

Determining RF Angle of Arrival Using COTS Antenna Arrays: A Field Evaluation

Hsieh-Chung Chen, Tsung-Han Lin, H. T. Kung, Chit-Kwan Lin, Youngjune Gwon
School of Engineering and Applied Sciences
Harvard University

Abstract—We are interested in estimating the angle of arrival of an RF signal by using commercial-off-the-shelf (COTS) software-defined radios (SDRs). The proposed COTS-based approach has the advantages of flexibility, low cost and ease of deployment, but—unlike traditional phased antenna arrays in which elements are already phase-aligned—we face the challenge of aligning individual SDRs during field deployment in order to ensure coherent phase detection. We propose a strategy to relax the requirement of tight phase synchronization between distributed oscillators by using a novel phase difference of arrival mechanism based on a field-deployable reference transmitter. This approach enables flexible and inexpensive COTS phased-array designs. We evaluate our method in an outdoor, 20m×20m open field and observe localization errors below 3m. We conclude that a COTS-based approach to RF source localization is amenable to rapid and low-cost deployment of sensing infrastructure and could potentially be of interest to the Intelligence, Surveillance and Reconnaissance (ISR) community at the tactical edge.

I. INTRODUCTION

Angle of arrival (AOA) estimation is fundamental to many wireless sensing and communications applications, especially RF source localization. Typically, AOA estimation is performed using an antenna array, in which the phase difference between the received signals at each antenna array element is mapped to the incident direction of the signal. This method generally gives two advantages. First, since the phase of the received signal is usually more stable than the received signal strength (RSS), AOA estimation can achieve higher accuracy than RSS-based localization approaches. Second, given an effective AOA estimation scheme, two antenna arrays suffice to achieve accurate target localization, while range-based approaches [1][2] require three or more sensor nodes.

However, antenna arrays are generally expensive and complex to build, since tight coordination among antenna elements is needed to achieve coherent phase detection. Furthermore, antenna arrays usually have fixed element configurations, meaning they are difficult to adapt to changing application needs in the field. In this work, we show that AOA estimation can be implemented using modular, commercial-off-the-shelf (COTS) software-defined radios (SDRs). This allows us to reap the benefits of flexibility, low cost and ease of deployment while still providing reasonable localization performance. We emphasize that the modular nature of such COTS components is key to enabling a flexible antenna array design; in particular, the modular COTS components we employ are readily

available Universal Software Radio Peripherals (USRPs) [3].

To our knowledge, no prior work takes the modular COTS approach to estimate AOA, in large part because the coordination problem has proven to be difficult to solve with COTS components such as USRPs. In particular, each modular COTS component in an antenna array assembly has an independent local oscillator with a phase offset that is different from and unknown to the other components. This can lead to incoherence in phase detection across antennas, which in turn prohibits phase difference estimation, and thus AOA estimation. The main contribution of our paper is a *phase difference of arrival* (PDOA) mechanism that allows us to relax the stringent requirement of coordination amongst array elements by using a reference transmitter to provide a common phase reference for all receive antennas. As a result, antenna modules within a COTS-based array can operate individually with their own local oscillators, sidestepping the need for complex hardware design or tight margins.

To evaluate our approach, we constructed an outdoor testbed using USRPs and performed localization field experiments in an 20m × 20m open field. In most of our experiments, we found the localization error (i.e., the Euclidean distance from the ground truth to the estimated locations) to be below 3m. This suggests that our COTS-based system can localize a target with high accuracy, without suffering from the inflexibility or design complexity of a typical antenna array system.

The rest of this paper is organized as follows. In Section II, we discuss the advantages of our modular COTS approach. Section III briefly reviews the theory behind AOA estimation based on PDOA measurements at multiple antennas. In Section IV, we present our design and implementation of the COTS antenna array system for measuring PDOA, and describe the field experiments we performed with our system. Section V evaluates the localization performance of our system and presents a comparative analysis of our scheme against a RSS-based method. In Section VI, we discuss related work, and Section VII concludes the paper.

II. ADVANTAGES OF THE MODULAR COTS APPROACH IN ESTIMATING AOA

Using COTS equipment—especially software-defined radios—to build antenna arrays has three major advantages:

- 1) **Flexible Configuration.** Using modular COTS SDRs to construct a antenna array allows for flexible configura-

tion. In contrast to traditional antenna arrays, which are typically tuned to a fixed, narrow frequency band, a SDR array gives us wide latitude to choose the frequencies over which to operate on a much larger frequency range (e.g. the WBX daughterboard from Ettus Research that we use operates over a range of 50MHz–2.2GHz). The modular nature of the COTS SDR allows for the separation distance between array elements to be easily changed as needed, based on the frequency band of interest. Furthermore, the modularity of each SDR unit allows us to install as many or as few array elements, in any array shape, as needed. For example, in situations demanding higher resolution, more elements can be installed, and vice versa.

- 2) **Low Cost.** The cost of COTS components to build an antenna array is relatively low compared to that of traditional fixed-band antenna arrays. The total cost of one of our COTS antenna arrays with three SDR-based antenna elements was less than \$5,000.
- 3) **Easy, Ad Hoc Deployment.** COTS components normally have high enough manufacturing quality and are small enough that they can be easily deployed in an ad hoc fashion. In our field experiments (see Section IV), deployment was surprisingly easy: the entire testbed took about an hour to set up (this included mounting the antenna arrays onto and erecting the towers). Also, because there were few faulty equipment issues, we found that we could easily manage a modestly-sized testbed with a simple command and control system based on general-purpose computers.

These advantages suggest that COTS antenna arrays may be suited for tactical edge surveillance applications that are of interest to the ISR community. Furthermore, we conjecture that our SDR implementation can be naturally migrated to a sensor array platform that instead uses cellular phones as both antenna elements and signal processors. Cell phones are ubiquitous and inconspicuous, making them an ideal platform for building ad hoc antenna arrays for ISR applications.

III. PHASE DIFFERENCE OF ARRIVAL

Estimating the AOA of an RF signal by an antenna array relies on detecting the signal’s phase when it arrives at multiple antenna elements. Due to the difference in propagation distances from the signal source to individual receive antennas, each antenna observes a different phase shift of the signal. For example, as shown in Figure 1, if the signal from transmitter A are assumed to propagate in parallel through space, then the phase observed by the two receive antennas, Φ_{A1} and Φ_{A2} , can be represented as a function of the angle of incidence θ and the distance separating the antennas d :

$$\Phi_{A1} - \Phi_{A2} = \frac{2\pi d \sin \theta}{\lambda} \quad (1)$$

where λ is the wavelength. Therefore, one only needs to know the *phase difference* in the antenna array to estimate the incidence angle.




Fig. 1: The angle of arrival θ is a function of the measured phase difference (red) and antenna separation distance d . The diagram depicts the ideal scenario of parallel rays from a target transmitter.

However, Equation (1) requires coherent phase detection by the antenna elements in order to compute the difference, meaning that the antenna elements need to be perfectly synchronized in both phase and clock rate. To ensure this, traditional antenna arrays usually are built on a single platform and with multiple antenna elements connecting to the same clock and oscillator. However, a COTS software-defined radio is generally equipped with only one receive antenna; thus, an antenna array is assembled by using multiple independent SDRs. This presents a problem which we experienced in our field deployments: while we can use an external reference clock to distribute a reference signal to synchronize the radios in clock rate, their individual local oscillators still have an unknown initial phase offset when down-converting the RF signal to baseband.

Fortunately, this phase offset is relatively stable over time, and we can use a phase difference of arrival mechanism (PDOA) to eliminate the effect of the unknown offsets. The basic idea behind this mechanism is to exploit a short reference signal sent from an additional reference transmitter at a known location. By taking the difference in *phase* between the target and the reference signal, the initial phase offsets can be eliminated.

To see how this PDOA mechanism works, consider a scenario in which two COTS components, each with one antenna, are used to estimate the incident angle of the signal from a target transmitter A . Equation (1) shows that AOA estimation depends only on the phase difference between the signals received at two antennas. Denoting Φ'_{A1} and Φ'_{A2} as the phase of A ’s signal measured at the two receive antennas, the measured phase difference can be written as:

$$\Phi'_{A1} - \Phi'_{A2} = (\Phi_{A1} + \gamma_1) - (\Phi_{A2} + \gamma_2) \quad (2)$$

where Φ_{A1} and Φ_{A2} are the true signal phase, and γ_1 and γ_2 are the initial phase offsets of the local oscillators in the two COTS antenna components.

It may seem that, without knowing individual initial phase offsets, the true phase difference $\Phi_{A1} - \Phi_{A2}$ cannot be derived. However, in Equation (2) we only need to know the difference between initial offsets $\gamma_1 - \gamma_2$, and this difference is stable over a sustained period of time. We thus can add an additional reference transmitter B with a known location

to send out a reference signal periodically and measure this offset difference. Similar to Equation (2), the measured phase difference of B 's signal can be written as:

$$\Phi'_{B1} - \Phi'_{B2} = (\Phi_{B1} + \gamma_1) - (\Phi_{B2} + \gamma_2) \quad (3)$$

Using Equation (2) and (3), we then can eliminate $\gamma_1 - \gamma_2$ and estimate the true phase difference by the following:

$$\Phi_{A1} - \Phi_{A2} = (\Phi'_{A1} - \Phi'_{A2}) - (\Phi'_{B1} - \Phi'_{B2}) + (\Phi_{B1} - \Phi_{B2}) \quad (4)$$

Equation (4) means that to obtain the correct phase difference estimate, one simply corrects the measured phase difference of the target signal by the measured phase difference of the reference signal. Note that the true phase difference $\Phi_{B1} - \Phi_{B2}$ of the reference signal needs to be known a priori, and by Equation (1) it can be computed given B 's location. Thus, the desired signal phase difference can be derived without measuring individual initial offsets γ_1 and γ_2 of the local oscillators.

Finally, several points about the PDOA mechanism are worth noting. First, the mechanism is independent from the size of the antenna array because the AOA estimation is built on the signal phase difference between pairs of antennas. Thus, the mechanism can be easily generalized to an array of arbitrary size. Second, there is no need to place the reference transmitter at a particular location or to have a sophisticated waveform design for the reference signal. Since the reference signal is solely for providing a phase reference, as long as the signal-to-noise ratio (SNR) of the reference signal is sufficiently high, the target signal's phase difference can always be correctly estimated. Third, if the initial phase offsets drift over time, the reference signal needs to be retransmitted periodically for re-calibration. Fortunately, we have not observed the drift to be serious in practice (it is sufficiently stable for at least one minute), meaning re-calibration can be infrequent. In the future, for even greater robustness, we could estimate the drift and compensate accordingly.

IV. FIELD EXPERIMENT WITH COTS EQUIPMENT

We validate our AOA estimation method through a simple field experiment. We first describe the equipment used for this experiment—stressing that all the components used are COTS—and then discuss our measurement methodology.

We constructed two COTS SDR antenna arrays (labeled R and G), each consisting of six components: three USRP N-200 software-defined radios (SDRs) manufactured by Ettus Research, Inc., each equipped with a 900MHz-band rubber duck omnidirectional antenna, an external clock module that provided a 10MHz synchronization signal to each directly-connected SDR, a standard desktop PC that hosted the SDR software and stored the measurement data, and a gigabit Ethernet hub that connected the SDRs to the PC. The SDRs and the external clock were installed inside a weatherproof and shockproof case (manufactured by Pelican Products, Inc.), and this entire package was mounted to a steel truss tower at a height of $\sim 6.1\text{m}$ (20ft). Figure 2 illustrates the design of




Fig. 2: A diagram of our COTS antenna array that uses three software-defined radios (SDRs). The photo inset shows one such antenna array consisting of three SDRs and a synchronization clock inside a weatherproof case. This entire package is mounted to the top of a 6m tower.

our COTS antenna array, mounted to the tower, with the photo inset depicting the arrangement of the SDRs and clock module inside the Pelican case. Note that even though each receiver had three SDRs, only two SDRs in each receiver were used during our experiments.

Our two transmitters (labeled A and B) were similarly constructed, except each node consisted of only a single USRP N-200 SDR directly connected to a host PC. These transmitters were also housed inside Pelican cases, but these were placed directly on the ground such that the antenna of the SDR was approximately 35cm from the ground.

Figure 3 shows the initial arrangement of transmitters A (the target) and B (the reference transmitter) with respect to the two COTS antenna arrays R and G . A is 15m from R and at a relative angle of 0° . B is also at 0° but is at the midpoint between A and R . Throughout our measurement campaign, B , R and G are fixed in their locations, but the location of A changes in 5° increments clockwise towards G (marked by black dots in Figure 3); thus, A travels along the blue arc in Figure 3, with the distance between A and R remaining fixed at 15m. At each location, R and G sample the channel as A and B take turns transmitting a signal at 916MHz. A measurement round at each location constitutes one experiment run; in total, we performed ten runs.

This setup provides us with the ground truth locations of A , against which we can compare the AOA derived from the measured signal phase. We discuss the performance of our method in the following section.

V. EXPERIMENTAL RESULTS

A. Maximum Likelihood AOA and Location Estimation

Instead of directly applying Equation (1) to estimate the incident direction of the target signal, for estimation robustness,




Fig. 3: Our field experiment setup. Transmitters (A and B) and antenna arrays (R and G) were placed as shown. Transmitter A was then moved by five-degree increments (locations marked by black dots). At each location, A and B took turns transmitting a signal while the antenna arrays R and G sampled the signal.

we take a maximum likelihood approach in which we overlay a grid onto the two-dimensional plane of the target and compute the likelihood that the target is at each grid location. In short, there is an expected phase difference at each location and by comparing this against the measured phase difference, we can evaluate the likelihood of the target being at each location. The location with the maximum likelihood would thus be our best estimate. Here, we focus on a single target since we only use two antennas on each receiver, but the system could be extended to handle multiple targets with more antennas [4].

We use $(\hat{\Phi}_{A1} - \hat{\Phi}_{A2})$ to denote the measured phase difference with some noise. The likelihood function for a potential target location x can be defined as:

$$\mathbf{L}(x) = p_x(\hat{\Phi}_{A1} - \hat{\Phi}_{A2}) \quad (5)$$

where $p_x(\Delta)$ is the probability of observing phase difference Δ at the antenna array when the target is at location x . Assuming Gaussian noise in the phase measurements,

$$p_x(\Delta) \sim \mathcal{N}(\Phi_{x1} - \Phi_{x2}, \sigma^2). \quad (6)$$

That is, the probability function follows a normal distribution with mean $(\Phi_{x1} - \Phi_{x2})$, the expected phase difference when the target is at x . We choose the variance σ^2 empirically by computing the variance of measurements $(\hat{\Phi}_{A1} - \hat{\Phi}_{A2})$ over a short period of time.

Figure 4 (left and center panels) shows two heat map plots of the likelihood computed from a representative experiment run using antenna arrays R and G , respectively, and where target A is located at 5° (see Figure 3). Since the locations on the same line of the incident angle share the same phase difference, they will have equal likelihood (the same color). Note that lines representing equal likelihood in Figure 4 are not exactly straight, but curve around the receive antenna array.

This is an artifact¹ of the three-dimensional geometry of our configuration, where the receive antennas are elevated at 6m and the target is at ground level. Were the receive antennas and the target both on the ground, the equal likelihood lines would be straight.

Next, we can combine the likelihood computed from the two antenna arrays for target localization. The joint likelihood can be calculated from

$$\mathbf{L}_{(G,R)}(x) = \mathbf{L}_G(x)\mathbf{L}_R(x), \quad (7)$$

and the resulting values are shown as a heat map plot in Figure 4 (right panel). From this, we can choose the location with the maximum likelihood as the location estimate for the target.

B. Localization Accuracy

The results of the ten experiment runs in localizing target A are summarized in Table I. From this, we make the following observations. First, most of the runs have a localization error (i.e., the Euclidean distance between the ground truth and estimated locations) below 3m (except for Runs 8 and 9), which is relatively accurate even when compared to a RSS-based localization scheme [5] that can address incorrect, outlier measurements. This accuracy is achieved despite the fact that we relied completely on geometry and did not include any environment or hardware-specific calibration. Second, for the less accurate Runs 8 and 9, most of the error is attributed to inaccurate phase measurement at R (note that R is farther from the target than G). As the phase measurement is stable for the two runs, we suspect that the inaccuracy comes from omitting the effects of uneven ground; this effect can be mitigated with additional calibration.

C. Comparison with RSS-based methods

A popular approach for RF localization is to use RSS as an indicator of distance [1][2], since the relationship between RSS and distance can be modeled by a path loss model. However, RSS-based methods have two major drawbacks. First, the path loss model often requires careful calibration in order to tune environment-dependent model parameters; such calibration can be difficult and time-consuming. Second, when the distance between the transmitter and the receiver increases, a small difference in RSS may correspond to a large change in distance due to decreased SNR, resulting in lower localization resolution.

We use the signal amplitude measured at R and G (shown in Figure 5) to extrapolate how a RSS-based method may perform in our experiment configuration. In particular, we observe that the decay of RSS in the figure is slower than $1/d^2$ (where d is distance) predicted by a free space path-loss model (see, e.g., [6]), which implies that the measured RSS is mostly

¹We calculate the likelihood with respect to the plane defined by a potential target location and the two antenna element locations and then project these values onto two dimensions as a heat map. This induces a curve in the projection.

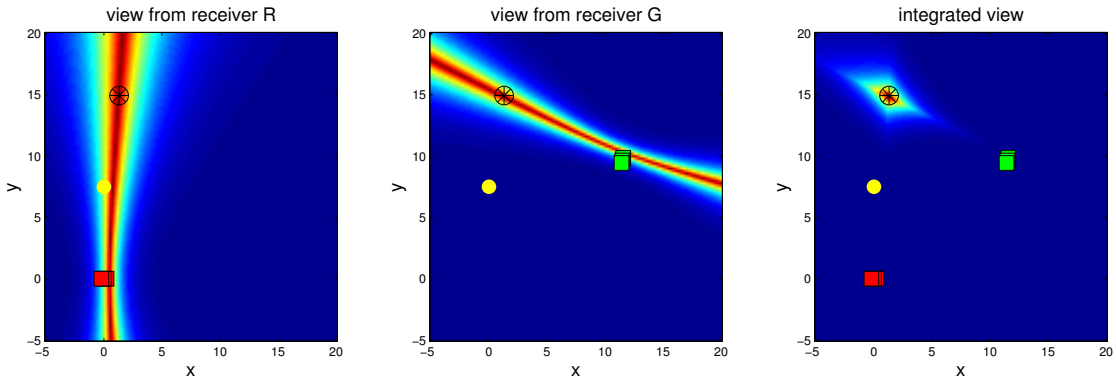


Fig. 4: Localization with two receivers R (red) and G (green). The reference transmitter B is marked in yellow and the target in black. This figure shows the results from Run 3, where the target is at 5° .

Run	1	2	3	4	5	6	7	8	9	10
Angle (degree)	0	0	5	10	15	20	25	30	35	40
Target to R distance (m)	16.15	16.15	16.15	16.15	16.15	16.15	16.15	16.15	16.15	16.15
Target to G distance (m)	14.02	14.02	12.95	11.88	10.83	9.81	8.84	7.94	7.16	6.54
Localization Error (m)	2.03	0.82	0.17	2.43	2.00	1.77	1.23	4.52	3.39	1.63
R Phase Error (%)	8.36	3.24	0.03	10.58	0.62	4.83	6.13	18.54	12.57	0.10
G Phase Error (%)	5.52	0.63	0.49	6.98	13.18	10.14	8.65	5.32	4.22	1.65
R measured phase variance (degree)	2.37	4.01	4.02	2.65	5.43	4.64	25.41	7.40	5.28	17.91
G measured phase variance (degree)	2.66	2.38	3.17	2.66	5.58	3.09	2.56	2.26	2.29	2.22

TABLE I: Localization performance. The first row is the target’s angle from R ’s perspective. Localization error is the distance between estimated target location and the ground truth location.




Fig. 5: Received signal strength (with variance) from transmitter A placed at different distances from receiver G , but at the same distance from receiver R (see Figure 3).

located within the flat tail region of the RSS-distance model. In this region, the RSS at different distances becomes difficult to distinguish and is more sensitive to noise, which would likely lead to poor localization performance. This can be seen in the figure as the signal amplitude only varies slightly when the distance is over 8m (and in our experiments, up to 16m). This means that errors as large as 8m would be inevitable.

VI. RELATED WORK

The use of antenna arrays for localization has already enjoyed significant attention in the literature. Blanco *et al.* [7]

examined the effectiveness of directionality in radio spatial reuse and localization for indoor environments. Niculescu *et al.* [8] used directional antennas for location estimation, while Sayraan-Pour and Kaspar [9] computed directional estimates with beamforming antennas. In contrast, our AOA estimation does not require beamforming or depend on antenna directionality but exploits signal phase difference. In addition, our system is passive, unlike phased array radar systems [10], which rely on active sensing elements.

In practical AOA implementation, SecureAngle [11] implements antenna arrays using the WARP platform, and takes a similar reference transmitter approach to resolve the initial phase offset problem. In that design, the reference transmitter is co-located with the antenna array and explicitly wired to the antenna array. A switch is thus required to connect all antennas to control the wire and wireless signal paths. In contrast, our COTS-based approach does not require additional wiring or hardware for the antenna array, and can easily scale in the size of the array.

Maróti *et al.* [12] have presented an implementation that also uses signal phase measurements for localization but instead based on radio interferometry. It requires a reference transmitter to transmit simultaneously with the target, but in a slightly different frequency to create a low frequency

envelope. The receiver then can measure the phase of the low frequency envelope and thus avoid tight phase and time synchronization among its antennas. However, this method not only requires special waveforms and additional controls in transmission timing, but also suffers from the potential inaccuracy of inferring the phase of a low frequency envelope from signal amplitudes.

The use of reference transmitters to avoid explicit synchronization can be found in time difference of arrival (TDOA) localization schemes. BeepBeep [13] employs a reference transmitter to resolve the difficulty of precise time synchronization amongst receivers. Our approach instead exploits the reference transmitter to sidestep the need for explicit phase synchronization.

Lastly, there are localization techniques that are based on received signal strength. As mentioned earlier, a problem for RSS-based localization is ranging (distance) estimation error, especially when SNR is low. It is a well-known fact that RSS ranging errors often do not closely follow a well-behaved distribution (e.g., Gaussian), implying that conventional least squares optimization schemes or model fitting will not be effective to mitigate ranging errors [14]. Methods such as SISR [5], SDP [15] and MDS-MAP [16] attempt to mitigate the negative effects of outliers in RSS ranging. In contrast, we avoid the issue of noisy RSS ranging by instead relying on AOA estimation, which is more stable, as we argued in Section V-C. Furthermore, we perform AOA estimation using COTS-based antenna arrays; to our knowledge, no prior work has taken this approach.

VII. CONCLUSIONS AND FUTURE WORK

In this paper we have shown that the modular COTS approach to building array antennas can provide important benefits including flexible configuration, low cost, and easy ad hoc deployment. We have taken this approach to implement an antenna array, using modular software-defined radios and demonstrated through field experiments that the approach supports rapid deployment of a practical system for accurately determining the AOA of RF signals. To our knowledge, no prior related work uses COTS components, perhaps because the reference transmitter approach had not been considered as a solution to the antenna coordination problem.

We anticipate that our experiment scenario will naturally evolve into an even more sophisticated one involving an unmanned aerial vehicle (UAV). In this new scenario, the UAV is expected to carry a COTS antenna array for localizing stationary signal sources on the ground transmitting on various frequency bands. Our current testbed setup can thus be considered as an approximation of the UAV scenario, since each tower approximates a way-point along the UAV flight path at which the airborne antenna array samples the emitted signal. Of course, the real UAV case will be more challenging because—unlike the tower configuration, where the locations of the receiver and reference transmitter are known—the location of the airborne receiver could be difficult

to determine precisely during flight (GPS does not provide sufficient accuracy). In this case, an UAV may need to get its precise location by receiving coordinate information from localization anchor nodes on the ground. We hope to tackle in the near future such challenges by leveraging our expertise in UAV flight experiments [17].

ACKNOWLEDGMENTS

This material is in part based on research sponsored by Air Force Research Laboratory under agreement number FA8750-10-2-0180. The U.S. Government is authorized to reproduce and distributed reprints for Governmental purposes not withstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies of endorsements, either expressed or implied, of Air Force Research Laboratory or the U.S. Government.

REFERENCES

- [1] C. Savarese, J. M. Rabaey, and K. Langendoen, "Robust Positioning Algorithms for Distributed Ad-Hoc Wireless Sensor Networks," in *USENIX ATC*, 2002.
- [2] A. Savvides, C.-C. Han, and M. B. Strivastava, "Dynamic Fine-grained Localization in Ad-Hoc Networks of Sensors," in *MobiCom*, 2001.
- [3] Ettus Universal Software Radio Peripheral (USRP). [Online]. Available: <http://www.ettus.com/>
- [4] R. Schmidt, "Multiple emitter location and signal parameter estimation," *Antennas and Propagation, IEEE Transactions on*, vol. 34, no. 3, pp. 276 – 280, March 1986.
- [5] H. T. Kung, C.-K. Lin, T.-H. Lin, and D. Vlah, "Localization with Snap-inducing Shaped Residuals (SISR): Coping with Errors in Measurement," in *MobiCom*, 2009.
- [6] T. Rappaport, *Wireless communications: principles and practice*. Prentice Hall, 1996, vol. 2.
- [7] M. Blanco, R. Kokku, K. Ramachandran, S. Rangarajan, and K. Sundaresan, "On the Effectiveness of Switched Beam Antennas in Indoor Environments," in *PAM, Volume 4979 of Lecture Notes in Computer Science*, 2008, pp. 122–131.
- [8] D. Niculescu and B. Nath, "VOR Base Stations for Indoor 802.11 Positioning," in *MobiCom*, 2004.
- [9] K. Sayrafian-Pour and D. Kaspar, "Source-assisted Direction Estimation Inside Buildings," in *IEEE INFOCOM*, 2006.
- [10] C.-F. Kuo, T.-W. Kuo, and C. Chang, "Real-time Digital Signal Processing of Phased Array Radars," *IEEE Trans. on Parallel and Distributed Systems*, vol. 14, no. 5, pp. 433–446, May 2003.
- [11] J. Xiong and K. Jamieson, "SecureAngle: Improving Wireless Security Using Angle-of-arrival Information," in *ACM Hotnets*, 2010.
- [12] M. Maróti, P. Völgyesi, S. Dóra, B. Kusý, A. Nádas, A. Lédeczi, G. Balogh, and K. Molnár, "Radio Interferometric Geolocation," in *SenSys*, 2005.
- [13] C. Peng, G. Shen, Y. Zhang, Y. Li, and K. Tan, "BeepBeep: a High Accuracy Acoustic Ranging System Using COTS Mobile Devices," in *SenSys*, 2007.
- [14] K. Whitehouse, C. Karlof, A. Woo, F. Jiang, and D. Culler, "The Effects of Ranging Noise on Multihop Localization: an Empirical Study," in *IPSN*, 2005.
- [15] P. Biswas, T. C. Liang, C. K. Toh, Y. Ye, and T. C. Wang, "Semidefinite programming approaches for sensor network localization with noisy distance measurements," *IEEE Trans. Autom. Sci. Eng.*, vol. 3, no. 4, pp. 360–371, 2006. [Online]. Available: <http://dx.doi.org/10.1109/TASE.2006.877401>
- [16] Y. Shang, W. Ruml, Y. Zhang, and M. Fromherz, "Localization from connectivity in sensor networks," *IEEE Trans. Parallel Distrib. Syst.*, vol. 15, no. 11, pp. 961–974, 2004.
- [17] D. Hague, H. T. Kung, and B. W. Suter, "Field experimentation of cots-based uav networking," in *MILCOM*, 2006.