



AFRL-RI-RS-TR-2013-198

WIRELESS COMPUTING ARCHITECTURE III

HARVARD COLLEGE

SEPTEMBER 2013

FINAL TECHNICAL REPORT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09. This report is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RI-RS-TR-2013-198 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/ S /

MICHAEL LITTLE
Work Unit Manager

/ S /

MARK H. LINDERMAN
Technical Advisor, Computing
and Communications Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE*Form Approved*
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) SEPTEMBER 2013		2. REPORT TYPE FINAL TECHNICAL REPORT		3. DATES COVERED (From - To) MAY 2010 – MAR 2013	
4. TITLE AND SUBTITLE WIRELESS COMPUTING ARCHITECTURE III				5a. CONTRACT NUMBER FA8750-10-2-0180	
				5b. GRANT NUMBER N/A	
				5c. PROGRAM ELEMENT NUMBER N/A	
6. AUTHOR(S) H.T. Kung				5d. PROJECT NUMBER WCNA	
				5e. TASK NUMBER WC	
				5f. WORK UNIT NUMBER H3	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Harvard College Office of Sponsored Programs 1350 Mass Ave Suite 600 Cambridge, MA 02138-3846				8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/RITB 525 Brooks Road Rome NY 13441-4505				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RI	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-RI-RS-TR-2013-198	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited. This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT We have developed new theory and novel systems on wireless networking, computing and sensing architectures. Our research has led to a methodology for achieving high throughput ground-to-UAV data transport via parallel links; a model of performing collaborative compressive spectrum sensing in a UAV environment \; a distributed architecture testbed for spectrum sensing with compressive sensing \; a design for concurrent channel estimation in scalable multiuser MIMO networking; and novel networking protocols based on machine learning technology. In addition, we have conducted field evaluation of using COTS antenna arrays to determine RF angle of arrival.					
15. SUBJECT TERMS Network, Antenna Arrays, UAV networking, Angle of Arrival, Localization MIMO, Access Point, Channel State Information, Compressive Sensing					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 31	19a. NAME OF RESPONSIBLE PERSON MICHAEL LITTLE
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) N/A

TABLE OF CONTENTS

Section	Page
1.0 SUMMARY	1
2.0 INTRODUCTION.....	3
2.1 Concurrent Channel Estimation in Scalable Multiuser MIMO Networking.....	3
2.2 Use of COTS Antenna Arrays to Determine RF Angle of Arrival	3
2.3 Wireless Inference-based Notification (WIN) without Packet Decoding	4
3.0 METHODS, ASSUMPTIONS, AND PROCEDURES	5
3.1 Concurrent Channel Estimation in Scalable Multiuser MIMO Networking.....	5
3.2 Use of COTS Antenna Arrays to Determine RF Angle of Arrival	7
3.3 Wireless Inference-based Notification (WIN) without Packet Decoding	12
4.0 RESULTS AND DISCUSSION	16
4.1 Concurrent Channel Estimation in Scalable Multiuser MIMO Networking.....	16
4.2 Use of COTS Antenna Arrays to Determine RF Angle of Arrival	17
4.3 Wireless Inference-based Notification (WIN) without Packet Decoding	20
5.0 CONCLUSIONS.....	21
5.1 Concurrent Channel Estimation in Scalable Multiuser MIMO Networking.....	21
5.2 Use of COTS Antenna Arrays to Determine RF Angle of Arrival	21
5.3 Wireless Inference-based Notification (WIN) without Packet Decoding	22
6.0 REFERENCES.....	23
7.0 LIST OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS.....	25

LIST OF FIGURES

Figure	Page
Figure 1. Two access strategies for multiuser MIMO (MU-MIMO) networks. Shaded areas denote packet preambles. Staggered access means only partially parallelized data transmissions, resulting in low channel utilization. In contrast, concurrent access can realize MIMO capacity gain by fully parallelizing data transmission.	6
Figure 2. 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.	9
Figure 3. 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.	11
Figure 4. 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.....	12
Figure 5. Conventional approach vs. WIN. Time slots labeled by time are shown at the bottom. Solid bars denote normal (black) and event (red) packets transmitted at various time slots.....	14
Figure 6. When detecting packets in the first layer, we use max-pooling with a sliding-window to address variations in packet delay due to multipath.	15
Figure 7. The WIN receiver can receive notification from the sender at a distance beyond the packet decoding range. In contrast, a conventional receiver can receive notification only in the packet decoding range.....	15
Figure 8. MIMO decoding performance using CSI estimated from concurrent preambles in 4x4 MIMO. Taking 13 taps is sufficient for reconstructing accurate CSI. Using fewer taps results in degradation in decoding performance, especially when the signal SNR is high.	17
Figure 9. 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 degree	19
Figure 10. WIN exhibits approximately a 13dB gain over conventional approach under AWGN channel.	20

LIST OF TABLES

Table	Page
Table 1. 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.	19

1.0 SUMMARY

We have conducted research in theory and systems on wireless networking, computing and sensing architectures. The work included establishment of theoretical models, development of system testbeds based on wireless local area networks, as well as system experimentation in the lab and in the field. It has led to a number of results and findings, including:

- A methodology for achieving high throughput ground-to-UAV transport via parallel links;
- A model of performing collaborative compressive spectrum sensing in a UAV environment;
- A distributed architecture and testbed for spectrum sensing with compressive sensing;
- A design for concurrent channel estimation in scalable multiuser MIMO networking;
- A system prototype and field evaluation on the use of COTS antenna arrays to determine RF angle of arrival; and
- An energy-efficient wireless inference-based notification (WIN) protocol for sensor networking.

We have documented these results in a number of papers, including:

- “Achieving High Throughput Ground-to-UAV Transport via Parallel Links,” 20th IEEE International Conference on Computer Communication and Networks (ICCCN 2011), August 2011.
- “Collaborative Compressive Spectrum Sensing in a UAV Environment,” Military Communications Conference (MILCOM 2011), November 2011.
- “Compressive Sensing with Optimal Sparsifying Basis and Applications in Spectrum Sensing,” IEEE Global Telecommunications Conference (GLOBECOM 2012), December 2012.
- “Compressive Sensing Medium Access Control for Wireless LANs,” IEEE Global Telecommunications Conference (GLOBECOM 2012), December 2012.
- “Determining RF Angle of Arrival Using COTS Antenna Arrays: A Field Evaluation,” Military Communications Conference (MILCOM 2012), October 2012.
- “A Chip Architecture for Compressive Sensing Based Detection of IC Trojans,” IEEE Workshop on Signal Processing Systems (SiPS 2012), October 2012.
- “Output Compression for IC Fault Detection Using Compressive Sensing,” Military Communications Conference (MILCOM 2012), October 2012.

- “Performance Gains in Conjugate Gradient Computation with Linearly Connected GPU Multiprocessors,” 4th USENIX Workshop on Hot Topics in Parallelism (HotPar ’12), Poster Session, June 2012.
- “Parallelization Primitives for Dynamic Sparse Computations,” 5th USENIX Workshop on Hot Topics in Parallelism (HotPar ’13), June 2013.
- “Scaling Network-based Spectrum Analyzer with Constant Communication Cost,” 32nd IEEE International Conference on Computer Communications (INFOCOM 2013), April 2013.
- “Concurrent Channel Access and Estimation for Scalable Multiuser MIMO Networking,” 32nd IEEE International Conference on Computer Communications (INFOCOM 2013) Mini-Conference, April 2013.
- “Wireless Inference-based Notification (WIN) without Packet Decoding,” 10th International Conference on Autonomic Computing (ICAC ’13), June 2013.

2.0 INTRODUCTION

This project had the goal of providing on-demand computing resources over wireless computing, communication and sensing infrastructures. We aimed to develop system architectures for the tactical edge that integrate networking, sensing and computing.

Our research has led to a methodology for achieving high throughput ground-to-UAV data transport via parallel links; a model of performing collaborative compressive spectrum sensing in a UAV environment; a distributed architecture testbed for spectrum sensing with compressive sensing; a design for concurrent channel estimation in scalable multiuser MIMO networking; and novel networking protocols based on machine learning technology. In addition, we have conducted field evaluation of using COTS antenna arrays to determine RF angle of arrival.

In the rest of this report, we highlight three of these results.

2.1 Concurrent Channel Estimation in Scalable Multiuser MIMO Networking

We have designed a scheme called MIMO/CON (“MIMO with concurrent channel access and estimation”), which is a PHY/MAC cross-layer design delivering throughput scalable to many users for multiuser MIMO wireless networking. By allowing concurrent launches of multiple data transmissions from multiple users, MIMO/CON can fully realize the capacity gain of a multi-antenna MIMO system. Using compressive sensing, MIMO/CON simultaneously estimates channel state information (CSI) of multiple channels from concurrently received preambles. Furthermore, MIMO/CON can boost channel utilization by allowing concurrent transmissions to exceed receive antennas momentarily. MIMO/CON has been implemented and evaluated on a lab testbed with software-defined radios. Further, simulation results suggest that MIMO/CON can achieve an improvement by up to 210% in MAC throughput over existing staggered access protocols in a 4×4 MIMO scenario.

2.2 Use of COTS Antenna Arrays to Determine RF Angle of Arrival

We have developed a scheme for 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 have built a prototype system and evaluated it in an outdoor, 20m×20m open field and observed 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.

2.3 Wireless Inference-based Notification (WIN) without Packet Decoding

We have designed an ultra-energy-efficient wireless protocol for transmitting notification messages in sensor networks. We argue that the usual practice where a receiver decodes packets sent by a remote node to acquire its state or message is suboptimal in energy use. We propose an alternative approach where a receiver first (1) performs physical-layer matched filtering on arrived packets without actually decoding them at the link layer or higher layer, and then (2) based on the matching results infers the sender's state or message from the time-series pattern of packet arrivals. We show that hierarchical multi-layer inference can be effective for this purpose in coping with channel noise. Because packets are not required to be decodable by the receiver, the sender can reach a farther receiver without increasing the transmit power or, equivalently, a receiver at the same distance with a lower transmit power. We call our scheme Wireless Inference-based Notification (WIN) without Packet Decoding. We demonstrate by analysis and simulation that WIN allows a sender to multiply its notification distance. We show how senders can realize these energy-efficiency benefits with unchanged systems and protocols; only receivers, which normally are larger systems than senders and have ample computing and power resources, need to perform WIN-related processing.

3.0 METHODS, ASSUMPTIONS, AND PROCEDURES

In this section, we describe assumptions and methodologies we used in the three highlighted areas of work in this report.

3.1 Concurrent Channel Estimation in Scalable Multiuser MIMO Networking

MIMO technologies enable an opportunity of linear increase in wireless channel capacity from the additional degrees-of-freedom created by multiple antennas. However, in single-user MIMO, the capacity gain is limited by the relatively small diversity offered by transmit antennas co-located on the same user platform. Multiuser MIMO (MU-MIMO) [5] removes this limitation with geographically separated users, resulting in rich spatial diversity. This allows further boosting of channel capacity.

We consider a MU-MIMO scenario where an access point (AP) is equipped with many antennas and every user possesses one antenna. We focus on the uplink case where multiple indoor users (i.e., “senders”) concurrently transmit data to a multi-antenna AP. With MU-MIMO, one would expect a throughput speedup factor of K with K receive antennas on the AP given sufficient spatial diversity; however realized throughput in realworld systems can be substantially less due to the difficulty of fully parallelizing channel access. Fundamentally, concurrent transmissions must be coordinated for proper MIMO decoding, but not hampered by inefficiency in access control.

Existing MU-MIMO systems (e.g., [7] , [10]) allow random access, but suffer from staggered data transmissions in order to avoid preamble collisions. In particular, the resulting loss of efficiency increases with the number of senders. For example, consider 1500-byte packets transmitted with 39Mbps data rate. Note that each packet transmission spans $300\mu\text{s}$. With an average access delay of $100\mu\text{s}$ [7] , there can be no more than 3 concurrent transmissions. Further, they are only partially parallelized as depicted in Figure 1(a). One may use frame aggregation [12] to send longer payload and amortize the access overhead. However, frame aggregation is not practical for delay of sensitive traffic such as VoIP or HTTP. Such protocols are thus not scalable to large K .

We argue that a more efficient approach to coordinate distributed senders is to launch multiple data transmissions concurrently thereby allowing the transmissions fully parallelized (shown in

Figure 1(b)). We call this access strategy *concurrent access*.

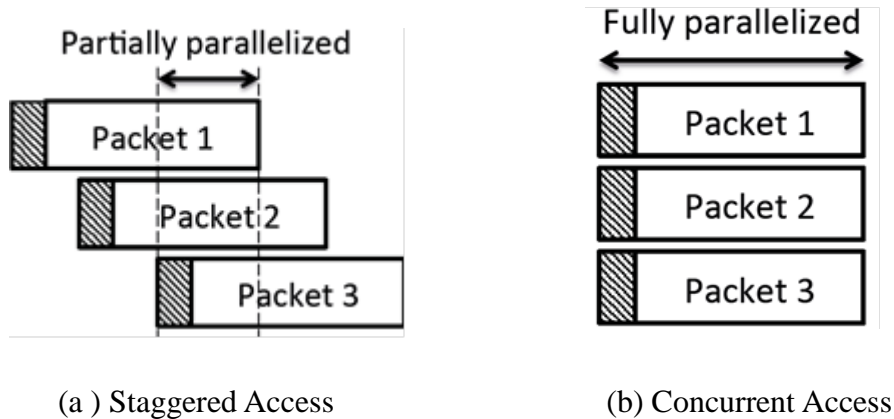


Figure 1. Two access strategies for multiuser MIMO (MU-MIMO) networks. Shaded areas denote packet preambles. Staggered access means only partially parallelized data transmissions, resulting in low channel utilization. In contrast, concurrent access can realize MIMO capacity gain by fully parallelizing data transmission.

Our design, MIMO/CON, supports concurrent access with the following two features:

- MIMO/CON can obtain accurate channel state information (CSI) from relatively short *concurrent preambles*. Therefore, senders do not need to stagger their transmissions, and consequently eliminate inefficiency due to sequential channel access. We show that the derived channel estimates achieve similar MIMO decoding performance as interference-free, serially transmitted preambles, and that the total preamble transmission time is near the minimum required for channel estimation.
- MIMO/CON can boost channel utilization by allowing concurrent transmissions to exceed the number of receive antennas momentarily, as long as this does not sustain over time. Therefore the senders can still randomly access the network (like in 802.11 DCF) and avoid being tightly scheduled for channel access. In short, MIMO/CON can achieve high and scalable MAC efficiency to take advantage of an increased number of receive antennas on the AP, and is amenable to the future trend of massive MIMO designs (e.g., 802.11ac suggests up to 8 antennas on the AP, and an unlimited number of antennas scenario is depicted in [9] for cellular networks).

A key observation that MIMO/CON exploits is that the CSI, i.e., the channel impulse response, is

expected to be *sparse* and constituted of only a few significant taps over the indoor environment we are interested in. This is due to the small delay spread of wireless signals relative to the OFDM symbol length. The received signal of concurrent preambles can be viewed as a linear combination of multiple sparse impulse responses from different senders. Estimating CSI then can be formulated as a sparse identification and recovery problem that finds all the spikes in different impulse responses. MIMO/CON leverages the recent advance in compressive sensing to tackle the sparse recovery problem: every sender uses a random code as its preamble sequence, and the AP treats the received concurrent preambles as random linear measurements of the unknown CSI. That is, under the compressive sensing setting, these random preambles form the sensing matrix. The measurement length only needs to be approximately proportional to the total number of nonzero taps. Since the number of active senders involved in a concurrent transmission is limited by the MIMO degrees-of-freedom, the measurement length can still be kept small under MU-MIMO.

Due to its ability in concurrent channel estimation, MIMO/CON can efficiently handle collisions when the number of active senders exceeds the number of AP antennas. Note that a compressive sensing solution to the sparse recovery problem yields the measured CSI vectors, and also identifies the associated senders. Suppose that a sender will retransmit the same packet if the packet is not acknowledged due to collision. When the retransmission is received, MIMO/CON can use it to recover not only the lost packet, but also can use the received CSI to decode other previously undecodable packets that would otherwise be discarded. We call this strategy *delay packet decoding*. As a result, MIMO/CON can tolerate demand fluctuations better and relax the access control by concerning only average use of the medium and realize statistical gains over random access.

3.2 Use of COTS Antenna Arrays to Determine RF Angle of Arrival

Angle of arrival (AOA) estimation is a capability fundamental to many wireless sensing and communications applications, and especially to RF source localization. Typically, AOA estimation is performed using an antenna array, in which the phase difference between the received signal at each antenna array element is mapped to the incident direction of the signal. This method generally gives two substantive 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, just *two* antenna arrays suffice to achieve accurate target localization, while range-based approaches (see, e.g., [17] 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 have fixed element configurations, meaning they are difficult to adapt to changing applications 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) to reap the benefits of flexibility, low cost and ease of deployment while still yielding reasonable localization performance. In particular, the COTS components we employ are the readily available Universal Software Radio Peripherals (USRPs) [19] manufactured by Ettus Research, Inc. We emphasize that the modular nature of such COTS components is key to enabling a flexible antenna array design.

To our knowledge, no prior work takes the 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. Thus, the main contribution of our work 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.

Phase Difference of Arrival. It is well-known that the angle of arrival of an RF signal can be estimated by an antenna array. Due to the difference in propagation distance from the signal source to individual receive antennas, each antenna will observe a different phase shift of the signal. For example, as shown in Figure 2, if the signal waves are assumed to propagate in parallel through space, then the phase observed by the two receive antennas, Φ_1^A and Φ_2^A , can be represented as a function of the angle of incidence θ and the distance separating the antennas d :

$$\Phi_{(12)}^A = \Phi_1^A - \Phi_2^A = \frac{2\pi d \sin\theta}{\lambda} \quad (1)$$

where λ is the wavelength.




Figure 2. 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.

Note that Equation (1) requires coherent phase detection by the antennas. 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 SDRs simultaneously. This presents a problem which we experienced in actual field deployments: while we can use an external reference clock to distribute a reference signal to synchronize the radios, their individual local oscillators still have an unknown initial phase offset when down-converting the RF signal to baseband.

Fortunately, experience with real-world environments often shows that this phase offset is relatively stable over time, and we can use a phase difference of arrival (PDOA) mechanism to eliminate the effect of unknown offsets. The basic idea behind this mechanism is to employ an additional reference transmitter at a known location to send a short reference signal. 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, let us assume a scenario with two receive antennas, a target transmitter A , and a reference transmitter B . Denoting $\hat{\Phi}_1^A$ and $\hat{\Phi}_2^A$ as A 's signal phase measured at the two receive antennas, the measured phase difference can be written as:

$$\hat{\Phi}_1^A - \hat{\Phi}_2^A = (\Phi_1^A + \gamma_1) - (\Phi_2^A + \gamma_2) \quad (2)$$

where Φ_1^A and Φ_2^A are the true signal phase, and γ_1 and γ_2 are the initial phase offsets on the two receive antennas.

Without knowing the initial phase offsets, the true phase difference $\Phi_{(12)}^A = \Phi_1^A - \Phi_2^A$ that we are interested in cannot be obtained. However, this problem can be sidestepped by using signals from a reference transmitter B . The measured phase difference for B can be written similarly, assuming the initial phase offset remains constant between measurements of A 's and B 's signals:

$$\widehat{\Phi}_1^B - \widehat{\Phi}_2^B = (\Phi_1^B + \gamma_1) - (\Phi_2^B + \gamma_2) \quad (3)$$

We then can estimate the true phase difference by the following equation:

$$\Phi_{(12)}^A = (\widehat{\Phi}_1^A - \widehat{\Phi}_2^A) - (\widehat{\Phi}_1^B - \widehat{\Phi}_2^B) + (\Phi_1^B - \Phi_2^B) \quad (4)$$

Note that in Equation (4), the true phase difference $\Phi_1^B - \Phi_2^B$ of the reference signal can be computed given B 's location. Thus the desired signal phase difference can now be measured without knowing the individual initial offsets of local oscillators.

Several points about the PDOA mechanism are worth noting. First, the mechanism is independent from the size of the antenna array because for AOA estimation, we are only interested in the signal phase difference between pairs of antennas. Second, there is no need to place the reference transmitter at a particular location or to have a sophisticated waveform design of 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 for it.

Field Experiment with COTS Equipment. We have validated 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 32 illustrates the design of 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.




Figure 3. 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 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 4 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 4); thus, *A*

travels along the blue arc in Figure 4, 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.




Figure 4. 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.

This setup provides us with the ground truth locations of A , against which we can compare the AOA derived from the measured signal phase.

3.3 Wireless Inference-based Notification (WIN) without Packet Decoding

We consider a common sensor network scenario where remote senders, such as sensors, transmit notifications about event detected as well as their operational conditions (e.g., device operating normally, and remaining battery power) to some designated receivers over wireless channels. In such a scenario, it is often desirable that nodes draw only a small amount of power in transmitting such notifications. This would allow transmitters to survive for a long time like years even operating on a small coin battery, in applications such as industrial monitoring and home automation.

Under a conventional approach (e.g., [20]), we will adopt a low-power wireless network, e.g., Bluetooth or ZigBee, to send notifications. A sender will periodically transmit normal packets to report that it is in a normal state, and start transmitting event packets when it enters an event state

upon noticing events of interest. A receiver will decode each received packet to determine if it is a normal or event packet, and in the latter case, may also examine packet payload to obtain further information about the event. In real-world applications, we expect that the bulk of the transmission is for normal packets and transmission of event packets is relatively infrequent. This means that it is especially important for the sender to minimize transmission energy for normal packets, while being able to quickly alert the receiver when events of interest occur.

We note that for many sensor applications this conventional approach is suboptimal in terms of energy use. For example, there is no need for the sender to transmit at a relatively high transmit power to ensure all these normal packets transmitted can be decoded by the receiver, if the time series of packet arrivals can already reveal that the sender is in the normal state. Upon noticing events of interest a sender merely need to seek attention from the receiver about the new situation. To this end, the sender can just transmit packets with a different pattern in time series. The receiver can then use a robust inference method to classify the sender being in a normal or event state based on patterns in the time series of packet arrivals, without having to decode packets.

We have explored such inference-based approaches where no packet decoding is required. This would enable the receiver to operate at a lower signal-to-noise ratio (SNR), and, in turn, allow the sender to reach receiver at the same distance with lower transmit power or, equivalently, farther receivers with the same transmit power.

A key issue with such approaches is their accuracy in classifying the current state of the sender in low SNR situations when the receiver is quite far away, and/or the wireless channel is noisy. We have shown how a two-layer hierarchical inference can be effective in providing robust and reliable classification based on the packet arrival patterns, even when some packets may have distorted symbols or may be completely lost. We call our approach Wireless Inference-based Notification without Packet Decoding, or for short, WIN.

In the following, we describe the conventional approach of transmitting notifications, and then describe how our proposed WIN approach can accomplish the same task with lower energy consumption. As depicted in Figure 5, under the conventional approach a sender periodically transmits normal packets (black) to a receiver to report that the sender is alive and it is in a normal state. Upon noticing events of interest, the sender enters the event state and starts transmitting event packets (red). The receiver will attempt to decode every received packet to determine the state of the sender.

Under a corresponding WIN approach, the sender in the normal state will periodically transmit normal packet like in the conventional approach. When the sender enters the event state, it will transmit event packets periodically under a different arrangement about the length of packet burst or gap. Figure 5 depicts an example of such a WIN scheme based on the following time series of packet transmissions:

Normal state: burst =1 and gap = 3

Event state: burst = 2 and gap = 6

Note that in supporting WIN, a conventional sender does not need to change its protocol stack; all it needs to do is to change packet transmission patterns during the event state. Thus existing sensor transmission systems are readily useable.



Figure 5. Conventional approach vs. WIN. Time slots labeled by time are shown at the bottom. Solid bars denote normal (black) and event (red) packets transmitted at various time slots.

The receiver employs physical-layer matched filters to determine whether each time slot has an arriving packet. Based on the matching results from multiple time slots, the receiver uses inference methods to infer the state of the sender, as depicted in Figure 6. By making use of aggregated matching results from multiple time slots and leveraging the designed-in separation between the time series of packet transmissions for the normal vs. event state, a WIN receiver can operate at a lower SNR. As a result, a distant receiver may still be able to infer the state of the sender even it cannot decode normal or event packets. This is illustrated in

Figure 7. When a receiver determines that the sender is in the event state, should the receiver happen to be mobile, it could move itself closer to the sender to decode the event packet and

learn about the event. Alternatively, the receiver may dispatch other agents for the task.




Figure 6. When detecting packets in the first layer, we use max-pooling with a sliding-window to address variations in packet delay due to multipath.




Figure 7. The WIN receiver can receive notification from the sender at a distance beyond the packet decoding range. In contrast, a conventional receiver can receive notification only in the packet decoding range.

4.0 RESULTS AND DISCUSSION

In this section, we present highlights of our results in the three focus areas.

4.1 Concurrent Channel Estimation in Scalable Multiuser MIMO Networking

We use a 4×4 MIMO scenario to evaluate the performance of concurrent channel estimation in a lab environment. The performance is compared against a baseline case where interference-free preambles are transmitted sequentially. In the setting, we assume there are 100 senders but only 4 of them transmit at any given time. The distance between the transmitters and the receivers is around 2 to 3 meters. We vary the transmission power and the distances to get different SNR values.

For the baseline scheme to which MIMO/CON compares, we apply the standard least squares method [12] to interference-free preambles. The channel is estimated by solving the following equation:

$$\hat{\mathbf{h}} = \Phi^{-1}\hat{\mathbf{y}} \quad (5)$$

where $\Phi = \text{diag}(\mathbf{a}_i)$ and \mathbf{a}_i is the known preamble sequence and $\hat{\mathbf{h}}$ is the channel estimation.

In both cases, the obtained channel estimate is then used to decode 4 MIMO data streams immediately followed by the preamble with the standard zero-forcing method and successive interference cancellation [10]. The FFT size of both preamble and data symbols are set to 128 points. We repeat each experiment 300 times with different random preambles.

Figure 5 shows the scatter plots of the decoded SNR of the subsequent data transmission decoding with the channel estimated from interference-free preambles versus that from concurrent preambles. The experimental results reveal the following: first, taking 13 taps (6 on each side of the significant tap) is sufficient for channel estimation in all SNRs. Taking fewer taps can result in a degradation in decoded SNR because the recovered CSI is less accurate. The number of taps required also determines the preamble length for sufficient measurement. Second, with a sufficient number of taps such as 13 taps in Figure 5(a), the decoding performance with concurrent preambles is better than interference-free preambles. This shows that MIMO/CON in fact can help filter out noise in channel estimation. This is because a nonzero value appearing in

a large-delay tap is suppressed because it violates the sparsity property presented in channel delay statistics. Third, when the signal SNR is high, more taps are required to achieve relatively good decoding performance. For the case with low signal SNR, the accuracy of channel estimation is limited by the noise and thus taking fewer taps is sufficient.

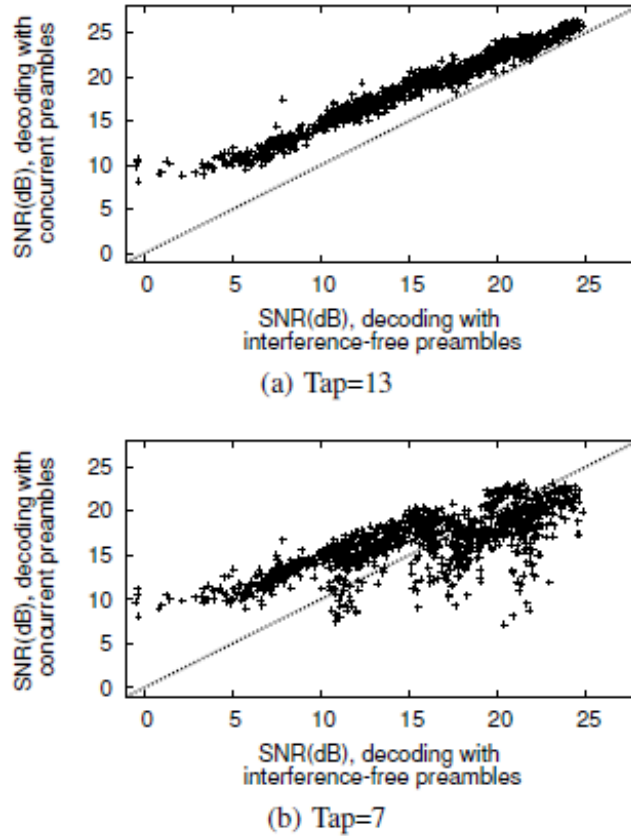


Figure 8. MIMO decoding performance using CSI estimated from concurrent preambles in 4x4 MIMO. Taking 13 taps is sufficient for reconstructing accurate CSI. Using fewer taps results in degradation in decoding performance, especially when the signal SNR is high.

4.2 Use of COTS Antenna Arrays to Determine RF Angle of Arrival

Maximum Likelihood AOA. Instead of directly applying Equation (1) to estimate the incident direction of the target signal, for estimation robustness, 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 presented at each grid location. The basic idea is that, at each location,

there is an expected phase difference and that by comparing it against the measured phase difference, we can evaluate the likelihood of the target being at a particular location. Then, the location with the maximum likelihood would be our best estimate. For this work, we focus on the single target case since we only use two antennas on each receiver, but the system could be extended to handle multiple targets if more antennas are added [14].

Recall that with Equation (4), we can measure the target signal's phase difference $\Phi_{(12)}^A = \hat{\Phi}_1^A - \hat{\Phi}_2^A$ at the two receive antennas. We can then compute the likelihood of each location x with a likelihood function defined as follows:

$$L(x) = \mathcal{N}(\Phi_{(12)}^A | \Phi_{(12)}^x, \sigma^2) \quad (6)$$

assuming the measurement error follows a zero mean Gaussian distribution with the variance $\sigma^2 = \text{var}(\Phi_{(12)}^A)$.

Figure 9 (left and center panels) shows 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 4). Since the locations on the same line of the incident angle share the same phase difference, they will have equal likelihood (i.e., the same color). Note that lines representing equal likelihood in Figure 9 are not exactly straight, but curve around the receive antenna array. This is an artifact 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 written as:

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

The joint likelihood results for the same example run are shown as a heat map plot in Figure 9 (right panel). From this, we can choose the location with the maximum likelihood as the location estimate for the target.




Figure 9. 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 degree

Localization Accuracy. The results of the ten experiment runs in localizing target A are summarized in Table 1. From this, we can 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 [15] 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.

Table 1. 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.

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

4.3 Wireless Inference-based Notification (WIN) without Packet Decoding

We have compared the performance of WIN and the conventional approach using analysis and simulation. In this work, we assume that the number of total packets (R) in a complete transmission is 20, and the number of bits per packet (n) is 80. Since CRC error becomes more likely when the packet size is larger, we select the smallest packet size for a wireless network to avoid bias against the conventional method. This size is 80 bits according to the specifications of Bluetooth LE [21]. We simulate with two channel models: AWGN channel and AWGN channel with uniform random packet delay. The results are shown in Figure 5. Under AWGN channel, WIN achieves error rates lower than 1% as long as the received SNR is greater than -10 dB (see blue curve), while the conventional method has more than 1% error at 3 dB. In other words, there is roughly a 13 dB gain for WIN. Note that our analytic estimations match closely to the results obtained by simulation.

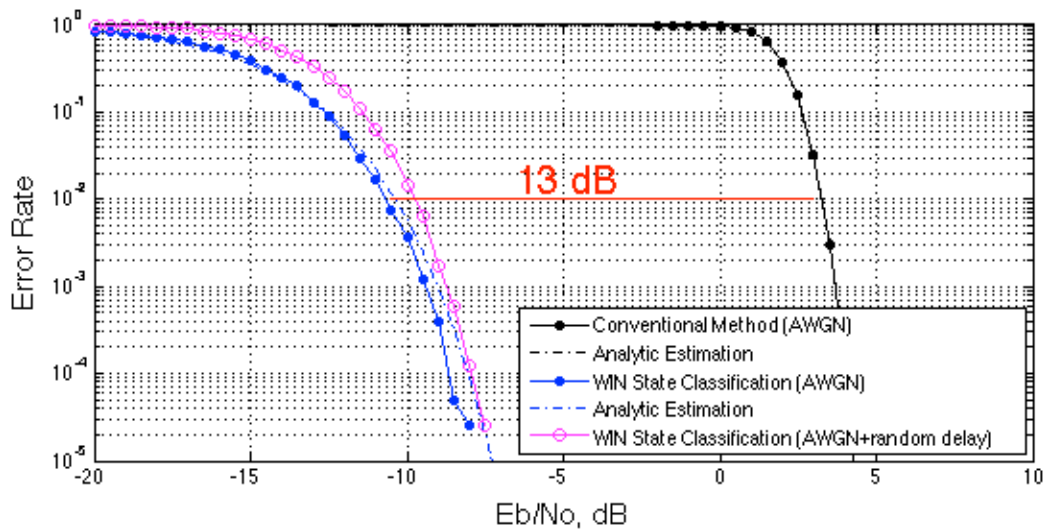


Figure 10. WIN exhibits approximately a 13dB gain over conventional approach under AWGN channel.

5.0 CONCLUSIONS

Here we summarize our results and findings in the three focus areas of this report.

5.1 Concurrent Channel Estimation in Scalable Multiuser MIMO Networking

We have proposed an ambitious scheme for the purpose of achieving full utilization of uplink capacity offered by an AP equipped with many receive antennas. The key to our scheme, called MIMO/CON, is a novel compressive sensing based decoding method which can estimate channel state information and identify users from concurrently received preambles. By exploiting sparsity in the target signal and use environment, preambles can be kept short while allowing high-quality decoding for MIMO. We have demonstrated the working of the proposed concurrent channel access and estimation method using hardware testbed based on software-defined radios. In addition, the proposed decoding method can be applied to opportunistically decode conflicting transmissions over time and boosts channel utilization. Overall, our work has offered a comprehensive study of MIMO/CON and we believe concurrent access will be an important component for future multiuser MIMO networks.

5.2 Use of COTS Antenna Arrays to Determine RF Angle of Arrival

We have shown that the COTS approach to build phased array antennas can provide important benefits including flexible configuration, low cost, and easy ad hoc deployment. We have taken such an 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 phased 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 future such challenges by leveraging our expertise in UAV flight experiments [1] .

5.3 Wireless Inference-based Notification (WIN) without Packet Decoding

Conventional network layering is provided to support modular design principles, but it is at the expense of losing information in each layer. For example, in the physical layer we loss information from demodulation and in the link layer we loss information when we toss the entire packet upon CRC errors. Furthermore, conventional design avoids utilizing prior knowledge because it is not always available. Such information loss and underutilization means a substantial drawback for applications that have stringent low-energy requirements. Via interference technology based on machine learning, WIN aims at making use of all information resulting from physical-layer matched filtering operations. In addition, WIN leverages designed-in separation between traffic patterns of different states of the sender, so the state classification can be tolerant to channel noise. For these reasons, we have shown that WIN can achieve 13 dB gains in terms of robustness against channel noise. Lowering the required signal strength at receiver by 13 dB translates to 4.5x range in free space.

We can view WIN as a beginning of a new class of low-power coding methods based on packet arrival patterns learned by receiver. These new protocols could be especially useful for future ultra-low power designs for notification transmission over wireless channels.

6.0 REFERENCES

- [1] J. Andersen, T. Rappaport, and S. Yoshida. Propagation measurements and models for wireless communications channels. *IEEE Commun. Mag.*, 33(1):42–49, 1995.
- [2] J. Andrews et al. *Fundamentals of WiMAX: understanding broadband wireless networking*. Prentice Hall PTR, 2007.
- [3] W. Bajwa et al. Compressed channel sensing: A new approach to estimating sparse multipath channels. *Proc. IEEE*, 98(6):1058–1076, 2010.
- [4] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *Selected Areas in Communications, IEEE Journal on*, 18(3):535–547, 2000.
- [5] D. Gesbert et al. Shifting the MIMO paradigm. *IEEE Signal Process. Mag.*, 24(5):36–46, 2007.
- [6] A. Gudipati and S. Katti. Strider: automatic rate adaptation and collision handling. In *ACM SIGCOMM 2011*.
- [7] K. Lin et al. Random access heterogeneous MIMO networks. In *ACM SIGCOMM 2011*.
- [8] E. Magistretti et al. WiFi-Nano: reclaiming WiFi efficiency through 800 ns slots. In *ACM MobiCom 2011*.
- [9] T. Marzetta. Noncooperative cellular wireless with unlimited numbers of base station antennas. *IEEE Trans. Wireless Commun.*, 9(11):3590–3600, 2010.
- [10] K. Tan et al. SAM: enabling practical spatial multiple access in wireless LAN. In *ACM MobiCom 2009*.
- [11] D. Tse and P. Viswanath. *Fundamentals of wireless communication*. Cambridge Univ Pr, 2005.
- [12] J. Van de Beek, O. Edfors, M. Sandell, S. Wilson, and P. Borjesson. On channel estimation in OFDM systems. In *IEEE VTC 1995*.
- [13] IEEE 802.11n-2009. Amendment 5: Enhancements for higher throughput. October 2009.
- [14] R. Schmidt. Multiple emitter location and signal parameter estimation, *Antennas and Propagation, IEEE Transactions on*, vol. 34, no. 3, pp. 276 – 280, March 1986.
- [15] 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 *ACM MobiCom 2009*.
- [16] D. Hague, H. T. Kung, and B. W. Suter. Field experimentation of COTS-based UAV

- networking. In MILCOM 2006.
- [17] C. Savarese, J. M. Rabaey, and K. Langendoen. Robust Positioning Algorithms for Distributed Ad-Hoc Wireless Sensor Networks. In USENIX ATC 2002.
 - [18] A. Savvides, C-C Han, and M. B. Strivastava. Dynamic Fine-grained Localization in Ad-Hoc Networks of Sensors. In ACM MobiCom 2001.
 - [19] Ettus Universal Software Radio Peripheral (USRP).
 - [20] J. V. Capella, A. Perles, A. Bonastre, and J. J. Serrano. Historical Building Monitoring Using an Energy-Efficient Scalable Wireless Sensor Network Architecture. In Sensors 2011.
 - [21] Bluetooth specification core version 4.0,
www.bluetooth.org/Technical/Specifications.

7.0 LIST OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS

AOA	Angle of arrival
AP	Access Point
AWGN	Additive White Gaussian Noise (AWGN)
COTS	Commercial-off-the-shelf
CRC	Cyclic Redundancy Check
CSI	Channel State Information
dB	Decibel
DCF	Distributed Coordination Function
HTTP	Hypertext Transfer Protocol
IEEE	Institute of Electrical and Electronics Engineers
ISR	Intelligence, Surveillance and Reconnaissance
MAC	Media Access Control
MIMO	Multiple-Input and Multiple-Output
MIMO/CON	MIMO with concurrent channel access and estimation
MU-MIMO	Multiuser MIMO
OFDM	Orthogonal Frequency Division Multiplexing
PDOA	Phase Difference of Arrival
PHY	Physical layer
RF	Radio Frequency
RSS	Received Signal Strength
SDR	Software-Defined Radios
SNR	Signal-to-Noise Ratio
UAV	Unmanned Aerial Vehicle
VoIP	Voice over Internet Protocol
WiFi	Wireless Fidelity
WIN	Wireless Inference-based Notification