

Room-Level Wi-Fi Location Tracking

Joshua Correa, Ed Katz, Patricia Collins, Martin Griss
Carnegie Mellon Silicon Valley
{joshua.correa, ed.katz, patricia.collins, martin.griss}@sv.cmu.edu
MRC-TR-2008-02
November 2008

Abstract

Context-aware applications for indoor intelligent environments require an appropriately accurate and stable interior positioning system to adapt services to the location of a mobile user or mobile device in a building. Different technologies provide a varying mix of resolution, accuracy, stability and challenges. In this paper we report on our experience using an existing Wi-Fi infrastructure without specialized hardware added to support location tracking. There are several approaches to track the location of Wi-Fi enabled devices within a building such as signal propagation models and signature matching. We found signature matching most effective in our environment. Signature matching is accomplished by storing Wi-Fi signatures (signal strengths observed for several detectable access points) for each room and comparing the current signature on the device to stored signatures to find the closest match. In this paper we explain experiments we conducted to explore and optimize Wi-Fi location tracking in one building. While we had hoped for more accurate positioning, we found that only room-level granularity was consistently and reliably achieved. The accuracy of Wi-Fi location tracking is improved as more signature points are stored, but is significantly reduced by the presence of people moving in the area. It also appears that strategically placed access points within a building can contribute to optimum room-level disambiguation of location. Use of a histogram of signal strengths for signatures at a single location may offer a good compromise between a single average and storing a large number of signatures needed for improved accuracy.

Background

An important goal for our group, which is researching smart spaces [SmartSpaces] to support home care for elders and other home and neighborhood intelligent environments, is to develop a reasonable indoor position system (IPS) similar in spirit to GPS. Given sufficient accuracy and stability, this capability can inform a variety of context-aware services, such as tracking location indoors, offering location-aware reminders, and even guiding robotic appliances. In the past we have explored infra-red, RFID and acoustic technologies [Katz 2008]. Each of these requires specialized sensors installed solely for location tracking purposes in the environment and on the user or the user's mobile device, and experience a variety of difficulties such as noise and occlusion. On the other hand, Wi-Fi-based location tracking makes use of existing Wi-Fi access points (APs) and the user's Wi-Fi enabled mobile devices, such as a mobile phone or laptop computer. The cost of installing and maintaining these increasingly ubiquitous Wi-Fi devices is already covered by the organization desiring Wi-Fi access, and not solely by the location tracking need.¹

There are two primary approaches to Wi-Fi location tracking indoors²: The signal propagation model approach is easier to set up but yields potentially less accurate results. The core of this approach is a model that relates signal strength to

¹ Our research suggests that successful location tracking requires multiple APs be detectable in each room. If this is the case, then the hypothesis that accurate room-level location tracking can be achieved with typical cost-effective AP configurations may need to be modified, particularly at home when users might only have a single AP. However, in many cases, Wi-Fi signals from neighbour's APs may be useful.

² Other approaches such as triangulation or time-of-arrival are feasible outdoors with clear line of sight, but do not work well indoors with multiple reflection paths.

distance from the Wi-Fi AP. Typically this model is based on empirical signal propagation data for frequencies in the Wi-Fi band, and does not take into account obstacles such as walls and furniture.³ The only set up required for this approach is the mapping of each AP on a floor plan. The user's mobile device then captures signal strengths received from each of the APs at the current location, and the location on the floor plan is determined using the signal propagation model. This approach typically yields an accuracy of around 3 to 10 meters [MOTOROLA 2008]. Signal propagation (RSSI) based approaches to locationing are ideal for environments that have established core use cases for Wi-Fi such as data, voice and video, such as tracking slowly moving assets. This is the least disruptive, most scalable and lowest cost approach, since it does not require the installation of a special client on the mobile object, nor does it require extensive training.

The signature matching approach requires more time spent setting up, but potentially yields much more accurate results. During the training phase of this approach, one or more sets of signal strengths (from the several detectable APs) are captured and stored for each of many sampled reference locations in a building. This builds a database of signal strength signatures for each reference location. Once this is done, the user's location can be determined by capturing the current signal strength signature and comparing it to the stored signatures for the closest match or matches. In the single match case, the user's location is assigned to be that of the closest matched reference location; in the multiple closest match case, the user's location might be determined to be the average of the reference locations of those matches. Signatures are stored for each reference location in multiple orientations (e.g., facing north, south, east, west), and perhaps at several times and in a variety of situations (e.g., people present or not). A room may be divided into multiple locations each about 1 m x 1 m and a signature stored for each of these locations. This means that a 5 m x 5 m area requires signatures to be captured for 25 locations, each with 4 orientations for a total of 100 signatures. This process is very time consuming but optimally it yields an accuracy of 2 to 3 meters [RADAR 2000].

Unlike the signal propagation approach, the signature matching approach works well indoors because it automatically takes into account obstacles such as walls and furniture (and indeed benefits from changes in signatures from room to room due to obstructions such as walls). However, any change in the layout of these obstacles requires a new set of signal strength signatures to be captured and stored. The signature matching approach is also sensitive to changes in hardware, of either mobile device or AP. So if an AP is replaced with one that is a different model or from a different manufacturer, this also requires a new set of signal strength signatures to be captured and stored. Another drawback of the signature matching approach is that at a particular location, the signal strengths received by a mobile device such as a phone are significantly different from those received by a laptop. Signal strengths received on different laptop models differ as well, though to a lesser extent than two different types of devices. This means that a signature database must be built for each device that must be tracked if the level of accuracy is to be maintained, though it may be possible in some cases to automatically adjust for different hardware [Kjærgaard 2006].

Since the RADAR paper published by Microsoft 8 years ago, indoor location tracking accuracy has hovered around 2 meters [RADAR 2000, HORUS 2007]. Although this is good enough for most applications, these systems have not been practical for various reasons, as demonstrated by the lack of widely deployed Wi-Fi location tracking systems. For example, it is very time consuming to train a signature-based Wi-Fi location tracking system. To try to mitigate some of these factors, our initial goal has been to determine location at a room level with sufficient reliability to allow a subset of interesting applications to be built and evaluated; other applications (such as robotic guidance) will wait for improved approaches, possibly using additional hardware. For this reason, the signature matching approach is more suited to our goals than the signal propagation approach. Room-level accuracy is also good enough for several of our SmartSpaces use cases involving location-aware reminders (e.g., [Collins 2008]). For example, if a user is scheduled for an upcoming meeting in a specific conference room, a location-aware calendar could send him a reminder if he is not already in the conference room. In the context of elder care, we could determine that the user has not moved from the bedroom to the

³ Some signal propagation researchers have augmented their models by using a count of intervening walls to model signal attenuation.

kitchen, as they usually do each morning, and could send an alert to a caregiver or family member so they could call or check up in person. By focusing on room-level accuracy, we significantly reduce the training overhead of the signature matching approach, while maintaining a level of accuracy appropriate to our use cases.

While not the focus of this phase of the work, different approaches have differing security and privacy issues, depending on where the collected signatures are processed to determine the location. Where the processing occurs may be determined by the size of the reference database, as to whether location requests can be tracked by a server, or performed completely on the device.

Related Work

We explored many existing products and open-source solutions, each of which addresses a different set of use cases. For example, PlaceLab is an open-source Java-based solution which accesses a vast database of Wi-Fi signatures to determine location at street-level accuracy. This solution is based on the signature matching approach and is similar to the solution used by products such as Google Maps when GPS signals are not available [PLACELAB 2005]. This solution is not well suited to indoor location tracking because there is not a high enough level of location accuracy to distinguish between adjacent rooms.

We also tried Motorola's RFS7000 solution, which is more suited to indoor location tracking. It is based on the signal propagation approach however. Due to its 3 to 10 meter accuracy, it frequently located the user in an adjacent room. This product determines the location of a device every time the device sends out special 'probe' packets. Typically these packets are only sent when the device is trying to connect to the APs. However, once the device is connected, no probes are sent, and the location of the device is not updated.⁴ This solution also requires proprietary Motorola APs and switches. For these reasons, this product was not well suited to our use cases. This product is better suited to determining the general location of a device at a large campus with several buildings. It could also be used to help secure Wi-Fi networks by restricting locations from where a device can connect to the network. For example, it can be used to deny access to devices trying to connect from buildings across the street from the APs. A similar product we hope to explore is Cisco's Mobile Services Environment (MSE) that also uses a specialized switch and AP's to add locationing to ordinary Wi-Fi access. [CISCO 2008].

In general, the signature matching approach can be quite good, especially for some use cases in (home) healthcare and education, but doesn't work for all environments since RF is dynamic. In particular, auto learning is not suitable for large production environments like hospitals since it involves a training period and also requires active participation from unskilled participants. Note that the transmit power of the APs can and will vary in a Wi-Fi installation due to dynamic RF management etc. and thus the signature can change on a regular or irregular basis. Most commercial Wi-Fi locationing use cases involve active RFID tags with refresh rates ~20sec and are well suited for Real Time Locationing Services (RTLS). Also, Wi-Fi client software is a required for the signature based schemes and is not practical in large deployments since the logistics of deploying a client on say 100,000 handhelds can be daunting, and access to low level radio functionality is required for these clients to extract the RSSI. Wi-Fi devices need a special software client to support this refresh rate and this is true even with signature based systems. Motorola's client can be easily ported to the Nokia N95 (Symbian) and solely relies on the 802.11 APIs.⁵

The final software application we tried is RedPin, which is an open source application designed to run on a Nokia N95 phone [RedPin 2008] with a supporting server. This product uses an extension of the pure RSSI signature matching

⁴ It may be possible to increase the probe frequency using Wi-Fi utilities such as NetStumbler: <http://www.netstumbler.com/downloads/>.

⁵ Motorola, private communication.

approach⁶, but eliminates the initial training overhead. Rather than acquire all the signatures up front, the product is designed to allow the users to input their current signature and label it with a location. As more users submit signatures, the system's accuracy improves. This is the most appropriate solution we have found so far as it is designed for a mobile device, and also addresses the issue of training overhead by gradually completing the training. However, we will need to run some experiments to determine whether the signatures obtained by one Nokia N95 can be used by another (model of) Nokia N95 device.

Experiments

In order to better understand the characteristics of our building and Wi-Fi setup, and to explore how accurate and stable a result we might expect, we performed several experiments. As part of this experimentation, we implemented a simplified version of signature matching, collecting signatures and using K-nearest neighbor⁷ matching and calibration for our environment.

Experiment #1 – effect of averaging

Set up three Wi-Fi transmitters (access points or APs) in a single area. Capture signal strength signatures at several positions in an area of a building. Compare signatures to see which positions can be reliably differentiated.

The experiment was run in the second floor open space of Carnegie Mellon Silicon Valley building 23. The measuring device was a Mac Book Pro. At the time data were collected, there were no people in the room and very few people in the building. However, there were stationary tables, couches, whiteboards, and other furniture in the open space. The three Motorola AP300 APs were very close to the device (within 10m); this is not a typical real-world setup. Data were collected for a single orientation at each location. The reason most signature matching approaches capture multiple orientations for each location is to account for the fact that a person would be holding the device as signal strengths were captured, and the human body is particularly good at obstructing Wi-Fi signals. These different orientations then represent the position of the user in relation to the device and the APs, rather than the actual orientation of the device [RADAR]. Since our device captured the signature using a time-delayed script, there were no people between the device and the APs. We also ensured that the measuring device was always facing the same direction to make sure that orientation did not interfere with our analysis.

⁶ Rather than using only RSSI strengths from several Wi-Fi APs, RedPin uses positive weighted contributions from a count of matching detected APs, a negative weighted contribution from un-matched APs as well as a weighted contribution from RSSI “distance”; furthermore, weighted contributions from recognized or matching Bluetooth and GSM cell IDs can also be included.

⁷ Using K=1, we used the Euclidean distance and the algorithm described by http://en.wikipedia.org/wiki/K-nearest_neighbor_algorithm




Figure 1: Location of Access Points in 4.8 m. by 9.9 m Open Space

Analysis

One hundred readings were taken at each of 18 locations (0,0; 0,1; 0,2; ... 2,5).

Ninety of the readings for each location were stored as signatures; the remaining ten were used to simulate a request to compute the current location and to test the accuracy of the location-tracking algorithm. Only signals from the 802.11bg radios on the APs were used, the 802.11a signals were discarded because they are not usually enabled on standard, widely deployed APs.

Two versions of the algorithm were used-- one where the 90 readings were averaged, resulting in one stored signature per location, the other where the 90 readings for each location were stored separately, resulting in 90 stored signatures for each location. The algorithm did not remove outlier data points when averaging. In both versions, there was no averaging of the remaining ten readings, which are used directly as test data.

The range of detected signal strengths was -19dB to -94dB, with 1dB granularity.

With Averaging:

average distance to nearest matching signature in signal space: 5.92 dB
standard deviation of distance to nearest matching signature in signal space: 4.88 dB
average distance to nearest matching signature in physical space: 1.12 m.

Without Averaging:

average distance to nearest matching signature in signal space: 0.85 dB
standard deviation of distance to nearest matching signature in signal space: 1.17 dB
average distance to nearest matching signature in physical space: 0.59 m.

The algorithm used to calculate distance between two signatures in signal space is:

Signature 1 (stored signature)

AP1: S_{1a1} dB

AP2: S_{1a2} dB

AP3: S_{1a3} dB

Signature 2 (current signature)

AP1: S_{2a1} dB

AP2: S_{2a2} dB

AP3: S_{2a3} dB

$$\text{Distance in signal space} = ((S_{1a1} - S_{2a1})^2 + (S_{1a2} - S_{2a2})^2 + (S_{1a3} - S_{2a3})^2)^{1/2}$$

The signature that yields the lowest distance in signal space is the closest match (K-nearest neighbor with K=1). In the case where there are multiple signatures with the same signal space distance, then the one with the greatest physical distance is chosen to depict a worst case scenario for accuracy. In the version of the algorithm without averaging, there were often multiple signatures with the same signal space distance to the current signature, some of which were from the same physical location. For example, four signatures from location (1,2) and two signatures from location (2,3) could have the same signal space distance from the current signature. In this case the location with the most matching signatures was chosen, i.e., location (1,2).

The reason for the high accuracy (< 1 dB) is most likely the proximity of the APs to the Wi-Fi device, and the lack of people moving around the building.

The comparison of averaging vs. no averaging in the stored signatures shows that storing multiple signatures for each location is clearly a better choice. The reason for this is that averaging assumes a uni-modal distribution of signal strengths for a particular AP and location. However, as our data show, there are often several distinct signal strengths received by the device from the same AP at the same location. The histogram in Figure 2 shows the primary signal strengths received and secondary peaks, which we assume are from signal reflections.




Figure 2: Signal Strength vs. Frequency of Occurrence at location (2,5)

There was no movement in the room during the data acquisition so the multiple signal strengths received from an AP are not due to a change in the environment. These multiple signal strengths could be caused by a pattern of interference between signals from the different APs, since there are several other APs in the building broadcasting on the same channel, as well as people moving in other rooms, unbeknownst to us.




Figure 3: Signal Strength vs. Frequency of Occurrence at location (0,0)

This leads us to believe that the multiple signal strengths are an effect of the multi-path phenomenon. Multi-path describes a situation where the signal from an AP reaches a device by several different paths. For example, it could travel the line-of-sight path to the device, as well as bounce off a wall and be reflected to the device. This would result in two different signal strengths since the signal that traveled the line of sight path would be considerably stronger than the signal that was reflected. When the device is in a corner, there is a smaller number of paths that a signal can take to the device and this is reflected in Figure 3.

Conclusions for experiment #1

It is thus possible to determine location with an average of 0.59 meters accuracy by storing a large number of signatures for each location. This level of accuracy is applicable when there are no people moving about in the room, and there are three APs within very close range (~10 m). There were multiple obstructions in the room (pillars, whiteboards, couches

etc.), so line of sight to the AP is not required for this level of accuracy, although it would probably be beneficial. The line of sight path is the cleanest signal path and typically fluctuates +/- 2dB over time.

Experiment #2 – reference device

For our second experiment, the hypothesis was that accuracy and reliability could be improved by placing a reference device in the room [PROXIMITY 2008]

We placed two identical laptops 0.5 m apart facing in the same direction in a room with six APs. We captured signal strengths on each laptop every few seconds for ~30 minutes. We then checked if fluctuations in signal strengths over time correlated between laptops. I.e., if one laptop sees a change in signal strength at a certain time, does the other laptop see a similar change in signal strength? There were two people in the room for the first few minutes; there were no people in the room for the rest of the 30 minutes.

Analysis:




Figure 4: Signal Strength for Two Laptops over Time for Access Point 1

For AP1 with MAC address 0:15:70:90:26:4, laptop B (red) shows a lot more jitter than laptop A (blue) even though they are only 0.5 m. apart. Also there is an event (circled) which causes a change at laptop B's received signal strength for AP1, but no corresponding change at laptop A's received signal strength. It is not clear what caused this change at laptop B, because there was no movement in the room. The AP is located in the same room as the laptops. It is possible that this event was simply a random period in which the signal strength did not vary quite as much as it did over the rest of the experiment.




Figure 5: Signal Strength over Time for Two Laptops for Access Point 2

For AP2 with MAC address 0:15:70:90:23:bc, laptop A (blue) shows a lot more jitter than laptop B (red), the opposite of the situation for AP1. AP1 and AP2 readings were taken simultaneously. The signal strength readings are clearly uncorrelated.




Figure 6: Signal Strength over Time for Two Laptops for Access Point 3

For AP3 with MAC address 0:15:70:90:23:d0, both laptops receive similar signal strengths, although the jitter for laptop A is higher than that for laptop B. However, at any point in time, the signals appear to be uncorrelated.

Conclusion for experiment #2

There is no obvious correlation between the signal strengths received by each of the laptops. The signals for each laptop appear to vary independently. Furthermore, we did not see significant signal strength changes between when there were two people in the room and when there was no one in the room. This may be due to positioning of APs in relation to the laptops and the people in the room. It may be worthwhile looking into the effect of a full room (~5 people) vs an empty room, as two people may not be enough to cause a significant difference. It is not clear why even when there are no people and no movement in the room, there is still a large difference between the signal strengths received at each of the laptops that were only 0.5 m. apart. Signals from AP1 and AP2 show large differences at the two laptops, while signals from AP3 show very little difference between the two laptops. These variations should be kept in mind when designing Wi-Fi-based location-tracking systems aimed at high positional accuracy. Some of these discrepancies between the signals received at each laptop may be explained by "small-scale variations" as described in [HORUS 2007]. It is also

interesting to note that [PROXIMITY 2008] seems to use a similar reference model approach via a statistical Gaussian Mixture Model to tell if two systems are nearby, yet we did not see enough consistent signatures to feel confident adopting this model.

Experiment #3 – signal propagation using specialized Wi-Fi equipment

Motorola Wi-Fi (RFS 7000, AP 300)

For the third experiment, our hypothesis was that we would be able to track Wi-Fi devices in a building, locating them with room-level accuracy, using the Motorola RFS7000 system. [MOTOROLA] The Motorola system is made up of a switch (RFS 7000) and several APs (AP 300) which report to the switch. The switch runs software that allows you to set up the location tracking software and displays coordinates of Wi-Fi devices.

To set up location tracking, you need to enter the dimensions of the area you are tracking (this version of the software does not support multiple floors), and the position of each AP. Once it is set up, you can check the coordinates of Wi-Fi devices within the tracking area. This system uses the signal propagation approach.

During initial tests, the software was not able to accurately differentiate between room 211 and room 212, two adjacent rooms measuring approximately 6 meters by 7.6 meters. However, this may be due to bad positioning of APs. The APs should be wall mounted at approximately 2 m. above the floor, in an asymmetrical layout; we did not wall-mount the APs. We had three APs located along the wall that separates the two rooms. Also, the software does not update the location of Wi-Fi devices very often. The location is calculated every time a Wi-Fi device sends out a 'probe' signal asking surrounding APs to identify themselves. This probing is typically done only when the device first tries to connect to the network. Motorola has a Microsoft Windows-based Java application that will continuously send out probes so that the location will be updated more often; but the application needs to be installed and running on a Windows-based Wi-Fi device, and will not work on our Nokia N95 devices. Due to its slow refresh rate and low accuracy, this system is not well suited to our needs and we did not pursue this approach further. Possible next steps would be to arrange the APs so that they are optimally located. We would also benefit from using the Motorola Java utility (or a Wi-Fi utility such as NetStumbler) on the Wi-Fi device to increase the refresh rate.

Conclusions and Next Steps

AN existing Wi-Fi infrastructure potentially provides a low-cost way to track Wi-Fi enabled devices in a building. Since Wi-Fi is a widely deployed technology, this method of location tracking requires no additional hardware. However, several factors can significantly alter the Wi-Fi environment. Therefore, we need more work to develop a system that can robustly handle these variations.

We are currently working with RedPin software on the Nokia N95 phone to build a system that is reliably able to track location at room-level granularity [RedPin 2008]. RedPin deals with changes in the Wi-Fi environment by accumulating location-labeled signatures submitted by users over time. As more signatures are submitted, this system will contain multiple signatures for a location, each representing the location in a different state. For example one signature may represent a location with few people, whereas another signature may represent the same location with multiple people. As the system accumulates signatures, the Wi-Fi environment is more accurately modeled in the system, and location tracking becomes more accurate.

Signature-based Wi-Fi location tracking systems may benefit from storing histograms of signal strengths rather than the full set of multiple signal strength readings or only their averages. Storing multiple signatures for each location is a good way to deal with the fact that multiple signal strengths could be received from a single AP due to the multi-path

phenomenon or other noise effects. However, a large number of signatures need to be stored for each location to capture all the permutations of signal strengths reliably from each AP at a particular location. This significantly increases the space required to store signatures and the number of signatures that must be compared to find the closest match.

A better approach may be to store a histogram of signal strengths in each signature rather than a single reading. Accuracy may also be improved by building a histogram signature for the current location and comparing that with stored histogram signatures. This is because a histogram represents various levels of signal strength received at a location due to the multi-path phenomenon, and therefore captures a more accurate description of the signal strength environment at that location than a single reading.

The histogram signatures may also help solve the problem of multiple or moving people in the space, since a person would obstruct only certain paths of the signal, which would alter part of the histogram but leave the rest unchanged. To test this, we should perhaps repeat experiment #1 with people on the room; if collecting additional signatures with people present can help, then we should do the same. To some degree, Redpin accounts for this, since measurements are collected as people go about their business, either alone or interacting with others.

Two similar devices that are very close to each other may receive significantly different signal strengths even when there is no movement in the room. This may be attributed to "small-scale variations" as described in [HORUS 2007]. "These variations happen when the user moves over a small distance (order of wavelength). This leads to changes in the average received signal strength. For the 802.11b networks working at the 2.4 GHz range, the wavelength is 12.5 cm and we measure a variation in the average signal strength up to 10 dBm in a distance as small as 7.6 cm (3 inches)" [HORUS section 2.3.2]

Although the signal propagation approach is much easier to set up and scales much more effectively to many mobile objects without needing a special client, existing systems such as the Motorola RFS7000 have currently significantly lower accuracy than signature-based systems. The Motorola system we experimented with also had a low refresh rate, making the system unsuitable for our use cases⁸. It should be noted that the Motorola system we experimented with used early beta versions of their location tracking firmware. Future models should incorporate additional algorithms, mechanisms and heuristics to account for walls and reflections, and address other issues.

Our overall conclusion is that for our use cases we should use a signature based system that provides an upfront calibration option (in addition to auto learning), and are modifying the basic RedPin system to do this. Future work will also explore ways for the client software and server to cooperate with the Wi-Fi infrastructure to mitigate the effects of dynamic RF management. Finally, security related to the server access by other parties to track a user's location requests must be addressed.

Acknowledgements

This work was supported in part by grants and equipment from Motorola, Nokia and SAP. We thank Tony Lin and Ilya Landa for assistance and comments, as well as useful suggestions from our colleagues at Motorola.

References:

- [CISCO 2008] Cisco 3300 Series Mobility Services Engine, <http://www.cisco.com/en/US/products/ps9742/index.html>
- [Collins 2008] Collins, Patricia. "Aging Independently." Enhancement to the Continua Health Alliance elder care scenario,

⁸ Though the use of NetStumbler or other specialized client software could increase the frequency of AP probing and thus improve the tracking of moving objects.

internal report, October 2008, MRC-TR-2008-01.

- [HORUS 2007] Youssef, Moustafa and Ashok Agrawala, "The Horus location determination system", Springerlink, 4 Jan. 2007, <http://www.springerlink.com/content/m67wm58j47n51115/>.
- [Katz 2008] Katz, Ed and Harry Sameshima. "Experiences with Crickets/Ultrasound Technology for 3-Dimensional Locationing within a Smart Environment." Mobility Research Center, Carnegie Mellon Silicon Valley, MRC-TR-2008-03.
- [Kjærsgaard 2006] Kjærsgaard, Mikkel Baun "Automatic Mitigation of Sensor Variations for Signal Strength Based Location Systems," Proceedings of 2nd International Workshop on Location- and Context-Awareness (LoCA 2006).
- [MOTOROLA 2008] "Motorola Real Time Location System (RTLS) Engine." Motorola. Sep. 2008. http://www.motorola.com/staticfiles/Business/Products/Wireless%20LAN%20Devices/RTLS%20Engine/Document/Static%20File/RTLS_SS_0908.pdf.
- [PLACELAB 2005] LaMarca, Anthony et al. "Place Lab: Device Positioning Using Radio Beacons in the Wild." In Proceedings of Pervasive Computing 2005, <http://www.placelab.org/publications/pubs/pervasive-placelab-2005-final.pdf>.
- [PROXIMITY 2008] Carlotto, Alessandro et.al., "Proximity Classification for Mobile Devices Using Wi-Fi Environment Similarity," Proceedings of MobiCom'08 MELT'08 Workshop, Sept 2008.
- [RADAR 2000] Bahl, Paramvir and Venkata N. Padmanabhan, "RADAR: An In-Building RF-based User Location and Tracking System." Proceedings of IEEE Infocom 2000, Tel-Aviv, Israel March 2000.
- [RADAR2 2000] Bahl, Paramvir, Venkata N. Padmanabha, and Anand Balachandran, "A Software System for Locating Mobile Users: Design, Evaluation, and Lessons." Technical Report MSR-TR-2000-12, Microsoft Research, February 2000.
- [REDPIN 2008] Bolliger, Phillipp, "Redpin - Adaptive, Zero-Configuration Indoor Localization through User Collaboration." MELT'08: Proceedings of the first ACM international workshop on Mobile entity localization and tracking in GPS-less environments, Sep. 2008.
- [SmartSpaces] Griss, Martin, "Smartspaces." Carnegie Mellon Silicon Valley. 2008. <http://sv.cmu.edu/research/smartspace>.