STROBE: Actively Securing Wireless Communications using Zero-Forcing Beamforming

Narendra Anand
Rice University
Houston, USA
Email: nanand@rice.edu

Sung-Ju Lee
Hewlett-Packard Laboratories
Palo Alto, USA
Email: sjlee@hp.com

Edward W. Knightly
Rice University
Houston, USA
Email: knightly@rice.edu

***Abstract*—We present the design and experimental evaluation of *S*imultaneous *TR*ansmission with *O*rthogonally *B*linded *E*avesdroppers (STROBE). STROBE is a cross-layer approach that exploits the multi-stream capabilities of existing technologies such as 802.11n and the upcoming 802.11ac standard where multi-antenna APs can construct simultaneous data streams using Zero-Forcing Beamforming (ZFBF). Instead of using this technique for simultaneous data stream generation, STROBE utilizes ZFBF by allowing an AP to use one stream to communicate with an intended user and the remaining streams to orthogonally "blind" (actively interfere with) any potential eavesdropper thereby preventing eavesdroppers from decoding nearby transmissions. Through extensive experimental evaluation, we show that STROBE consistently outperforms Omnidirectional, Single-User Beamforming (SUBF), and directional antenna based transmission methods by keeping the transmitted signal at the intended receiver and shielded from eavesdroppers. In an indoor Wireless LAN environment, STROBE consistently serves an intended user with an SINR 15 dB greater than an eavesdropper.**

I. Introduction

The broadcast nature of wireless communication necessitates the development and use of robust security protocols to thwart eavesdroppers from intercepting transmissions directed toward an intended user. While encryption mitigates this vulnerability, even industry standard encryption methods such as WEP and WPA could be compromised [1] and readily available software packages enable malicious users to defeat secure networks.

One method of enhancing the security of wireless transmissions is to prevent the eavesdropper from receiving or decoding the transmitted signal. A candidate solution is a directional transmission scheme that focuses signal energy toward an intended receiver using a directional antenna, switched-beam, or a single-target adaptive beamforming transmission. However, in practice, such techniques that depend on the predictable behavior of transmitted beam patterns or that are agnostic to the entire eavesdropper environment fail to prevent eavesdropping as confirmed by our own experiments and also in [2].

To address this problem, we design a new, multi-antenna, 802.11-compatible scheme that adaptively sends a beam toward an intended user while "blinding" (actively interfering with) potential eavesdroppers, STROBE (**S**imultaneous **TR**ansmission with **O**rthogonally **B**linded **E**avesdroppers). STROBE leverages the potential of a Zero-Forcing Beamforming (ZFBF) transmitter to send a signal toward an intended user while simultaneously transmitting "orthogonally blinding" streams (defined in § II-B) everywhere else.

ZFBF is a precoding method that enables a multi-antenna access point (AP) to create multiple simultaneous spatial streams [3]. Recent wireless standards such as 802.11n or the upcoming 802.11ac[1] employ physical layers (PHYs) that can implement ZFBF to construct multiple parallel transmission streams to a single user (11n) or simultaneously to multiple users (11ac). Because such existing technologies are already able to create multiple parallel streams, STROBE can be easily implemented in these systems with minor AP modification and no client modification. STROBE is orthogonal to WEP or WPA encryption methods and can be deployed jointly to further enhance wireless security.

This paper has the following main contributions: First, we design and implement STROBE in an FPGA-based software defined radio platform. Moreover, for comparative evaluation, we implement (*i*) Omnidirectional, (*ii*) Single-User Beamforming (SUBF), (*iii*) Directional Antenna, and (*iv*) Cooperating Eavesdropper (CE) schemes. CE is an unrealistic scheme in which eavesdroppers "cooperate" by providing their channel information to the transmitter. While, in practice, eavesdroppers would *never* aid the transmitter (rendering this method wholly unrealistic in practice), CE provides a best case benchmark for "blinding" eavesdroppers via ZFBF.

Second, we evaluate STROBE's performance against the aforementioned schemes in a baseline WLAN scenario. Our testbed experiments show that STROBE better controls leaked signal energy by actively thwarting eavesdroppers by transmitting simultaneous interference streams, severely diminishing their ability to eavesdrop. Additionally, we show that even when compared against the (unrealistic) Cooperating Eavesdropper scheme, STROBE realizes a greater signal energy difference between the intended user and the eavesdropper. This is accomplished even though the unfeasible CE scheme is aware of eavesdropper channel information.

Third, we show that despite the use of beamforming in our system design, eavesdropper proximity or orientation relative to the intended user has a *negligible effect* on STROBE's ability to serve an intended user while blinding potential eavesdroppers. STROBE exploits multi-path effects by harnessing signal reflections to reach the intended user. In fact, at a relative eavesdropper proximity of a quarter wavelength (3.25 cm) from the intended user, STROBE still serves the intended user with at least a 10 dB stronger signal than the eavesdropper.

Fourth, we explore STROBE's dependence on multi-path reflections by performing experiments in an open, outdoor environment. Because the environment contains no physical obstacles to cause reflections, STROBE must use the direct, line-of-sight (LOS) path to serve the intended user. We find a marked detrimental effect on STROBE's efficacy as eavesdroppers can easily overhear signal energy at close-by locations, i.e.,

This research is sponsored in part by HP Labs' Innovation Research Program and by NSF Grants CNS-1012831 and CNS-1126478.

[1]See standards.ieee.org and mentor.ieee.org for 802.11n and 802.11ac.

STROBE requires a multi-path rich, WLAN type environment.

Finally, we consider a nomadic eavesdropper that traverses an environment attempting to find a location to successfully eavesdrop. We show that even if the eavesdropper exhaustively traverses the room, it is still thwarted by STROBE. In contrast, eavesdroppers can very easily find suitable eavesdropping locations for the other transmission schemes considered, including when using a directional antenna.

The rest of the paper is organized as follows: § II gives background on ZFBF and the design of STROBE's orthogonal blinding method. § III describes our experimental platform and evaluation methodology. § IV presents a baseline evaluation of STROBE. § V explores the effects of eavesdropper proximity and location with relation to the intended user and transmitter. § VI evaluates STROBE in an open, outdoor environment with fewer multi-path effects. § VII evaluates the robustness of STROBE against a nomadic eavesdropper. § VIII describes related work and § IX concludes the paper.

II. STROBE Design

In this section, we first describe the mechanics of Zero-Forcing Beamforming, the core technique behind STROBE. We then detail the mechanism of "Orthogonal Blinding," the key component of STROBE that enhances wireless security.

System Model. STROBE is a downlink transmission technique. We consider a system consisting of a multi-antenna AP and several single antenna users. We believe this system is typical to current WLANs as APs have the ability to support complex, multi-antenna technologies whereas users (e.g. smartphones) are limited to singular antenna methods by constraints such as size, computational ability, and power consumption.

Of the single-antenna users, we call the user to which the transmission is intended the "Intended User" (IU). We call the unintended users who are attempting to overhear the "Eavesdroppers" (E).

Notation. The following describes the notation used throughout the paper. Further definition and description will follow in the appropriate sections. N refers to the number of transmit antennas at the AP. M refers to the number of concurrently served, single-antenna users.

The row vector h_m is a $1 \times N$ channel state vector for user m. Each element of h corresponds to the complex exponential gain between one of the transmitter's antenna and the user. The matrix $H = [h_1; h_2; \ldots; h_M]$ is the $M \times N$ channel matrix constructed using each user's h as its rows.

The column vector w_m is an $N \times 1$ beamsteering weight vector for user m. Each element of w corresponds to the complex exponential gain used by each transmitting antenna. The matrix $W = [w_1\ w_2\ \ldots\ w_m]$ is the $N \times M$ steering weight matrix consisting of each user's w as its columns.

A. ZFBF Overview

ZFBF is a precoding transmission method that enables an AP to construct multiple, concurrent spatial streams that can transmit data to multiple users in parallel. The basic principle is to first take each user's view of the channel, h, and construct a corresponding w for each h. Each user's data stream is then multiplied by its corresponding w, summed together and transmitted over the AP's antenna array. Careful selection of w is required for the construction of concurrent spatial streams and parallel transmission of multiple users' data.

The optimal method of constructing W from H to concurrently serve multiple users is known as Dirty Paper Coding (DPC) [4], [5]; however, in practice this method is difficult to implement due to its complexity. Instead, ZFBF is a sub-optimal W construction method that is simpler to implement yet still achieves a performance almost equivalent to DPC [3].

ZFBF selects weights that cause zero inter-user interference (the effect of one beamformed stream on another is "forced" to zero). The authors of [6] have shown that the optimal selection of W to satisfy this zero-interference condition is the pseudo-inverse of H as shown in Eq. (1).

$$W = H^{\dagger} = H^{*}(HH^{*})^{-1} \tag{1}$$

The use of the pseudo-inverse is how the zero-interference condition is achieved: if $W = H^{\dagger}$ then $h_i w_j = 0$ for $i \neq j$. Additionally, note that the matrix multiplication in Eq. (1) places a limit on the maximum number of concurrent users (or spatial streams). The number of concurrent streams (M) must be less than or equal to the number of transmit antennas (N).

In our implementation, we feed back channel state information (CSI), an h vector, in a manner analogous to the RTS/CTS exchange in compliance with 802.11ac and 802.11n.

B. Orthogonal Blinding

The key mechanism of STROBE is "orthogonal blinding" that occurs in parallel with transmission to the intended user. "Blinding" is the method of actively concealing the intended user's signal by overwhelming any potential eavesdroppers with garbage transmissions. These "blinding streams" are transmitted concurrently with the intended user's signal by the ZFBF enabled transmitter using its remaining available streams. The blinding streams are constructed orthogonally to the intended user's signal to ensure that these streams cause the least possible decrease of the intended user's signal.

The streams used for the intended user and for blinding correspond to different w vectors, which come from the pseudo-inverse of H. Thus, to construct orthogonal blinding streams, we generate h vectors orthogonal to the intended user's h and then perform ZFBF on the constructed H matrix.

To construct these orthogonal h vectors, we use the Gram-Schmidt process [7]. First, we take the intended user's CSI (h_1) and pad h_1 with a truncated $(M\text{-}1) \times N$ identity matrix to build a preliminary H. Finally, we construct the CSI matrix with orthogonal rows, $\widehat{H}$, by using the Gram-Schmidt process shown in Eq. (2) on the preliminary H. Since the resulting $\widehat{H}$ is unitary, the calculation of its pseudo-inverse is trivial: $W = \widehat{H}^{\dagger} = \widehat{H}^{*}$.

$$\widehat{h_k} = h_k - \sum_{j=1}^{k-1} \frac{\langle h_k, \widehat{h_j} \rangle}{\|\widehat{h_j}\|^2} \widehat{h_j}, \quad 1 \leq k \leq M. \tag{2}$$

The Gram-Schmidt process is simple to integrate into a ZFBF enabled 802.11n/ac AP. The first step of calculating the required pseudo-inverse is implemented in hardware using QR decomposition [8], an operation that decomposes a matrix into an upper triangular (R) and a unitary matrix (Q). The Gram-Schmidt process can also be executed using the QR method. Thus, the silicon in the PHY of a ZFBF enabled 802.11n/ac AP already exists to perform this algorithm; the only change necessary is how the input matrix is loaded.

III. Experimental Methodology

A. Experimental Platform

We conduct our experiments using WARPLab [9], a framework that integrates the versatility of MATLAB with the capabilities of an FPGA-based software defined radio platform (WARP, see Fig. 1). WARPLab gives the ability to rapidly prototype physical layer algorithms in MATLAB while using the WARP nodes to perform over-the-air (OTA) characterizations.

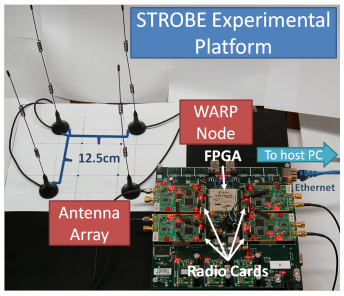

Fig. 1. STROBE Experimental Platform.

The WARPLab flow consists of two main parts: MATLAB on the host PC and the WARP node. All physical layer baseband processing occurs in MATLAB while the OTA transmission and reception are handled by the WARP nodes.

A single host PC can be connected to up to 16 WARP nodes through an Ethernet switch. MATLAB on the host PC processes a given bit stream using the implemented algorithm and then downloads the generated I/Q samples to the transmitting WARP node. The host PC then sends a sequence of control signals that primes the state machines of the connected nodes and triggers the transmission and reception of the OTA signal.

Each node contains four large sample buffers connected to four 2.4 GHz radio cards and antennas. These buffers either accept data over Ethernet for transmission or through the radio card for reception. Receiving nodes can then upload the received samples for decoding along with received RSS (Received Signal Strength) values in dBm.

For the evaluation of STROBE, the transmitter must sound the channel in order to obtain the relevant CSI (h vector) for the necessary receivers. This channel sounding followed by transmission is accomplished by implementing CSI feedback through the switch, H and W calculation in MATLAB, and beamforming weight multiplication in the WARP node. For the characterization of STROBE, we use the ZFBF experimental framework, built on top of WARPLab by the authors of [10].

We employ one transmit antenna for the receivers and all four for the transmitter (thus $N = 4$ for our characterization). Our antenna array is circular with the antennas spaced one wavelength apart for the 2.4 GHz wavelength ($\lambda = 12.5$ cm).

B. Compared Schemes

Omnidirectional Transmission. The initial baseline scheme for comparison is the Omnidirectional transmission. Because this is the most commonly used transmission method, it is important to observe where the intended user's transmitted energy is sent. This method reflects the status quo conditions under which existing encryption protocols operate.

Single-User Beamforming. The fundamental adaptive signal energy direction technique is Single-User Beamforming (SUBF). SUBF employs an antenna array to steer a beam toward an intended user based on that user's CSI (h vector). SUBF can be considered a subset of ZFBF in that the number of "concurrent" users is one. Because there is only one intended user, the zero-interference condition does not exist (since there is no other stream to interfere with) so the weight selection results in the maximum possible received signal energy at the intended user (for a ZFBF type scheme). Because the H matrix consists of only one vector, the SUBF steering weight is simply $W = (H_{1\times N})^{\dagger} = h^{\dagger} = h^{*}$. Thus, the intended user's steering weight for SUBF is its complex conjugate transpose, which is equivalent to the intended user's weight for STROBE.

However, to ensure a fair comparison, the power allocation to the steering weights of SUBF and STROBE differs, which contributes to a difference in intended user's received signal energy. This difference (discussed in § III-C) results in a $4\times$ greater transmit power allocated to the intended user's steering weight when using SUBF compared to STROBE.

Cooperating Eavesdropper. Finally, we compare STROBE against the unfeasible baseline, the Cooperating Eavesdropper (CE). This scheme explores the unrealistic scenario where the eavesdroppers actively aid in their blinding by providing the transmitter with their channel estimates. While this scenario would never occur in practice, it is essentially an upper limit of a ZFBF-based scheme's potential blinding performance.

With knowledge of eavesdropper's CSI (their h vectors), the transmitter has access to the "true" H matrix. This allows the transmitter to precisely blind eavesdroppers due to the zero-interference condition. Specifically, this condition signifies that the intended user's stream will have zero energy (zero interference) at the cooperating eavesdroppers' locations. Thus, even if the transmitter does not use the additional streams for blinding, the eavesdroppers will still be unable to overhear the intended user's signal. It was shown in [10] that a four antenna transmitter serving four concurrent users causes less than 1 dB of inter-user interference. Although this implies that CE can construct a beam to the intended user that cannot be overheard by the (cooperating) eavesdroppers, we still use the remaining three streams for blinding to ensure a fair comparison.

C. Measurement Procedure

Performance Metric. Our performance metric is Signal to Noise Ratio (SNR) or Signal to Interference plus Noise Ratio (SINR) expressed in dB. As mentioned in § III-A, the WARPLab platform allows us to measure RSS in terms of dBm; however, the inherent differences in RSS measurements between radio transceivers result in dBm readings that cannot be fairly compared.

To overcome this, we calculate SNR and SINR values from the difference between two back-to-back measurements performed at the intended user and each eavesdropper. The first RSS measurement indicates the overall received signal strength and the second contains only the noise (or noise plus interference). Thus the difference between the two is the signal to (interference plus) noise ratio. This measurement method is further detailed in [10]. For the remainder of this paper, we refer to this received signal energy as the SINR.

Power Allocation To ensure a fair comparison, we set the net transmit power for all schemes equivalent regardless of the number of antennas or streams used. Omnidirectional (one antenna) and SUBF (four antennas) transmissions use equivalent power to serve the intended user. STROBE and CE

generate N transmit streams so each stream is allocated $1/N^{\text{th}}$ the overall transmit power. This net transmit power control is implemented by the normalization of the W matrix.

Data Collection Each data point presented is an average of 30 OTA transmission measurements with standard deviations of 1 dB or less (shown for each data point as a confidence interval). All experiments were conducted in an interference-free, unused channel (802.11 channel 14).

IV. Baseline WLAN Scenario

First, we evaluate STROBE using a baseline WLAN topology. Namely, we explore STROBE's ability to serve an intended user while blinding the intended signal to the eavesdroppers by exploiting a multi-path rich (indoor) environment.

A. Experimental Setup

To realize a typical WLAN scenario, we create a conference room topology. Specifically, as depicted in Fig. 2, four receivers are placed along the far edge of a large table. The room itself is in the shape of a long rectangle filled with metal chairs and surrounded by metal blinds making it a multi-path rich environment. The receivers are separated by 1.25 m and the AP is separated from the group of users by a 5 m distance. We set one receiver to be the intended user (labeled IU) and the other three receivers to be eavesdroppers (labeled $E_{1\text{-}3}$). We also perform the experiments where the intended user is in the other three eavesdropper locations and obtain similar results, and hence not shown.

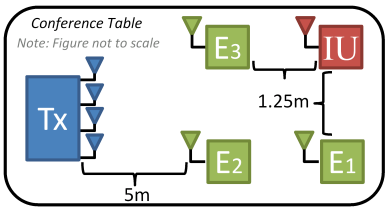

Fig. 2. Basic WLAN Topology.

B. Experimental Results

Fig. 3 shows the received SINR at each of the four receivers when data was transmitted toward the intended user. The SINR at the intended user indicates the received signal energy of the intended transmission whereas the SINR at $E_{1\text{-}3}$ shows the overheard signal energy.

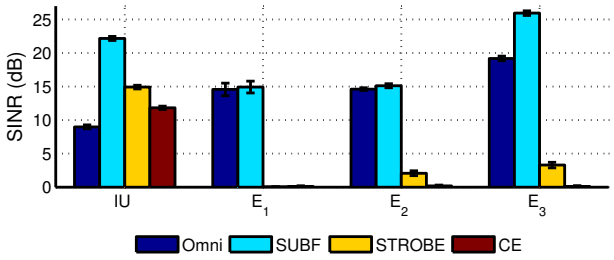

Fig. 3. SINR at IU and Overheard SINR at $E_{1\text{-}3}$.

Omnidirectional Transmission. The use of an Omnidirectional transmission yields a greater SINR at the eavesdropper locations than the intended user for this topology. This highlights the high vulnerability of existing systems to transmissions overheard by eavesdroppers, a critical security issue when encryption protocols are unused or defeated.

Single-User Beamforming. In contrast to an Omnidirectional transmission, SUBF directs a signal's energy toward the intended user. However, a large amount of energy is still available at other locations for the eavesdroppers to overhear and intercept. In fact, in this topology, the energy at the eavesdropper locations is greater than that of the Omnidirectional transmission. This behavior is not a flaw but a result of the design. SUBF's goal is to maximize the SINR at the intended user, but does so completely agnostic to other locations. For this reason, the signal transmitted by SUBF does not remain solely at the intended user but instead spills over to various other areas causing the scheme to be as vulnerable as an Omnidirectional transmission. This specific consequence of SUBF is further highlighted by the transmit power allocation method. As mentioned in § III-C, in order to ensure a fair comparison, the net transmit power for all schemes is equivalent regardless of the number of antennas used. Thus, the eavesdroppers overhear the intended user's signal just as well (if not better as at E_3) as the Omnidirectional transmission even with equivalent transmit powers and SUBF's inherently directional nature.

STROBE. Unlike Omnidirectional and SUBF transmissions, STROBE blinds the eavesdroppers and mitigates the possibility of overhearing the signal to the intended user. Actively blinding potential eavesdropper locations results in a maximum 3 dB eavesdropper SINR (at E_3) for this topology. Although the intended user's steering weight is identical to SUBF (as discussed in § III-B), the use of additional, orthogonal blinding weights enables STROBE to diminish the SINR at the eavesdroppers while still maximizing the SINR at the intended user. Again, this effect is further emphasized by the power allocation scheme; not only does the net transmit power remain fixed regardless of the number of antennas used but also regardless of the number of spatial streams employed. The use of four streams in our evaluation forces a $4\times$ decrease in the transmit power allocated to the intended user. However, in spite of this resulting 6 dB SINR decrease at the intended user, STROBE causes a *15-22 dB decrease* in SINR at the eavesdroppers.

To estimate how STROBE's measured energy levels would perform on a commodity device, we employ the theoretical Gaussian model for BPSK (for a conservative estimate) and compute approximate bit error rates (BERs) for the observed SINRs. The intended user's served SINR (14 dB) and maximum eavesdropper SINR (4 dB) correspond to approximate BERs of 5.4×10^{-5} and 2.3×10^{-2} respectively. This conservative estimate shows that STROBE serves an intended user with a BER three orders of magnitude lower than the eavesdropper, thus significantly decreasing the likelihood of an eavesdropper decoding an intended user's transmission.

Cooperating Eavesdropper. As a baseline for evaluating STROBE, we examine the unrealistic scenario of the cooperating eavesdropper where the eavesdroppers provide their channel information to the transmitter. The extra information provided to the transmitter allows for precise eavesdropper blinding. This additional accuracy manifests as eavesdropper SINR equaling approximately 0 dB, yet this is still only a 3 dB decrease in overheard signal energy at E_3 from STROBE. However, this decrease comes at the cost of a 3 dB decrease in the intended user's SINR so the relative gain of CE at the intended user over the eavesdropper is equivalent at E_3 and less at $E_{1\text{-}2}$ than STROBE. We further explore this effect next.

C. STROBE vs. CE

For our comparisons between STROBE and CE, we purposely set the number of overall receivers to four, a decision that results in the "best case" results for CE. This decision and the observed results highlight a subtle yet important difference in the two schemes' mechanisms.

The precision of any ZFBF based transmission scheme is dependent on the number of transmit antennas. The number of spatial streams that may be constructed is equal to the size of the transmit antenna array. Although both CE and STROBE can create an equivalent number of streams, STROBE creates its blinding streams solely based on the channel state of the intended user whereas CE considers all users. The consequence of this characteristic is that CE is only able to precisely blind as many eavesdroppers as one less than the number of transmit antennas. If we were to perform this experiment with additional eavesdroppers, our four antenna transmitter (when employing CE) would only be able precisely blind three of the eavesdroppers; the remaining eavesdroppers would overhear the signal with an SINR comparable to STROBE.

Both STROBE and CE construct multiple blinding streams using ZFBF. However, the manner in which STROBE constructs the channel matrix around the intended user's channel state (h vector) guarantees a maximum served SINR to the intended user because the constructed matrix is unitary and the resulting intended user's $w = h^*$ (as detailed in § II-B). This resulting w vector is equivalent to the SUBF weight and is the best that ZFBF can provide. The resulting W matrix still satisfies ZFBF's zero interference condition. However, CE's construction of the H matrix is based on all users' channel estimates. For CE to satisfy the zero interference condition, the intended user's weight loses magnitude.

Thus, the performance at an eavesdropper of STROBE and CE (after the maximal number of eavesdroppers) is equivalent; however, STROBE will always serve the intended user with a higher SINR. We verify this in Fig. 3 where CE results in less overheard signal energy but also less served energy to the intended user. Although CE is completely unfeasible in a real system and uses ZFBF in its intended manner for precise blinding, STROBE will always provide a higher SINR to the intended user and the benefits of this scheme will still function regardless of the number of eavesdroppers.

V. Relative Eavesdropper Location

We now evaluate the effect of eavesdropper position relative to the transmitter and the intended user. This analysis is done in two parts. § V-A examines the effect of eavesdropper proximity to the intended user. § V-B examines the effect of eavesdroppers inline with the intended user and transmitter. We conduct both experiments in the same setting as the experiment in § IV.

A. Eavesdropper Proximity

The purpose of this experiment is to quantify the effect of eavesdropper distance relative to the intended user. Because this is a spatially based transmission method, we examine how close the eavesdroppers can be to the intended user before the efficacy of STROBE begins to diminish. Specifically, the motivation for this experiment comes from the observed correlation in § IV between the proximity of the eavesdroppers to the intended user and the performance of beamforming based schemes. The results for that experiment (Fig. 3) showed an increased overheard SINR at E_3 for SUBF and STROBE (although only a slight increase for the latter).

In [10], the authors have shown that separation distance between receivers has a negligible effect on the served SINR when using a ZFBF based transmission scheme (such as CE). While we expect CE to cause low inter-user interference because the AP has knowledge of all users' channels, the efficacy of STROBE is unclear in this situation. We do not expect STROBE to match the intra-user interference reduction performance of CE (because STROBE only has the intended user's channel information); however, we do expect the blinding streams to compensate for this by overwhelming the overheard signal to the point where the overheard SINR is similarly minimal.

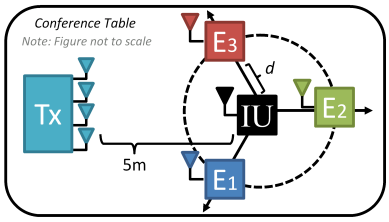

Fig. 4. Eavesdropper Proximity Topology.

1) Experimental Setup: We evaluate the effect of eavesdropper proximity on STROBE by placing the intended user at a fixed, 5 m distance from the transmitter with a direct line-of-sight (LOS) path and surround it by a circle of three eavesdroppers as shown in Fig. 4. For each measurement, we vary the radius of this circle (d) and express the distance in terms of λ. We place the eavesdroppers from a distance of 10λ (the separation distance for § IV's experiment) to $\lambda/4$ (the closest we can physically place the antennas together).

2) Experimental Results: Fig. 5 shows the SINR at the intended user and the three eavesdroppers at varying proximity distances for each transmission scheme.

Omnidirectional Transmission. As similarly observed in § IV, the eavesdroppers' overheard signal from the Omnidirectional transmission is relatively high as seen in Fig. 5(a). The only scenario in which the Omnidirectional scheme results in overheard signal energies substantially lower than the intended user is a combination of increased transmitter to eavesdropper distance and an obstructed eavesdropper LOS (E_2 at $d \geq 5\lambda$). The unpredictable (yet consistent[2]) behavior for all distances for this transmission scheme highlights the effect of multi-path signal propagation in indoor environments. The intended user's position at the center of the table and conference room (farthest away from walls, chairs, and other reflectors) allowed its SINR to remain consistent.

Single-User Beamforming. Similarly, indoor multi-path effects are observed in the received SINRs of the eavesdroppers when transmitting with SUBF (Fig. 5(b)). In fact, because the transmitted energy is being actively focused, a greater variation of overheard SINR occurs at the eavesdroppers. For example, the position difference from $\lambda/2$ to $\lambda/4$ at E_2 results in an 18 dB drop in signal strength even though the relative distance to the transmitter remains similar. The combination of a focused transmission and apparent multi-path randomness occasionally

[2]Single standard deviation confidence intervals are shown per data point.

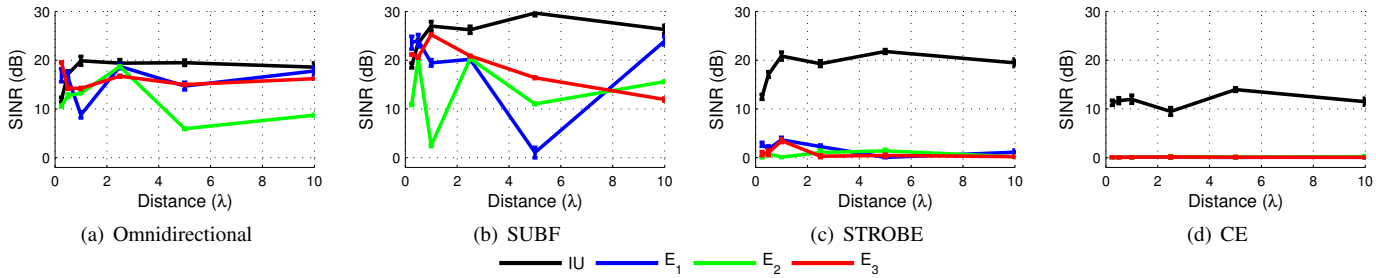

(a) Omnidirectional (b) SUBF (c) STROBE (d) CE

Fig. 5. SINR of Transmission to IU at varying λs for different transmission methods.

helps SUBF reduce the overheard signal energy (such as at E_1 for $d = 5\lambda$), but this accidental nulling is not deterministic.

STROBE. Regardless of eavesdropper proximity, STROBE's ability to consistently blind the eavesdroppers while still serving the intended user is shown in Fig. 5(c). The multi-path effects on the transmitted signal that cause the high variation in eavesdropper performance when using Omnidirectional or SUBF schemes have the opposite effect on STROBE. The overheard SINR range of a blinded eavesdropper from STROBE shown in Fig. 5(c) is 5 dB whereas Omnidirectional and SUBF transmission's ranges are 14 and 24 dB respectively. The ability of the multi-stream methods to separate receivers regardless of their relative distances observed in Fig. 5(c) and 5(d) matches the results shown in [10].

The only separation distance with an appreciable loss of SINR to the intended user for STROBE is $\lambda/4$ (12 dB); at all other proximity distances, the intended user is consistently sent a 20 dB signal. However, considering that this proximity distance is physically the closest our test antennas could be placed (the antenna bases were adjacent), the 12 dB SINR at the intended user and 10 dB SINR gain over the eavesdropper shows promise for STROBE. In fact, this result at a proximity distance of 3.125 cm ($\lambda/4$) implies that STROBE could potentially protect against covert eavesdropping devices secretly attached to the *intended user device itself*.

Cooperating Eavesdropper. Differences in eavesdropper blinding abilities between STROBE and CE confirm the findings of § IV. As detailed in § IV-C, full knowledge of the eavesdroppers' channels allows for the complete blinding of the eavesdroppers as shown in Fig. 5(d) (the eavesdroppers' lines are on top of one another). However, this precision comes at the cost of an SINR decrease for the intended user of 10 dB below STROBE. Although, CE's ability to serve the intended user remains constant even at $\lambda/4$, at that proximity distance, STROBE still serves the intended user with a stronger signal.

B. Inline Eavesdropper

In this section we evaluate the effect of eavesdroppers inline with the intended user. The goal is to quantify the effects of eavesdroppers blocking and placed along the LOS path from the transmitter to the intended user.

We expect indoor multi-path effects to aid STROBE in blinding the eavesdroppers and serving the intended user as hypothesized in § V-A. However, the major component of any transmitted signal is the LOS path so the potential for a beamforming based scheme to select this path and inadvertently serve an eavesdropper exists.

1) Experimental Setup: To evaluate the effect of eavesdroppers inline with the intended user, we set the transmitter a fixed distance (3 m) away from a line of receivers as shown in Fig. 6. We perform four iterations of the experiment setting each receiving node as the intended user and the remaining as eavesdroppers. Although the four iterations produce similar results, the topology shown in Fig. 6 results in the worst case performance for STROBE, which we present and analyze.

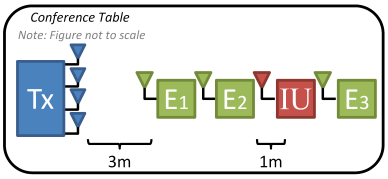

Fig. 6. Inline Eavesdropper Topology.

2) Experimental Results: Fig. 7 shows the received SINRs at the intended user and the inline eavesdroppers.

Omnidirectional Transmission. As observed in previous experiments, the Omnidirectional transmission's received SINR at the three eavesdropper positions is similar to the intended user. However, the received SINR at E_1 is the lowest although it is located closest to the transmitter with nothing blocking its LOS. This result and the similar deficits in E_2's SINR offer further examples of multi-path in an indoor environment.

Single-User Beamforming. The received SINR from the SUBF transmission at $E_{1\text{-}2}$ surpasses the SINR at the intended user. Unlike the Omnidirectional transmission, SUBF focuses energy toward the intended user and takes the direct LOS path. This results in the highest SINR at E_1 followed by E_2 and then by the intended user (the exact order in which they are located). This result shows that the best way to intercept a signal transmitted using SUBF is to simply eavesdrop in the LOS path of the intended user.

STROBE. STROBE serves the intended user with an SINR of 17 dB while allowing an eavesdropper to overhear, at most, an SINR of 4 dB (E_2). Unlike SUBF, STROBE handles the inline, LOS blocking eavesdroppers effectively blinding them even when given their positions. As previously stated, this intended user location did result in the worst case results for STROBE but even so, STROBE leverages multi-path and provides the intended user with 13 dB gain over the eavesdropper.

Cooperating Eavesdropper. As seen previously, CE precisely blinds the eavesdroppers. None of the intended user's signal energy is overheard while the intended user is served a 10 dB signal. Although the separation distances between the nodes were similar between this experiment and § IV, the SINR difference at the intended user between STROBE and CE is twice as much (3 to 6 dB). The increased difficulty in compensating for inline eavesdroppers blocking LOS paths causes a greater hit in the intended user's SINR when CE attempts to precisely blind the eavesdroppers. Even if using

this unrealistic scheme, the served SINR and relative gain over the eavesdropper SINR is still below that of STROBE.

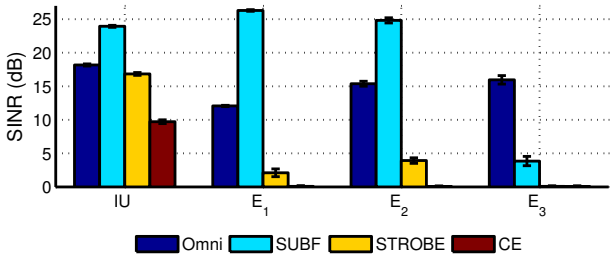

Fig. 7. SINR at IU and $E_{1\text{-}3}$ for inline receivers.

VI. Is Multi-Path Essential?

STROBE's efficacy relies on multi-path effects in an indoor environment as described in § V. The authors of [11] experimentally show that receiver separation distances of 70 m are required to serve users in parallel with ZFBF in outdoor environments. Multi-path is the hypothesized explanation for the ability of STROBE (along with CE) to function successfully regardless of eavesdropper proximity, relative position, or location. If this assumption can be validated, we can expose another benefit of STROBE. Increased multi-path effects are caused by "busier" environments (i.e. more physical obstacles). The "busier" an environment is, the larger the possibility for there to be eavesdroppers attempting to intercept an intended user's signal. Thus, if STROBE benefits from multi-path rich environments that support more eavesdroppers, such environments may actually help STROBE in securing a wireless transmission.

A. Experimental Setup

In order to evaluate the effects of decreased multi-path on STROBE, we redo the experiment described in § IV in an open space outdoors at considerable distance from buildings and other obstacles. The topology and relative distances between the nodes are identical to Fig. 2. Again, we perform four experiments setting each receiver as the intended user and the other three as eavesdroppers. All experiments produced similar results and hence for a direct comparison, we use the intended user location shown in Fig. 2.

B. Experimental Results

Fig. 8 shows the resulting SINRs when the transmitter sends a signal to the intended user in an open outdoor environment. The performance of Omnidirectional and SUBF transmissions are similar to the results from other indoor topologies.

However, the recreation of the initial experiment in an environment with far fewer multi-path effects results in drastic changes for the multi-stream methods. Observe the 2 dB served SINR when using CE to the intended user indicating the absolute failure of this multi-stream method. CE relies on its ability to separate the receivers channels in order to serve the intended user and precisely blind the eavesdroppers. Without multi-path, it is unable to do so.

In contrast, the served SINR at the intended user when using STROBE is almost 19 dB but the blinding abilities of STROBE completely fail at E_3 where a 13 dB signal is overheard. Recall from Fig. 2 that E_3 is located in front of the intended user and $E_{1\text{-}2}$ together on the opposite side. The eavesdropper SINR at $E_{1\text{-}2}$ is approximately 0 dB indicating that, without multi-path, STROBE is very susceptible to relative eavesdropper position and separation distance. Other results from setting the intended user at the different receiver positions confirm that without multi-path, STROBE becomes very directional and defeating the scheme simply requires approaching the intended user.

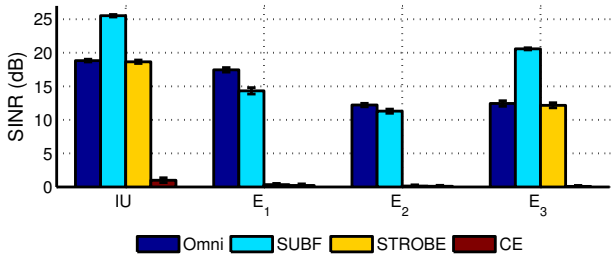

Fig. 8. SINR at IU over $E_{1\text{-}3}$ in outdoor environment.

VII. The Nomadic Eavesdropper

In this section we consider a nomadic eavesdropper that traverses throughout an indoor environment looking for the most opportune eavesdropping location. Previous experiments have demonstrated the wide variations in channel state from one position to the next due to multi-path effects. This randomness could permit a determined eavesdropper to exhaustively search an environment looking for such an opportune location.

A. Experimental Setup

1) Topology: To evaluate the potential of a location-based brute-force attack, we construct the topology shown in Fig. 9 in a large classroom (where each circle represents a seat). The classroom is filled with tables, chairs, and other objects that contribute to the rich multi-path characteristics of the environment. We place the transmitter at the front of the room and the intended user almost 6 m away on a direct LOS path. We transmit data to the intended user while placing the eavesdroppers at 24 different locations. The variety of different locations emulate the behavior of a determined eavesdropper searching for the optimal overhearing location.

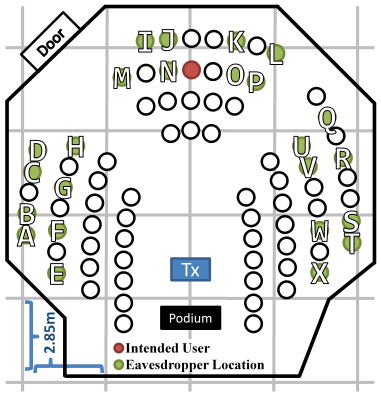

Fig. 9. Classroom Environment.

2) Compared Schemes: Unlike previous experiments, we compare the performance of STROBE against a directional antenna instead of the CE scheme. As described in § IV-C, CE is only capable of blinding three eavesdroppers (one less than the number of transmit antennas) and for this topology there

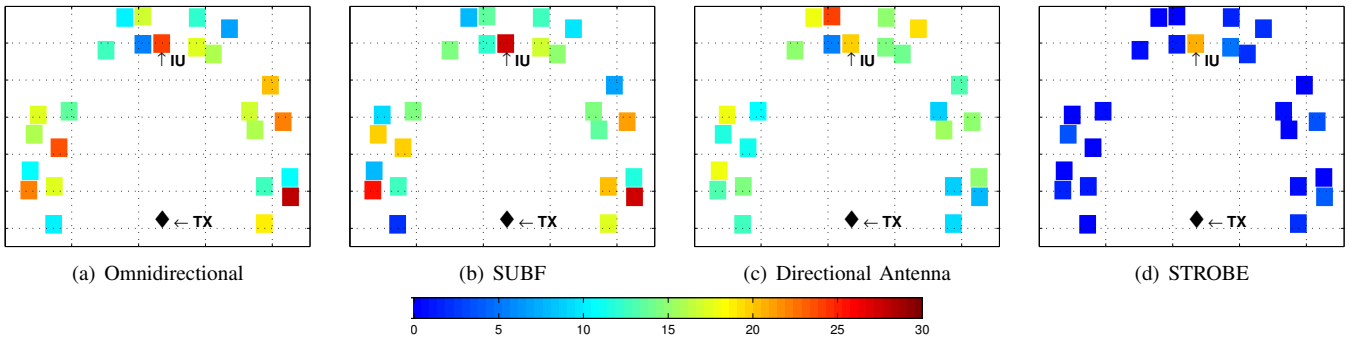

(a) Omnidirectional (b) SUBF (c) Directional Antenna (d) STROBE

Fig. 10. Overheard SINR(dB) at Eavesdropper Locations for different transmission methods.

are 24. Additionally, the simplest way to focus signal energy in a particular direction is to use a directional antenna. Regardless of the antenna's transmission angle, a directionally based transmission should put energy toward a particular location but not elsewhere. This makes such an antenna a promising candidate for directionally motivated security.

The antenna used is a Trendnet TEW-AO09D, a 9 dBi antenna with a 60^o radiation pattern. Although the width of the beam pattern precludes the possibility of perfectly removing overheard signal energy, the beam's shape suggests that regions of the environment can be spared from leaked signal strength.

B. Experimental Results

The results shown in Fig. 10 are presented as maps of received signal energy where colors correspond to the different received signal strengths. Dark blues represent lower SINR whereas dark reds represent high SINRs. The maps are to the scale of Fig. 9 with intended user and eavesdropper circles corresponding to the separate squares on the maps and all other lettered locations corresponding to eavesdropper locations.

Omnidirectional Transmission. As shown previously, observe in Fig. 10(a) the effect of multi-path randomness when agnostically transmitting energy with an Omnidirectional scheme. Note that the lowest received SINR for this scheme (location N, 5 dB) is located in the seat next to the intended user who received an SINR of 25 dB. Additionally, observe the wide color variation in eavesdropper SINR with regards to location. There is no correlation between distance from the transmitter and received SINR. However, other than the low measurements at locations N and L, all other eavesdroppers were able to overhear a signal of at least 10 dB while the intended user receives a signal of 24 dB.

The maximum received SINR from the Omnidirectional transmission is at location T (27 dB) while the average overheard SINR by the eavesdroppers is 16 dB. Thus, even with the inherent randomness of signal strength due to multi-path, it is relatively easy for an eavesdropper to find the opportune location to overhear an Omnidirectional transmission.

Single-User Beamforming. As shown in Fig. 10(b), SUBF serves the highest SINR to the intended user (27 dB) out of all schemes. Although the energy is focused toward a particular receiver, the high eavesdropper SINRs are located in places similar to that of the Omnidirectional transmission.

For example, location T overhears the signal with a 27 dB SINR, which is equal to the Omnidirectional transmission (and also equal to the intended user for this scheme). Locations R, G, and A also had similarly high received SINRs for both Omnidirectional and SUBF transmissions (20-25 dB). These locations receiving high SINRs relative to the intended user are understandable for an Omnidirectional transmission due to the shorter eavesdropper to transmitter distance. Although SUBF focuses energy toward the intended user, it still suffers from the generation of side lobes. When this effect is combined with multi-path, the result is a high overheard SINR at locations that are far away from the intended user.

Overall, the average overheard SINR for the SUBF transmission is 14 dB. This high energy average overheard signal combined with the lack of a correlation between eavesdropper location and overheard signal strength show that SUBF can be defeated by brute forcing different locations.

Directional Antenna. Unlike SUBF, the non-adaptive directional antenna transmission simply focuses energy where physically pointed. Although beamforming methods are aided by multi-path, the side effect is the potential for random signal reflections to increase SINRs at unintended locations (such as location T for SUBF). Observe in Fig. 10(c) the directional antenna's ability to passively focus energy in a particular direction allowing the directional antenna to better cope with multi-path induced randomness seen in previous schemes. Specifically, note that location T receives a strong signal reflection for Omnidirectional and SUBF schemes (27 dB) but a far weaker reflection for the directional antenna (8 dB). Other examples of this phenomenon include locations R, G, and A.

However, this ability does not make it immune to multi-path effects. The randomness caused by multi-path is simply constrained to the area where the antenna is aimed. Consider location N's received SINR of 6 dB. The intended user receives a 20 dB signal even though it is located in the adjacent seat. Additionally, location J receives the strongest signal overall (24 dB) by being behind the intended user and catching a stronger reflection from the back wall of the classroom.

Although the directional antenna reduces multi-path effects outside of its beampattern (sides of the classroom) it fails to do so where it is actually aimed. Additionally, the passive, directional transmission does not eliminate the overheard signal outside of its beampattern because of the constrained nature of the indoor environment. The average overheard SINR is still 13 dB and because there is a correlation between location and overheard SINR, it is feasible for an eavesdropper to move toward the intended user looking for favorable signal strength.

STROBE. Confirming our previous findings, STROBE successfully blinds eavesdroppers as depicted in Fig. 10(d). Observe that while the intended user receives a signal of 20 dB (shown as orange), all eavesdropper locations receive far less signal energy (all shades of blue). The maximum overheard

signal is only 5 dB at location O. Additionally, 60% of the overheard signal strengths are less than 1 dB. By employing orthogonal blinding, STROBE successfully and consistently diminishes eavesdropper SINR regardless of location.

STROBE's ability to handle multi-path randomness is also pronounced. Not only does STROBE consistently blind eavesdroppers, it also handles irregular reflection locations (such as T) far better than Omnidirectional and SUBF transmissions. Although location T was the second strongest overheard location, the eavesdropper SINR was limited to 4 dB, far less than the 27 dB overheard SINR for the first two schemes.

The average overheard SINR for STROBE was only *1.3 dB*, showing that STROBE outperformed the other three schemes by 12-14 dB. These consistently low overheard SINRs, regardless of eavesdropper location, show that STROBE can easily withstand a determined eavesdropper attempting to search for an opportune eavesdropping location.

VIII. Related Work

Wireless security exploiting CSI-based secret. One method of guarding a transmitter from an eavesdropper is CSI-based secret sharing. For example, the authors of [12] directly use intended user CSI as a secret generation method while the authors of [13] use intended user CSI to actively disrupt OFDM subcarriers to confuse eavesdroppers. In contrast, our implementation uses the CSI to beamform a signal toward the intended user while simultaneously blinding eavesdroppers. Because our method's use of CSI is independent of the aforementioned works, the two methods are complementary.

Beamforming-based multiple AP cooperation. Beamforming schemes that rely on groups of cooperating APs have also been proposed to secure wireless networks. The authors of [14] propose a method of securing wireless communications using a collection of phased arrays working in tandem to serve an intended user with a data stream. Each AP provides partial information of the overall data transfer and relies on precise directionally-based transmissions with tight AP synchronization to ensure the intersection of all partial streams at a particular geographic location. Additionally, the authors of [15] propose a set of multiple AP methods that allow energy to be focused toward an intended user but away from an eavesdropper. Although the authors propose using adaptive array beamformers, the weight construction technique employed is directly based on the physical shape of the constructed beam. However, such techniques can be unpredictable in indoor environments as shown in [2] and verified by our experiments in § VII. Furthermore, both of these works propose schemes that require multiple APs and custom hardware while STROBE works with a single AP and is compatible with 802.11n/ac. STROBE accomplishes the same goal with a single transmitter by leveraging the power of multi-stream transmission methods.

Information theoretic multi-antenna security. There have been a number of information theoretic studies that examine the theoretical performance of multi-antenna based security methods. In particular, these works define the fundamental limits of secrecy capacity. For example, [16] proves that a non-zero rate of communication can be guaranteed to be secret for any eavesdropper position. In contrast, our focus is on protocol design and experimental evaluation with alternate schemes. Likewise, [17] explores how eavesdroppers can be thwarted by a cooperative communication scheme.

IX. Conclusion

We design and experimentally characterize STROBE, a method of enhancing wireless security using ZFBF. We implement STROBE in the WARPLab experimental platform and thoroughly experimentally evaluate this transmission technique against Omnidirectional, Single-User Beamforming, Directional Antenna, and Cooperating Eavesdropper schemes. We show that STROBE achieves a higher SINR difference between the intended user and the eavesdropper than the single-target schemes because of its ability to blind. Additionally, we show that STROBE achieves a higher SINR difference than the unrealistic Cooperating Eavesdropper scheme because STROBE maximizes the SINR at the intended user while blinding eavesdroppers. We also demonstrate the performance of STROBE in a variety of indoor, multi-path rich environments and show its efficacy regardless of eavesdropper proximity or obstruction from the transmitter. We verify that STROBE's performance is due in part to the presence of multi-path in indoor environments by observing STROBE's diminished efficacy outdoors. Finally, we show STROBE's resilience to the nomadic eavesdropper that traverses an environment continuously searching for an opportune eavesdropping location. STROBE is a minimally invasive, viable method of augmenting wireless security using existing wireless technologies.

References

[1] E. Tews and M. Beck, "Practical attacks against WEP and WPA," in *Proc. ACM WiSec*, 2009.

[2] M. Buettner, E. Anderson, G. Yee, D. Saha, A. Sheth, D. Sicker, and D. Grunwald, "A phased array antenna testbed for evaluating directionality in wireless networks," in *Proc. ACM MobiEval*, 2007.

[3] T. Yoo and A. Goldsmith, "On the optimality of multiantenna broadcast scheduling using zero-forcing beamforming," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 3, Mar. 2006.

[4] M. Costa, "Writing on dirty paper," *IEEE Transactions on Information Theory*, vol. 29, no. 3, May 1983.

[5] H. Weingarten, Y. Steinberg, and S. Shamai, "The capacity region of the Gaussian multiple-input multiple-output broadcast channel," *IEEE Transactions on Information Theory*, vol. 52, no. 9, Sept. 2006.

[6] A. Wiesel, Y. Eldar, and S. Shamai, "Zero-forcing precoding and generalized inverses," *IEEE Transactions on Signal Processing*, vol. 56, no. 9, Sept. 2008.

[7] S. Friedberg, A. Insel, and L. Spence, *Linear Algebra*, 4th ed. Prentice Hall, 2003.

[8] A. Maltsev, V. Pestretsov, R. Maslennikov, and A. Khoryaev, "Triangular systolic array with reduced latency for QR-decomposition of complex matrices," in *Proc. IEEE ISCAS*, 2006.

[9] "Rice University WARP project," available at: http://warp.rice.edu.

[10] E. Aryafar, N. Anand, T. Salonidis, and E. Knightly, "Design and experimental evaluation of multi-user beamforming in Wireless LANs," in *Proc. ACM MobiCom*, Chicago, Illinois, Sept. 2010.

[11] F. Kaltenberger, M. Kountouris, D. Gesbert, and R. Knopp, "On the trade-off between feedback and capacity in measured MU-MIMO channels," *IEEE Transactions on Communications*, vol. 8, no. 9, Sept. 2009.

[12] Z. Li, W. Xu, R. Miller, and W. Trappe, "Securing wireless systems via lower layer enforcements," in *Proc. ACM WiSe*, 2006.

[13] S. Gollakota and D. Katabi, "Physical layer wireless security made fast and channel-independent," in *Proc. IEEE INFOCOM*, Shanghai, China, Apr. 2011.

[14] J. Carey and D. Grunwald, "Enhancing WLAN security with smart antennas: a physical layer response for information assurance," in *Proc. IEEE VTC*, Sept. 2004.

[15] S. Lakshmanan, C. Tsao, R. Sivakumar, and K. Sundaresan, "Securing wireless data networks against eavesdropping using smart antennas," in *Proc. IEEE ICDCS*, Beijing, China, Jun. 2008.

[16] S. Goel and R. Negi, "Guaranteeing secrecy using artificial noise," *IEEE Transactions on Communications*, vol. 7, no. 6, Jun. 2008.

[17] L. Dong, Z. Han, A. Petropulu, and V. Poor, "Improving wireless physical layer security via cooperating relays," *IEEE Transactions on Signal Processing*, vol. 58, no. 3, Mar. 2010.