

Towards Fine-Grained Radio-Based Indoor Location

Jie Xiong
Department of Computer Science
University College London
United Kingdom
j.xiong@cs.ucl.ac.uk

Kyle Jamieson
Department of Computer Science
University College London
United Kingdom
k.jamieson@cs.ucl.ac.uk

ABSTRACT

Location systems are key to a rich experience for mobile users. When they roam outdoors, mobiles can usually count on a clear GPS signal for an accurate location, but indoors, GPS usually fades, and so up until recently, mobiles have had to rely mainly on rather coarse-grained signal strength readings for location. What has changed this status quo is the recent trend of dramatically increasing numbers of antennas at the indoor AP, mainly to bolster capacity and coverage with multiple-input, multiple-output (MIMO) techniques. In the near future, the number of antennas at the access point will increase several-fold, to meet increasing demands for wireless capacity with MIMO links, spatial division multiplexing, and interference management. We thus observe an opportunity to revisit the important problem of localization with a fresh perspective. This paper presents the design and experimental evaluation of ArrayTrack, an indoor location system that uses MIMO-based techniques to track wireless clients in real time as they roam about a building. We prototype ArrayTrack on a WARP platform, emulating the capabilities of an inexpensive 802.11 wireless access point. Our results show that ArrayTrack can pinpoint 33 clients spread out over an indoor office environment to within a 36 cm location accuracy.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless communication*

General Terms

Design, Experimentation, Measurement

Keywords

Wireless, MIMO, AoA, 802.11, Multi-antenna system, ArrayTrack, Indoor Localization

This material is based on work supported by the European Research Council under Grant No. 279976.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

HotMobile'12, February 28–29, 2012, San Diego, California, USA.
Copyright 2012 ACM 978-1-4503-1207-3 ...\$10.00.

1. INTRODUCTION

The proliferation of mobile computing devices continues, with handheld smartphones, tablets, and laptops a part of our everyday lives. Outdoors, these mobile devices largely enjoy a robust and relatively accurate location service from Global Positioning System (GPS) satellite signals, but indoors where GPS signals don't reach, providing an accurate location service is quite challenging.

Furthermore, the demand for accurate location information is especially acute indoors. While the few meters of accuracy GPS provides outdoors are more than sufficient for street-level navigation, small differences in location have more importance to people and applications indoors: a few meters of error in estimated location can place someone in a different office within a building, for example. Applications such as information overlay on the screen of a camera phone or in-building navigation also demand sub-meter accuracy [12]. A solution that offers a centimeter-accurate location service indoors would enable these and other exciting applications in mobile and pervasive computing.

Using radio frequency (RF) for location has many challenges. First, the many objects found indoors near access points (APs) and mobile clients reflect the energy of the wireless signal in a phenomenon called *multipath propagation*. This forces an unfortunate tradeoff that most existing RF location-based systems make: either model this hard-to-predict pattern of multipath fading, or leverage expensive hardware that can sample the wireless signal at a very high rate. Most existing RF systems choose the former, building maps of multipath signal strength [4, 5, 28, 35], or estimating coarse differences using RF propagation models [10, 14], achieving an average localization accuracy of between 60 cm [35] and meters: too coarse for the applications at hand.

Systems based on ultrasound and RF sensors such as Active Badge [29], Bat [30], and Cricket [18] have achieved accuracy to the level of centimeters, but usually require dedicated infrastructure to be installed in every room in a building, an approach that is expensive, time consuming, and requires maintenance effort.

The key observation we make here is that in recent years, a new opportunity to design indoor location systems has presented itself: an ever-increasing number of antennas at an ever-increasing number of nearby APs. WiFi access point designers use multiple antennas per access point to bolster capacity and coverage with multiple-input, multiple-output (MIMO) techniques. A commercial wireless circular array with 16 antennas has been marketed by Xirrus [1] since 2010. We expect that in the future, the number of antennas




Figure 1: Number of production WiFi access points (APs) reachable from every client location in our experimental testbed. Transmissions from most locations reach seven or more production APs.

at the access point will increase several-fold, to meet the demand for MIMO links and spatial division multiplexing techniques [2, 26], both of which increase wireless capacity. Figure 1 shows the number of nearby 802.11 APs reachable from 33 locations in our testbed. With our WARP APs excluded, transmissions from most locations reach seven or more production APs, with all but about five percent of locations reaching five or more production APs.

ArrayTrack is a system that exploits the increasing number of antennas at commodity APs to provide fine-grained location for mobile devices in an indoor setting. Multiple ArrayTrack APs cooperate to use angle-of-arrival (AoA) information from clients’ incoming signals to determine the client’s location. While AoA techniques are in wide use in radar and acoustics, the specific challenge in using these techniques indoors is the strong multipath RF propagation present in most buildings. To address this problem, we introduce novel algorithms that eliminate the effects of multipath, even in the relatively-common situations when no energy arrives on the direct path between the client and one of the APs.

ArrayTrack advances the known state of the art in RF localization by combining novel algorithms that synthesize and cluster the results of many APs’ AoA information with “best of breed” algorithms for AoA determination at each access point. A key feature of our approach is that by operating within the physical layer, we can estimate location based on detecting part of a single packet’s preamble at -17 dB signal-to-noise ratio (SNR), an SNR far lower than that needed to actually receive a packet. This allows many ArrayTrack APs to cooperate to localize clients that are far away from some of them. It also allows ArrayTrack to determine a client’s location to within centimeters in real time with a relatively small computational demand, with each AP processing on the order of 100 16-bit baseband RF samples, and a backend database running a simple clustering algorithm on the order of 50 location coordinates.

We implement ArrayTrack on the Rice WARP FPGA platform, and evaluate in a 33-node wireless network deployed over one floor of a busy office space. Preliminary experimental results in this setting show that ArrayTrack can localize 802.11 clients to a median 25 cm accuracy, localizing 90% of all clients to within 72 cm of their true location. Furthermore, we expect error to decrease under more realistic experiments utilizing a more complete prototype,

planned for future work. To our knowledge, these are the most accurate location results to-date for an RF-based location system that does not require infrastructure apart from a normal density of nearby WiFi access points.

2. DESIGN

ArrayTrack’s design can be separated into three different stages: first, ArrayTrack leverages techniques to detect packets at very low signal strength, so that many access points can overhear a single transmission. Next, ArrayTrack generates a pseudospectrum, or estimate of likelihood versus bearing, at each access point. Finally, ArrayTrack combines these estimates, discarding outliers, to generate a final location estimate.

2.1 Packet detection

To obtain bearing information for a client, an ArrayTrack access point needs to overhear transmissions from this client. These transmissions can be of any type of frame: a control frame (RTS/CTS) or a data packet. For a wireless transmitted packet, the most robust part is the preamble as it is normally transmitted at base rate and furthermore, the preamble contains a known time domain sequence for the receiver to detect the existence of a packet.

The preamble part of 801.11a/g packets contains known short and long training symbols. We propose and implement a modified version of the Schmidl-Cox algorithm [21] to detect the short training symbols. As there are 10 short training symbols in the preamble, we apply a moving average filter with window size equal to 10 short training symbols to enhance the detection of the packet. Our modified version of Schmidl-Cox implementation can detect the preamble of a packet at SNR as low as -17 dB. The WiSpy equipment we used to scan the number of APs works at a minimum SNR of 1-3 dB. So the clients will be heard by even more APs with our proposed scheme. Once a packet is detected, multiple samples of the packet are recorded to process AoA spectrum for this transmission.

2.2 Pseudospectrum generation

In both indoor and outdoor wireless channels, a sender’s signal reflects off objects in the environment, resulting in multiple copies of the signal arriving at the access point; this phenomenon is known as multipath. A *pseudospectrum* of a client’s received signal at a multi-antenna access point is the estimation of the incoming signal’s power as a function of angle of arrival as shown in Figure 2. With this spectrum, we can find out the direct path and multi-path bearing of signal’s arrival. Because of strong multi-paths, the direct-path signal may not be stronger than the multi-path reflections. The highest peak may not correspond to the true direct path bearing which can be used for localization which makes the indoor localization challenging. The pseudospectrum generation step samples a single packet at the multiple antennas of an access point, and generates an estimate of power versus angle.

ArrayTrack only needs a very small part of the packet to generate a pseudospectrum, it works with any part of the packet, and does not need to decode the packet. So once ArrayTrack has detected the preamble, it records a small part of it. In principle, one time domain packet sample (a 1,000-byte packet transmitted at 6 Mbps sampled at 40 MHz contains around 50,000 samples) will work for our scheme.




Figure 2: A *pseudospectrum* of a client’s received signal at a multi-antenna access point estimates the incoming signal’s power as a function of its angle of arrival.

However, recorded packets will be affected by background noise and interference from other senders. We therefore capture multiple samples to obtain mean phase difference. For a 1000-byte packet, even receiving a fragment of the preamble as small as a single byte (a duration of $1.5 \mu\text{s}$), is enough data for ArrayTrack to work well. The channel coherence time in an indoor environment is on the order of milliseconds, and is much longer than the $1.5 \mu\text{s}$ required here. ArrayTrack therefore does not add any usability overhead to the WLAN.

The best known AoA estimation algorithms are based on the MUSIC [22] algorithm, which performs an eigenstructure analysis of a *correlation matrix* formed by samplewise-multiplying the raw signal from the l th antenna with the raw signal from the m th antenna, then computing the mean of the result. When the incoming signals are highly correlated, as is a direct path reception and its multipath reflections, the traditional MUSIC algorithm does not work well. We therefore apply a modified version of MUSIC with spatial smoothing [23]. For an eight-antenna array, the MUSIC algorithm here involves matrix multiplications of size 8-by-1 and 1-by-8. Because of the small size of these matrices, this process is not computational demanding and runs very fast.

2.3 Pseudospectra synthesis

In this step, ArrayTrack combines the pseudospectra of several APs into a location estimate. If N APs generate pseudospectra $P_1(\theta), \dots, P_N(\theta)$ and we wish to consider the likelihood of the client being located at position \mathbf{x} as shown in Figure 2, we compute the bearing of \mathbf{x} to AP i , and the likelihood of location \mathbf{x} , $L(\mathbf{x})$, as

$$L(\mathbf{x}) = \prod_{i=1}^N P_i(\theta). \quad (1)$$

With Equation 1 we search for the most likely location of the client with standard search techniques. This process takes much longer compared with the previous packet detection and pseudospectrum generation processes. In the future we will optimize this search process to further decrease the system latency of ArrayTrack.




Figure 3: The ArrayTrack prototype with a linear antenna arrangement.

2.4 Outlier rejection

With multiple APs available, we have the option to choose all or some of them to process location information for clients. Each AP has a pseudospectrum for a client as described in the previous section, and combining the pseudospectra gives us a location for the client. On average, more APs yield a more accurate location result. However, for one particular client, this may not always be true. One possible reason is that the linear array exhibits more accuracy for client angles facing the broad side of the array rather than sighting down the line of the array when we apply the MUSIC algorithm to compute the pseudospectrum. So including one more AP for a particular client may degrade the accuracy level rather than enhancing it. Also, some of the APs may experience strong multipath reflections generating spurious peaks in the pseudospectrum.

With $N \geq 3$ APs deployed, we can choose any number M (with $3 \leq M \leq N$) of APs to process location information, resulting in $\sum_{M=3}^N \binom{N}{M}$ subsets of APs. Choosing the maximum number of APs generally presents us a safe result but not always the best result. So we employ Mahalanobis distance [15] to compare all the locations generated by all the different AP combinations. By applying a threshold, we can separate these location results into two groups. We reject the group with greater Mahalanobis distance as outliers. Then we average the remaining to obtain the final location result.

3. IMPLEMENTATION

Our prototype access point, shown in Figure 3, uses two Rice WARP FPGA-based wireless platforms. Each WARP platform is equipped with four radio front ends and four omnidirectional antennas. The WARPs run a custom hardware design of the Rice WARPLab hardware and software. All eight antennas are phase-calibrated with respect to one of them as described in previous work [33]. The clients in our experiments are Soekris computers with 802.11g WiFi interfaces, each equipped with two antennas.

4. EVALUATION

In this section, we present testbed results to show how ArrayTrack performs in an indoor office environment. First we present summary statistics on client location in our testbed and evaluate the efficacy of the AP clustering algorithm presented in the previous section. Then, we show the impact that the number of APs participating in ArrayTrack has on accuracy.

To measure ground truth in the location experiments presented in this section, we used scaled architectural drawings




Figure 4: Testbed environment: Soekris clients are indicated with small dots, and access point locations are numbered 1–6.

of our building combined with measurements taken from a Fluke 416D laser distance measurement device, which has an accuracy of ± 1.5 mm over 60 m.

Due to budget constraints, we used one WARP AP, moving it between the different locations marked on the map in Figure 4 and receiving new packets to emulate many APs receiving a transmission simultaneously. Consequently, the results we report next *overestimate* the magnitude of the location error that our system has described, by the following reasoning. The time interval between moving the WARP AP from location to location and taking measurements was approximately one day, well outside the wireless channel coherence time, or the time it takes for the wireless channel to change because of motion of objects nearby. Assuming that the nearby object motion between our measurements at different APs is random and uncorrelated with the wireless channel, such random changes in the wireless channel can only on average *add* error to our system.

The eight antennas attached to the WARP were placed in a linear arrangement, with antennas spaced at a half wavelength distance (6.1 cm). We placed the prototype AP at points numbered in our testbed floorplan, shown in Figure 4. We placed the 33 clients roughly uniformly over the floorplan, covering areas both near to, and far away from the AP. Some of them were in the same room with the AP while most are in other rooms. We put some soekris boxes near metal, wood and plastics to make our experiments more comprehensive. We also placed some clients behind concrete pillars in our office so that the direct path between the AP and client is blocked, making the situation more challenging.

4.1 Location accuracy

We first evaluate how accurately the pseudospectrum computation and pseudospectrum combining elements of ArrayTrack localizes clients by capturing packets from each of the 33 Soekris clients at the APs. Figure 5 shows the cumulative distribution of location error across clients using measurements taken at three to six APs and combined with Equation 1. The curves labeled “3 APs,” “4 APs,” and “5 APs” in the figure show raw location error results computed with Equation 1 across all different AP combinations. We see that the general trend is that average error decreases with an increasing number of APs. The median errors decrease

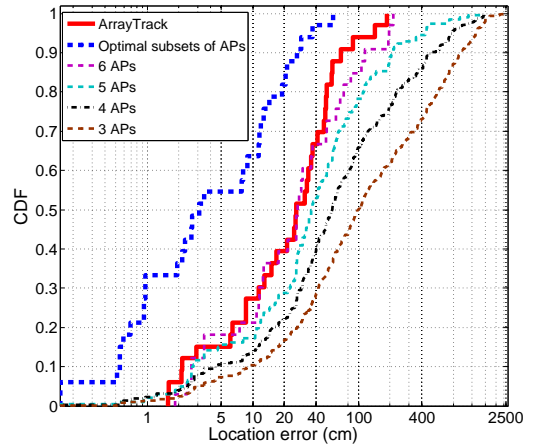


Figure 5: Cumulative distribution of location error across clients using measurements taken at all combinations of three, four, five, and six APs. The “Optimal subset of APs” curve picks (*a posteriori*) the subset of APs that comes closest to the ground truth location of the client.

from 100.0 cm for 3 APs to 56.7 cm for 4 APs to 36.1 cm for 5 APs to 26.5 cm for 6 APs respectively. The average error for 6 APs we achieved is 41cm in strong multipath indoor environments.

The curve labeled “ArrayTrack” includes the outlier detection algorithm presented in Section 2.4. With this algorithm, we further improve the average value to 36.4 cm and median value to 25.6 cm. We measure 90%, 95% and 99% of clients to be within 72 cm, 145 cm and 195 cm respectively of their actual positions.

The curve labeled “Optimal subset of APs” shows the error obtained by an omniscient scheme that picks the subset of APs that comes closest to the ground truth location of the client. This curve shows that with the right postprocessing (in future work), our scheme might yield a 5 cm median location error. We speculate further about how this might be realized in Section 6.

4.2 Effects of number of APs

Next, we show the effect of number of APs on the location accuracy level with one typical example selected from one of the 33 nodes. For all our experiments, linear array is applied which can only differentiate signal from one array side which is 0–180 degrees. This inherent ambiguity requires more APs to be present to locate clients. If a circular array, which can differentiate signal from 0–360 degrees was applied, the number of APs required would be decreased. However, a circular array requires a greater number of antennas to achieve the same accuracy level as a linear array.

In Figure 6, we present heat maps of the location of client A in Figure 4, with the number of APs varying. The heat map indicates the probability of the location of the client, with a darker hue denoting higher probability. The true client location is also marked on all the heatmaps with a small square. We see that with four APs, the location of the client is identified, but so is another area, causing possible false positives for our classification. Adding more APs removes this ambiguity. Note that not in all the scenarios we need 5 APs to identify the location. When the number of

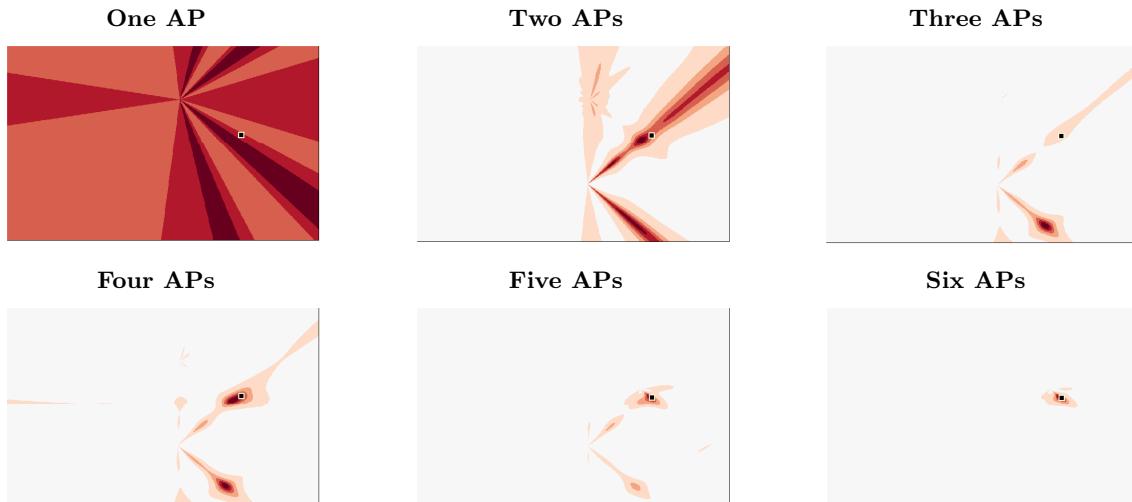


Figure 6: Heat maps showing the location likelihood of client A in Figure 4 with differing numbers of APs computing its location. We denote the ground truth location of client A in each heat map with a small dot.

lobes on the spectrum is small and some of the high probability false positive areas are far out of the floorplan, we may need only 3 APs to identify the clients. With the circular array, this can even be possible with 2 APs.

5. RELATED WORK

The most widely used RF-based approach for location uses average received signal strength (RSS) from packets, usually measured in units of whole decibels. While readily available from commodity WiFi hardware at this granularity, the resulting RSS measurements are very coarse compared to the physical-layer information we use in ArrayTrack, and so incur an amount of quantization error, especially when few readings are present.

There are two main lines of work using RSS; the first, pioneered by RADAR [4, 5] builds “maps” of signal strength to one or more access points, achieving an accuracy on the order of meters [20, 25]. Later systems such as Horus [35] use probabilistic techniques to improve localization accuracy to an average of 60 centimeters when an average of six access points are within range of every location in the wireless LAN converge area, but require large amounts of manual calibration. While some work has attempted to reduce the calibration overhead [11], mapping generally requires significant calibration effort. Other map-based work has proposed using overheard GSM signals from nearby towers [28], or dense deployments of desktop clients [3]. In contrast to map-based techniques, the experimental results we show here achieve better location accuracy from very small numbers of detected packets, with no calibration steps required.

The second line of work using RSS are techniques based on mathematical models. Some of these proposals use RF propagation models [19] to predict distance away from an access point based on signal strength readings. By triangulating and extrapolating using signal strength models, TIX [10] achieves an accuracy of 5.4 meters indoors. Lim *et al.* [14] use a singular value decomposition method combined with RF propagation models to create a signal strength map (overlapping with map-based approaches). They achieve a localization error of about three meters indoors. EZ [6] is a

system that uses sporadic GPS fixes on mobiles to bootstrap the localization of many clients indoors. EZ solves these constraints using a genetic algorithm, resulting in a median localization error of two meters indoors, without the need for any explicit pre-deployment calibration.

AoA-based approaches. Niculescu and Nath [17] use a mechanically-rotated directional antenna to triangulate clients’ locations from packet-level RSS readings as base stations rotate their antennas. Their system achieves a 2.1 m median error with seven participating base stations. However, it requires an additional rotating antenna to be added to the base station, and needs to overhear hundreds of packets from each client in order to get enough RSS data to achieve that accuracy.

Wong *et al.* [31] investigate the use of AoA and channel impulse response measurements for localization. While they have demonstrated positive results at a very high SNR (60 dB), typical wireless LANs operate at significantly lower SNRs, and the authors stop short of describing a complete system design of how the ideas would integrate with a functioning wireless LAN as ArrayTrack does. Niculescu *et al.* [16] simulate AoA-based localization in an ad hoc mesh network. AoA has also been proposed in CDMA mobile cellular systems [34], in particular as a hybrid approach between TDoA and AoA [8, 32], and also in concert with interference cancellation and ToA [27].

Geo-fencing [24] utilizes directional antennas and a frame coding approach to control the indoor coverage boundary. Compared to Geo-fencing, ArrayTrack provides a location service which does not impact the flow of traffic in the wireless network.

Image processing based approaches. These approaches match features extracted from images from a mobile’s camera to localize a device. Examples include work by Hile *et al.* [12] and vSLAM [13]. The approach is computationally demanding, with a backend server performing a significant amount of image processing, but accuracy is high, typically on the order of 10–30 cm. In general, however, visual occlusion or poor lighting conditions pose challenges to image processing based approaches, decreasing accuracy.

A reference table comparing several selected localization techniques was included in [7] by Jaewoo Chung *et al.*

6. CONCLUSIONS AND FUTURE WORK

We have presented ArrayTrack, an indoor location system that uses angle-of-arrival techniques to locate wireless clients indoors in a wireless local area network. ArrayTrack combines best of breed algorithms for AoA based direction estimation and spatial smoothing with novel algorithms for suppressing non-line of sight reflections that occur frequently indoors. ArrayTrack achieves median 25 cm location accuracy when clients are stationary indoors.

While these initial results are promising, we believe that ArrayTrack's accuracy can be increased substantially. In future work, we plan to investigate novel angle-of-arrival estimation algorithms that leverage the fact that the multipath components arriving at an AP are phase-coherent with each other. We will investigate algorithms that weight APs whose line of antennas faces a client less in the combination step. And finally, we will investigate hill-climbing algorithms that can refine the location search quicker than fine-grained sampling. We speculate that realization of the latter two techniques alone might come close to the omniscient scheme presented in the evaluation.

We also plan to investigate the problem of real-time mobile tracking using Bayesian techniques [9], which we believe is ideally matched to ArrayTrack's low processing overhead and extremely small system latency. 3D tracking is also one of our future interests which will require the use of a two dimensional array for localization.

7. REFERENCES

- [1] Xirrus. <http://www.xirrus.com/Products/Wireless-Arrays>.
- [2] E. Aryafar, N. Anand, T. Salonidis, and E. Knightly. Design and experimental evaluation of multi-user beamforming in wireless LANs. In *Proc. of ACM MobiCom*, 2010.
- [3] P. Bahl, J. Padhye, L. Ravindranath, M. Singh, A. Wolman, and B. Zill. DAIR: A framework for managing enterprise wireless networks using desktop infrastructure. In *Proc. of ACM HotNets*, 2005.
- [4] P. Bahl and V. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *Proc. of IEEE Infocom*, pages 775–784, 2000.
- [5] P. Bahl, V. Padmanabhan, and A. Balachandran. Enhancements to the RADAR location tracking system. Technical Report MSR-TR-2000-12, Microsoft Research, Feb. 2000.
- [6] K. Chintalapudi, A. Iyer, and V. Padmanabhan. Indoor localization without the pain. In *Proc. of ACM MobiCom*, 2010.
- [7] J. Chung, M. Donahoe, C. Schmandt, I. Kim, P. Razavai, and M. Wiseman. Indoor location sensing using geo-magnetism. In *Proceedings of ACM MOBISYS*, 2011.
- [8] L. Cong and W. Zhuang. Hybrid TDoA/AoA mobile user location for wideband CDMA cellular systems. *IEEE Trans. on Wireless Communications*, 1(3):439–447, 2002.
- [9] D. Fox, J. Hightower, L. Liao, D. Schulz, and G. Borriello. Bayesian filtering for location estimation. *IEEE Pervasive*, 2(3), 2003.
- [10] Y. Gwon and R. Jain. Error characteristics and calibration-free techniques for wireless LAN-based location estimation. In *ACM MobiWac*, 2004.
- [11] A. Haeberlen, E. Flannery, A. Ladd, A. Rudys, D. Wallach, and L. Kavraki. Practical robust localization over large-scale 802.11 wireless networks. In *Proc. of ACM MobiCom*, 2004.
- [12] H. Hile and G. Borriello. Positioning and orientation in indoor environments using camera phones. *IEEE Computer Graphics and Applications*, 2008.
- [13] N. Karlsson, E. Di Bernardo, J. Ostrowski, L. Gonclaves, P. Pirjanian, and M. Munich. The vSLAM algorithm for robust localization and mapping. In *Proc. of the IEEE ICRA*, 2005.
- [14] H. Lim, C. Kung, J. Hou, and H. Luo. Zero configuration robust indoor localization: Theory and experimentation. In *Proc. of IEEE Infocom*, 2006.
- [15] R. D. Maesschalck, D. Jouan-Rimbaud, and D. Massart. The Mahalanobis distance. 2000.
- [16] D. Niculescu and B. Nath. Ad-hoc positioning system (APS) using AoA. In *Proc. of IEEE Infocom*, 2003.
- [17] D. Niculescu and B. Nath. VOR base stations for indoor 802.11 positioning. In *Proc. of ACM MobiCom*, 2004.
- [18] N. Priyantha, A. Chakraborty, and H. Balakrishnan. The Cricket location-support system. In *Proc. of the ACM MobiCom Conf.*, pages 32–43, Aug. 2000.
- [19] T. S. Rappaport. *Wireless Communications: Principles and Practice*. Prentice-Hall, 2nd edition, 2002.
- [20] T. Roos, P. Myllymaki, and H. Tirri. A probabilistic approach to WLAN user location estimation. *International J. of Wireless Information Networks*, 9(3), 2002.
- [21] T. M. Schmidl and D. C. Cox. Robust Frequency and Timing Synchronization for OFDM. *IEEE Trans. on Communications*, 45(12):1613–1621, Dec. 1997.
- [22] R. Schmidt. Multiple emitter location and signal parameter estimation. *IEEE Trans. on Antennas and Propagation*, AP-34(3):276–280, Mar. 1986.
- [23] T.-J. Shan, M. Wax, and T. Kailath. On spatial smoothing for direction-of-arrival estimation of coherent signals. *IEEE Transactions on Acoustics, Speech and Signal Processing*, (4), 1985.
- [24] A. Sheth, S. Seshan, and D. Wetherall. Geo-fencing: Confining Wi-Fi Coverage to Physical Boundaries. In *Proceedings of the 7th International Conference on Pervasive Computing*, 2009.
- [25] A. Smailagic, D. Siewiorek, J. Anhalt, D. Kogan, and Y. Wang. Location sensing and privacy in a context aware computing environment. In *Pervasive Computing*, 2001.
- [26] K. Tan, H. Liu, J. Fang, W. Wang, J. Zhang, M. Chen, and G. Voelker. SAM: Enabling practical spatial multiple access in wireless LAN. In *Proc. of ACM MobiCom*, 2009.
- [27] A. Tarighat, N. Khajehnouri, and A. Sayed. Improved wireless location accuracy using antenna arrays and interference cancellation. 4, 2003.
- [28] A. Varshavsky, E. Lara, J. Hightower, A. LaMarca, and V. Otsason. GSM indoor localization. In *Pervasive and Mobile Computing*, 2007.
- [29] R. Want, A. Hopper, V. Falcao, and J. Gibbons. The active badge location system. *ACM Trans. on Information Systems*, 10(1):91–102, Jan. 1992.
- [30] A. Ward, A. Jones, and A. Hopper. A new location technique for the active office. *IEEE Personal Communications*, 4(5):42–47, 1997.
- [31] C. Wong, R. Klukas, and G. Messier. Using WLAN infrastructure for angle-of-arrival indoor user location. In *Proc. of the IEEE VTC Conf.*, pages 1–5, Sept. 2008.
- [32] Y. Xie, Y. Wang, P. Zhu, and X. You. Grid-search-based hybrid ToA/AoA location techniques for NLOS environments. *IEEE Comms. Letters*, 13(4):254–256, 2009.
- [33] J. Xiong and K. Jamieson. Secureangle: Improving wireless security using angle-of-arrival signatures. In *Proc. of ACM HotNets*, 2010.
- [34] L. Xiong. A selective model to suppress nlos signals in angle-of-arrival (AoA) location estimation. In *Proc. of the IEEE PIMRC*, 1998.
- [35] M. Youssef and A. Agrawala. The Horus WLAN location determination system. In *Proc. of ACM MobiSys*, 2005.