennas. The signal strength measurements were conducted by a number of standard notebooks equipped with IEEE 802.11b network interface cards. The samples were retrieved using the Wireless Research API (WRAPI) [14] developed by Microsoft Research and the University of California, San Diego. Measurements were carried out for stationary and mobile nodes as described in the following.

3.1 Stationary Measurements

The crosses indicate the 52 locations where the received signal strength (RSS) for immobile nodes was measured. The locations within rooms were chosen opportunistically, depending on the available space for the measurement equipment; in the hallway measurements were conducted systematically every meter. The measurements were carried out on four days during the late afternoon and early evening hours; at these times few (sporadically moving) people populated the environment. Three different notebooks equipped with Wireless LAN interface cards were used to measure the RSS. It was measured in dBm with a resolution of 1 dBm. Samples were taken for at least 10 minutes each with a frequency of one sample for second, resulting in more than 600 samples per access point and per location. The three access points were permanently visible at all locations, except for AP_1 which lacks samples for two locations in the large room in the lower left corner.

3.2 Mobile Measurements

For the mobile measurements a notebook equipped with a Wireless LAN card was carried at walking speed down the hallway from left to right and back. The mobile measurements took place at noon and the environment was slightly populated. As before, RSS samples of all visible access points




Figure 2: Profile of the walk along the hallway

were recorded once a second. The node's true position was measured using a laser range scanner placed in the doorframe below AP_3 . The range measurements were conducted roughly every 200 ms with a distance error of less than one centimeter. The profile of one such walk is depicted in Fig. 2.

4 MOTION DETECTION

In some application domains, e.g. mobile robotics, specialised hardware motion detectors are commonplace and thus motion detection is considered a minor issue. Autonomous robots are also typically in control of their actuators, i.e. their wheels or tracks. Consequently, when a robot uses its actuators to move and the motion detector returns corresponding measurements, there is a high probability that movement is actually taking place.

Of course special equipment for sensing movement can not always be assumed or required from users or their devices. Moreover, many applications do not foresee the control of actuators such as the user's feet. Hence, reliable detection of node movement solely based on RSS samples is a challenging problem, when one considers the high fluctuations in signal strength even when a terminal is stationary.

4.1 Requirements on motion detection

Since only RSS measurements with respect to the visible base nodes can be made use of, the mobility status of a mobile node must be derived through some filter process; the implementation of this process is called *motion detector* in the following. In order to be effective any such detector has to fulfil the following three requirements:

Low latency. The motion detector should detect node movement as quickly as possible, i.e. instantaneously at best. It should not rely on sample sets covering extended periods of time.

Location independence. The detection accuracy should be independent of the node's location, i.e. of the reception characteristics and number of visible access points.




Figure 3: RSS samples taken once per second from mobile and stationary terminal.

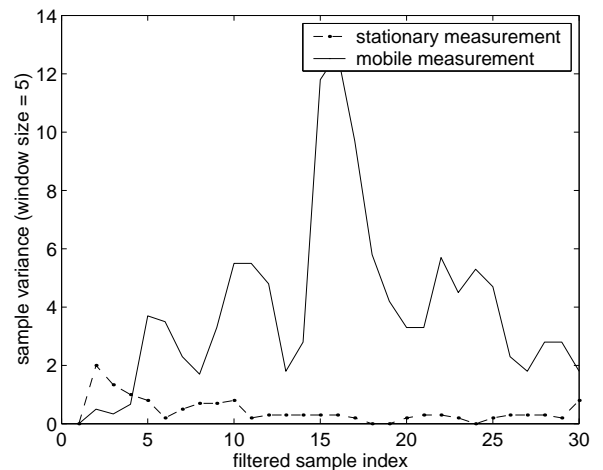


Figure 4: Sample variance of a sliding window with five samples.

Low error probability. The probability for falsely detecting motion, when a node is actually stationary should be as low as possible. Likewise, the probability for falsely reporting immobility, when a node is actually in motion should be as low as possible.

In addition, it would be instrumental for modelling movement, if the chosen filter provided an indication of the nodes's speed and direction.

4.2 Initial observations

Through experiments it was discovered that the sample variance of a series of measurements can be used to detect absolute movement. The sample variance of a set of RSS samples $\{r_i\}_{i=1}^n$ is defined as

$$(1) \quad s^2 = \frac{1}{n-1} \cdot \sum_{i=1}^n (r_i - \bar{r})^2$$

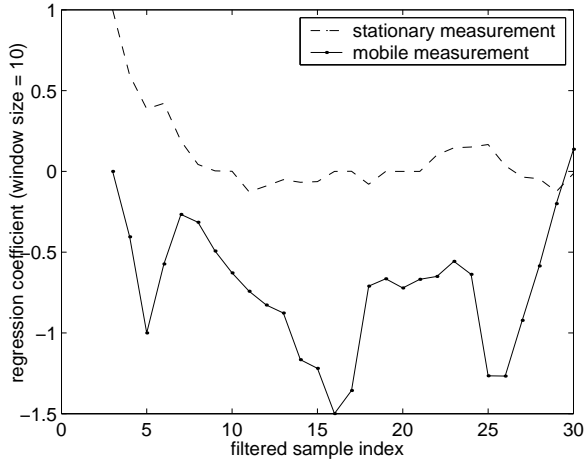


Figure 5: Regression coefficient for a sliding window with ten samples.

where \bar{r} is the sample mean. It is the best unbiased estimate for the true variance assuming that samples are normal distributed. Signal strength samples generally do not follow a normal distribution, although this is often assumed for reasons of simplicity. At this point, however, the true distribution function is irrelevant, since the aim is to detect the motion of a mobile node and not to estimate the true variance. For a fixed node the sample variance of a sliding window of samples seldom exceeded unity, even in highly dynamic situations. This value was well overstepped at normal walking speed. This phenomenon was first observed in [15] and later independently described in [8]. However, the former publication does not describe the details of the detection process while the latter publication uses a more complex algorithm based on hidden Markov models. Furthermore, the latter scheme only uses a single access point. It will later be shown, that the results improve considerably when measurements from multiple access points are employed.

Figure 3 shows the course of RSS measurements over a 30 s time period for a fixed and a moving node. In both cases samples were taken once a second. The corresponding filtered values are depicted in Fig. 4. The variance filter in this example uses a backward window n of five samples. (Note that values for sample indices less than the window size do not reflect the true performance of the filter.) The figure indicates that motion detection based on the use of the sample variance is a promising approach. Yet, some questions remain to be answered, namely how to set the threshold and the window size to ensure timely and reliable detection. The fraction of false positive and false negative detections given a set of parameters is also of interest.

4.3 Detecting direction of motion

Before the parameters and the quality of the variance-based approach are determined, a potentially more powerful technique is examined. The variance filter apparently only detects motion as such and does not indicate the mobile node's speed and direction. The curves in Fig. 3, however, suggest that




Figure 6: Ratio of false direction estimates depending on the size of the sliding window.

a tendency of the signal strength over time can be derived even when using few recent samples. Decreasing signal levels indicate motion away from the access point and vice versa, whereas stable RSS implies immobility. A method to capture the prevailing trend is to interpret a series of recent samples as a linear function of time. In other words, it is assumed that the sample set $\{r_i\}_{i=1}^n$ follows

$$(2) \quad r_i = a + b \cdot i + \varepsilon_i$$

where the ε_i are random errors. The gradient b is estimated by performing a linear regression analysis, through a least squares fit of the sample data. When the regression coefficient, i.e. the estimate for the gradient, is positive, there is motion towards the respective access point, whereas a negative coefficient indicates movement away from the AP. Figure 5 shows the course of the regression coefficient for the sample series depicted in Fig. 3. The technique apparently works well in this example – alas, it was found, that over the entire reference data set it is not able to provide reliable and timely results at the same time.

As an example, consider the hallway of the testbed depicted in Fig. 1. When moving from left to right the filter should yield a negative regression coefficient for the samples from AP_3 and a positive coefficient for AP_1 ; the opposite is true for movement from right to left. Figure 6 shows the ratio of false direction estimates depending on the sliding window size and the respective AP. For the data from AP_1 the ratio decreases as the sliding window gets larger. This could be expected, when assuming constant movement in one direction. However, even when using ten samples roughly one tenth of direction estimates are wrong. Such a large window conflicts with the requirement for low latency. Assuming one sample a second and an average walking speed of $1.2 \frac{m}{s}$, the node would have moved 12 m before the motion detector yielded a reasonable estimate for its direction. The results are even worse when using the data from AP_3 , which can be attributed to the non-monotonic behaviour of the received signal caused by a propagation phenomenon known as wave guiding. Em-

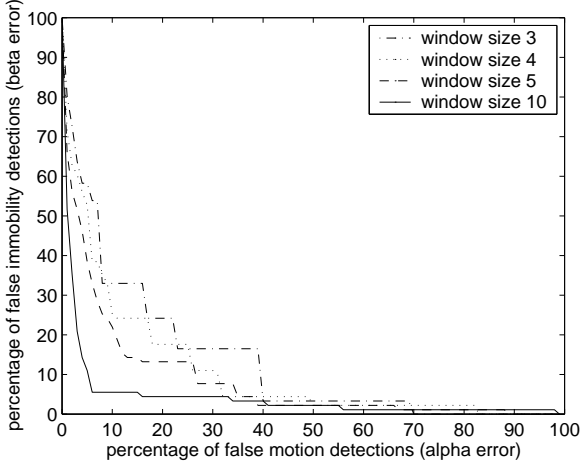


Figure 9: Beta error depending on alpha error for different window sizes.

playing a robust regression technique (e.g. as described in [4]) does not improve the results substantially. This approach to motion detection is thus not further investigated.

4.4 Motion detection with the variance filter

The requirement for timeliness demands that the motion detector can operate on very few samples. On the other hand, the requirement for accuracy calls for the detector to reliably distinguish motion and immobility of a node. Intuitively, these demands are to some extent contradictory. It can be expected that the higher the needed accuracy, the greater the number of samples. This supposition is substantiated by Fig. 7 and Fig. 8. The graphs show the dependency of the detector's accuracy on the number of samples and the fixed threshold; they are based on measurements conducted by a fixed and a moving node as described in Sect. 3. Sample variances above the threshold signal a moving node, whereas variances below the threshold indicate node immobility. Trivially, the higher the threshold, the lower the fraction of false mobility detections (Fig. 7) and vice versa (Fig. 8). The expectation of increased accuracy when using a larger number of samples is also confirmed. (Although in the case of the fixed node the accuracy surprisingly improves with smaller window sizes when the threshold is low.)

The interrelation between the two types of errors is summarised in Fig. 9. The curves show the dependency of the beta-error (i.e. falsely detected immobility) on the alpha-error (i.e. falsely detected motion) for different window sizes. It can be seen that both errors can be kept low, when choosing a large window size. Unfortunately, large delays conflict with the requirement for low latency, as has been laid out above. For a window of five samples the beta-error is about 22 percent, when setting the threshold such that the alpha-error is ten percent. This accuracy can be considered sufficient, yet it does leave room for improvement.

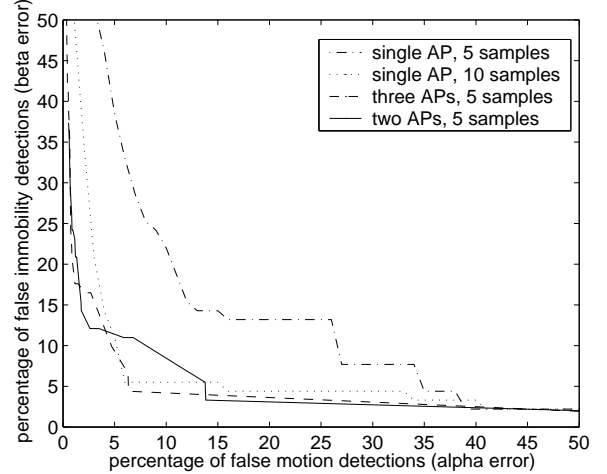


Figure 10: Accuracy of motion detection based on samples from a single AP vs. samples from multiple APs.

4.5 Using multiple access points

So far only samples with respect to a single access point were employed. The detector's accuracy can be well increased, when considering samples from multiple access points. Let $\{r_{ij}\}_{i=1}^{n_j}$ denote the sample set associated with AP_j . The function

$$(3) \quad \gamma_\tau(j) = \begin{cases} 1 & \text{if } \frac{1}{n_k-1} \cdot \sum_{i=1}^{n_j} (r_{ij} - \bar{r}_j)^2 \geq \tau \\ 0 & \text{otherwise} \end{cases}$$

yields unity, when the sample variance of AP_j 's sample set exceeds the threshold τ . Assuming k access points are visible, then the motion detection function

$$(4) \quad \Gamma_{\tau,k} = \prod_{i=1}^k \gamma_\tau(i)$$

equals unity, when the variances of *all* sample sets pass the threshold. This result is interpreted as node motion, whereas a value of zero signals immobility. The rationale behind $\Gamma_{\tau,k}$ is that true motion generally causes the sample values of all APs to fluctuate more strongly. Surely, a fixed node in a highly dynamic environment will experience great variations in the sample values associated with a subset of the visible APs. In practice, it is, however, unlikely that all sample sets are affected simultaneously.

Figure 10 compares the alpha- and beta-errors of the $\Gamma_{\tau,k}$ ($k > 1$) motion detector with those of the single AP scheme ($\Gamma_{\tau,1}$). (Note that the axes have been clipped off at 50 percent in order to provide more detail where it matters.) Using $k = 3$ APs and $n_j = 5$ recent samples from each AP yields significantly lower errors compared to the results based on data from one AP. Indeed, the accuracy corresponds to the single AP scheme with a window size of ten. Even the use of data from just two APs provides a great improvement. Often more than three access points are visible, yet it is recommendable to use only the three strongest APs to ensure that a sufficient number of samples is available.




Figure 7: Ratio of falsely detected motion, when the node is immobile.




Figure 8: Ratio of falsely detected immobility, when the node is in motion.

Based on these results the movement detector's threshold can be set to $\tau = 0.6$. With three access points and a sliding window of five samples per AP, this yields roughly five percent false positives and about ten percent false negatives. This represents a significant improvement over the more complex

scheme presented in [8], which reportedly shows about 13 percent false positive detections. (Results for false negatives are not given.)

Of course, these parameters are to some extent application-specific and can take other values. For example, an appli-




Figure 11: Percentage of falsely detected mobility for each location.

cation might trade accuracy for timeliness by decreasing the window size or it could modify τ to decrease the ratio of false positives or negatives. Also, note that these results are only valid for the particular environment where the measurements were conducted. However, the characteristics of the reference sample series are similar to those described in other publications, e.g. [10], [3], [6] and [7]. It is thus reasonable to assume that the scheme's detection accuracy will be comparable in other environments.

To conclude the discussion of the motion detector, the dependency of its accuracy on the particular location is investigated. Figure 11 visualises the percentage of false mobility detections for the locations where samples were collected. The circle radius corresponds to the error percentage. At positions with an error of five percent or more the numerical value is also shown. Apart from two exceptions all sample series leading to large errors were collected within rooms, i.e. with no line of sight to an access point. Two of these locations exhibit exceptionally large errors with 14 and 16 percent of false positives. True location independency of the accuracy demands that the errors are the same at all positions. This is certainly not the case when looking at the spread between the errors. Furthermore, similar sized errors seem to agglomerate, which can also be observed in the corridor. Although location independence of error is not given in the strict sense, from a practical point of view all errors can be regarded sufficiently low.

5 CONCLUSION

This paper has presented a simple motion detection scheme based on received signal strength measurements obtained from a Wireless LAN. The experimental evaluation has shown that the results provided by this approach are quite satisfactory, particularly when considering the scheme's simplicity. The employed reference data was sampled in a typical office en-

vironment; the data's statistical characteristics correspond to those reported by other research groups. It can thus be expected that the detection quality in other office environments is similar, though this will need to be verified.

An attempt to detect the direction of movement has been discussed as well. In principle, direction estimates are feasible, yet, a large backward window is required to keep the number of false direction estimates low. Since timeliness was identified as an important requirement for motion detection, the presented direction estimation approach was not further pursued.

Since the presented scheme is based on signal strength measurements it relies on the scanning procedures defined in the 802.11 standard. It has been mentioned that no communication can take place while the channels are scanned. Usually this is not considered a problem, as scanning to ensure link quality is carried out rather infrequently. The motion detector, however, works best with a continuous sample stream. Wireless LAN drivers thus need to be modified to conduct continuous passive scanning while no communication is taking place.

Additional future work on this topic includes transferring the scheme to other types of wireless networks, e.g. to mobile communication and mobile ad hoc networks. Furthermore, it should be investigated how the positioning and tracking performance of Wireless LAN location systems is actually affected when sample filtering is made dependent on the mobility status. It can be expected that the accuracy improves considerably for stationary clients since then the filter's sliding window can be made arbitrarily large. Still, quantitative results to support this supposition are necessary.

REFERENCES

- [1] P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system.

- In *Proceedings of IEEE INFOCOM 2000 Conference on Computer Communications. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 2, pages 775 – 84. IEEE, March 2000.
- [2] Keith Cheverst, Nigel Davies, Keith Mitchell, and Adrian Friday. Experiences of developing and deploying a context-aware tourist guide: the GUIDE project. In *Proceedings of the sixth annual international conference on Mobile computing and networking*, pages 20 – 31. ACM Press, August 2000.
- [3] A. F. de Toledo and A. M. D. Turkmani. Propagation into and within buildings at 900, 1800 and 2300 MHz. In *1992 IEEE 42nd Vehicular Technology Conference*, volume 2, pages 633 – 636, May 1992.
- [4] P. J. Huber. *Robust Statistics*. Wiley Series in Probability and Statistics. John Wiley and Sons Ltd., New York, first edition, January 2004.
- [5] IEEE 802.11 Working Group. *ANSI/IEEE Std. 802.11, IEEE Standards for Information Technology – Telecommunications and Information Exchange between Systems – Local and Metropolitan Area Network – Specific Requirements – Part 11: Wireless LAN Medium*, 1999.
- [6] G. J. M. Janssen and R. Prasad. Propagation measurements in an indoor radio environment at 2.4 GHz, 4.75 GHz and 11.5 GHz. In *Proceedings of IEEE 42nd Vehicular Technology Conference*, volume 2, pages 617 – 620, 1992.
- [7] K. Kaemarungsi and P. Krishnamurthy. Modeling of indoor positioning systems based on location fingerprinting. In *Proceedings of INFOCOM 2004. Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 2, pages 1012 – 1022. IEEE Computer Society, March 2004.
- [8] John Krumm and Eric Horvitz. LOCADIO: inferring motion and location from wi-fi signal strengths. In *Proceedings of Mobiquitous 2004, First Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services*, pages 4–14, August 2004.
- [9] A. LaMarca, Y. Chawathe, S. Consolvo, J. Hightower, I. Smith, J. Scott, T. Sohn, J. Howard, J. Hughes, F. Potter, J. Tabert, P. Powledge, G. Borriello, and B. Schilit. Place Lab: Device positioning using radio beacons in the wild. In Hans W. Gellersen, Roy Want, and Albrecht Schmidt, editors, *Pervasive Computing: Third International Conference, PERVASIVE 2005*, volume 3468 of *Lecture Notes in Computer Science*, pages 116–135. Springer Verlag, May 2005.
- [10] D. Molkdar. Review on radio propagation into and within buildings. *Microwaves, Antennas and Propagation*, 138(1):61 – 73, February 1991.
- [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*, 9(3):155 – 64, July 2002.
- [12] Bill Schilit, Anthony LaMarca, Gaetano Borriello, William Griswold, David McDonald, Edward Lazowska, Anand Balachandran, Jason Hong, and Vaughn Iverson. Challenge: Ubiquitous location-aware computing and the Place Lab initiative. In *Proceedings of The First ACM International Workshop on Wireless Mobile Applications and Services on WLAN (WMASH 2003)*, September 2003.
- [13] Maria Strimpakou, Ioanna Roussaki, Carsten Pils, Michael Angermann, Patrick Robertson, and Miltiades E. Anagnostou. Context modelling and management in ambient-aware pervasive environments. In Thomas Strang and Claudia Linnhoff-Popien, editors, *Location- and Context-Awareness: First International Workshop, LoCA 2005, Oberpfaffenhofen, Germany, May 12-13, 2005. Proceedings*, volume 3479, pages 2–15. Springer, 2005.
- [14] University of California, San Diego. Wireless Research API (wrapi) Homepage. <http://ramp.ucsd.edu/pawn/wrapi/>, June 2005.
- [15] M. Wallbaum and T. Wasch. Markov localization of wireless local area network clients. In R. Battiti, M. Conti, and R.L Cigno, editors, *Wireless On-Demand Network Systems. First IFIP TC6 Working Conference, WONS 2004. Proceedings.*, volume 2928 of *Lecture Notes in Computer Science*, pages 1 – 15. Springer Verlag, December 2003.
- [16] R. Want, A. Hopper, V. Falcao, and J. Gibbons. The Active Badge location system. *ACM Transactions on Information Systems*, 10(1):91 – 102, January 1992.
- [17] Greg Welch and Eric Foxlin. Motion tracking: No silver bullet, but a respectable arsenal. *IEEE Comput. Graph. Appl.*, 22(6):24–38, 2002.
- [18] M. Youssef, A. Agrawala, and A. Udaya Shankar. WLAN location determination via clustering and probability distributions. In *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications*, pages 143 – 50. IEEE Computer Society, March 2003.