# Vulnerability and Protection of Channel State Information in Multiuser MIMO Networks

Yu-Chih Tung, Sihui Han, Dongyao Chen, and Kang G. Shin
The University of Michigan
Email: {yctung,sihuihan,chendy,kgshin}@umich.edu

## ABSTRACT

Multiple-In-Multiple-Out (MIMO) offers great potential for increasing network capacity by exploiting spatial diversity with multiple antennas. Multiuser MIMO (MU-MIMO) further enables Access Points (APs) with multiple antennas to transmit multiple data streams concurrently to several clients. In MU-MIMO, clients need to estimate Channel State Information (CSI) and report it to APs in order to eliminate interference between them. We explore the vulnerability in clients' plaintext feedback of estimated CSI to the APs and propose two advanced attacks that malicious clients can mount by reporting forged CSI: (1) *sniffing attack* that enables concurrently transmitting malicious clients to eavesdrop other ongoing transmissions; (2) *power attack* that enables malicious clients to enhance their own capacity at the expense of others'. We have implemented and evaluated these two attacks in a WARP testbed. Based on our experimental results, we suggest a revision of the current CSI feedback scheme and propose a novel CSI feedback system, called the `CSIsec`, to prevent CSI forging without requiring any modification at the client side, thus facilitating its deployment.

## Categories and Subject Descriptors

C.2.0 [**Computer-Communication Networks**]: Security and Protection

## Keywords

Multiuser MIMO Networks; Physical Security; Channel State Information

## 1. INTRODUCTION

Multiple-Input Multiple-Output (MIMO) has the potential for solving the problem of insufficient bandwidth in WLANs, and has already been included in several wireless standards, such as IEEE 802.11n/ac [6, 45]. By exploiting spatial diversity, a transmitter with multiple antennas can either use its antennas to transmit the same stream to achieve power gain for the enhancement of the receiver's SNR or enable concurrent transmission of multiple (different) streams to achieve multiplexing gain. Moreover, to realize the potential benefit of multiple antennas, 802.11ac supports the multiuser MIMO (MU-MIMO) mode, in which multiple clients can be served concurrently without interfering with each other, achieving multiplexing gain even for those clients with a single antenna.

A key component in all MIMO technologies is Channel State Information (CSI). Due to the multipaths a wireless signal takes, the received signal is actually a combination of different, delayed and attenuated copies of each sent signal, and CSI can be regarded as the coefficient of this combination. This information is critical to MIMO networks since it is used by transmitters to precode signals either for boosting received signal strength at a client or removing interference to the other clients. Researchers have also shown that the network capacity can be improved via transmission with known CSI [38]. Note that only receivers know the CSI from transmitters to themselves, which is estimated by using a pre-defined known sequence. For example, in downlink transmission, only clients know their own CSI from the Access Point (AP) to themselves, which is estimated by dividing the received signal by the known sequence, and the clients are responsible for feeding back this information to the AP for precoding messages as shown in Fig. 1(a). Moreover, the freshness of CSI is critical to the performance of wireless networks, as stale CSI does not represent the current state of multi-path fading and may result in precoding errors. In MU-MIMO networks, when the CSI feedback delay is greater than 200ms, the achievable network capacity decreases by more than 50% [4]. Thus, receivers are required to report CSI in plaintext as quickly as possible.

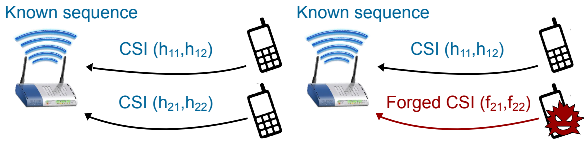

(a) Genuine CSI feedback flow (b) Forged CSI feedback flow

**Figure 1—Attack models based on forged CSI.**

We uncover the vulnerability caused by this plaintext CSI feedback, and propose two advanced attacks in MU-MIMO networks with forged CSI. As shown in Fig. 1(b), we assume malicious clients can first sniff other clients' CSI, modify/fabricate and then report it. In general, clients have no incentive to report a wrong estimation of CSI, as it does nothing but decreases its own received signal-to-interference-plus-noise ratio (SINR). However, in MU-MIMO networks, data can be transmitted to multiple clients concurrently who share the same wireless medium, and hence, clients can intentionally mislead the precoding process at the transmitter by reporting forged CSI for malicious purposes. Based on this observation, we first introduce a potential threat, called the *sniffing attack*, which enables a client to eavesdrop packets sent to others even under the protection of *physical security* which is theoretically proven to be immune to eavesdropping [12, 24]. The other potential threat we identified is the *power attack*, which manipulates the AP's power allocation to antennas based on the reported (forged) CSI. By reporting forged CSI, a client can receive a favorable power allocation for its own transmission. To the best of our knowledge, we are the first to discover, implement, and prove the possibility of malicious attacks on the MU-MIMO precoding process by reporting forged CSI. These attacks are very different from the others proposed before [14, 23, 32] — which also provide wrong metrics to fool systems — in that ours actively exploits the precoding procedure by forging CSI which is unique in MU-MIMO networks. The

*CCS'14,* November 3–7, 2014, Scottsdale, Arizona, USA.
Copyright 2014 ACM 978-1-4503-2957-6/14/11 ...$15.00.
http://dx.doi.org/10.1145/2660267.2660270.

vulnerabilities caused by forging CSI are expected to become a critical issue as MU-MIMO is becoming popular and deployed widely; it is thus important to prevent such attacks before they become rampant.

We propose `CSIsec` to protect existing MU-MIMO networks from the attacks with forged CSI. `CSIsec` is a novel CSI feedback system in which transmitters send a "cheated" known sequence instead of the genuine known sequence to mislead the CSI estimation process at clients before they forge CSI. Using this approach, no clients can estimate their own CSI correctly, and it is also impossible to know CSI of another client because even that client doesn't know its own CSI. Under `CSIsec` the transmitter is the only one who knows the CSI estimated at clients, and this information is acquired through a process similar to Diffie-Hellman key exchange [10]. It is important to note that `CSIsec` requires no modification at clients, and hence it provides backward compatibility and can be easily applied to existing systems. We also suggest adding randomization in existing MU-MIMO mechanisms to disincentivize clients from reporting forged feedback under `CSIsec`.

We implement the attack models and `CSIsec` in a WARP testbed [20], and validate that by reporting forged CSI, a client can successfully decode packets destined for other clients with less than 2% bit-error-rate (BER) on average even with physical security enabled. On the other hand, a client can also acquire an unfairly higher capacity (by 20%) only by reporting forged CSI.

This paper makes the following three main contributions:

- The first exploration of potential threats in MU-MIMO precoding process caused by reporting forged CSI;
- Implementation and proof of the possibility of sniffing and power attacks in a real testbed; and
- Development of `CSIsec`, a novel CSI feedback system that does not require modifications of clients but can prevent clients from forging CSI.

The remainder of this paper is organized as follows. Section 2 briefly introduces the principles of physical security in MU-MIMO networks and presents our attack model. Two plausible attacks are illustrated in Sections 3 and 4, respectively. Section 5 describes our implementation setting and presents our experimental evaluation. The countermeasures are discussed in Section 6 while Section 7 summarizes related work. We discuss future directions and conclude the paper in Sections 8 and 9.

## 2. Background and System Model

MU-MIMO networks [6] are an emerging communication technology for next-generation wireless communications thanks to their potential for enhancing receivers' capacity, even when each receiver is equipped with a single antenna. Security in MU-MIMO is a must since multiple receivers/clients are served concurrently, implying that malicious behavior of even a single client can affect the transmission of all others. Instead of ensuring security by traditional cryptosystems, *physical security* has been studied widely to thwart eavesdroppers [3, 12, 13] because of its guarantee of high security at a relatively low cost in MIMO networks [9]. In this section, we first introduce basic MU-MIMO techniques and the state-of-art physical layer security in MU-MIMO networks, and then illustrate our attack model.

In what follows, we use upper-case boldface letters to represent matrices while using lower-case boldface for vectors. $\mathbf{X}^T$ stands for the transpose of $\mathbf{X}$ and $\mathbf{X}^*$ for the conjugate transpose of $\mathbf{X}$. $||\mathbf{x_k}||$ represents the norm of a vector $\mathbf{x_k}$ and $|x_{ij}|$ represents the absolute value of a matrix element $x_{ij}$.

### 2.1 Beamforming in MU-MIMO

Consider a MU-MIMO system with one transmitter (the AP) with $N$ antennas, and $M$ receivers (clients), each with a single antenna. The downlink CSI from the transmitter's $j$-th antenna to the $i$-th receiver is characterized by a single frequency-domain complex coefficient $h_{ij}$. Therefore, the full CSI can be represented by an $M \times N$ matrix $\mathbf{H} = [\mathbf{h_1}^T, \mathbf{h_2}^T, \dots, \mathbf{h_M}^T]^T$ where the $i$-th row vector $\mathbf{h_i}$ represents the CSI of the link from the transmitter's $N$ antennas to the $i$-th receiver. In 802.11n/ac, data is modulated into different subcarriers based on Orthogonal Frequency-Division Multiplexing (OFDM) in which each individual $\mathbf{H_k}$ is used to represent the CSI of the $k$-th subcarrier. To keep the model succinct, we ignore the subscript $k$, and it can be easily extended to the multiple-subcarrier case by treating each $\mathbf{H_k}$ independently. Thus, the received signal $\mathbf{y}$ of transmitted signal $\mathbf{x}$ can be expressed as:

$$\mathbf{y} = \mathbf{Hx} + \mathbf{n}, \tag{1}$$

where the $N \times 1$ vector $\mathbf{x} = [x_1, x_2, \dots, x_N]^T$ represents the signals sent from the transmitter's $N$ antennas, the $M \times 1$ vector $\mathbf{y} = [y_1, y_2, \dots, y_M]^T$ represents the signals received at the $M$ concurrent receivers, and $\mathbf{n} = [n_1, n_2, \dots, n_M]^T$ represents an additive white Gaussian noise with variances $\sigma_1^2, \sigma_2^2, \dots, \sigma_M^2$.

Unlike single-user MIMO (SU-MIMO), the $i$-th receiver in a MU-MIMO system doesn't have knowledge of signals $y_{k,k \neq i}$ received at other receivers, so the received data cannot be jointly decoded. Thus, a precoding phase at the transmitter before sending the signal is necessary in MU-MIMO networks. Researchers have proven the optimality of Dirty Paper Coding (DPC) [8, 41] in MU-MIMO. However, the implementation of DPC incurs significant additional complexity, making it unsuitable for wireless protocols. Thus, we will instead focus on linear precoding schemes like zero-forcing beamforming (ZF-BF), where the received signals are expressed as:

$$\mathbf{y} = \mathbf{Hx} + \mathbf{n} = \mathbf{HCm} + \mathbf{n} \tag{2}$$

where $\mathbf{m} = [m_1, m_2, \dots, m_M]^T$ is an $M \times 1$ vector representing the messages clients expect to receive. For example, $m_i$ is the message that the $i$-th receiver expects. $\mathbf{C}$ is the $N \times M$ precoding matrix where $\mathbf{c_k}$ represents the $k$-th column of matrix $\mathbf{C}$. In this scheme, $\mathbf{Cm}$ represents the precoded signals being sent from the transmitter's $N$ antennas and meets the power constraint $||\mathbf{Cm}|| < P$, where $P$ is the total transmit power. After precoding, the received SINR at the $i$-th receiver is:

$$SINR_{m_i} = log(\frac{||\mathbf{h_i c_i} m_i||^2}{\sigma_i^2 + \sum_{k \neq i} ||\mathbf{h_k c_k} m_k||^2}) \tag{3}$$

where $||\mathbf{h_i c_i} m_i||^2$ represents the magnitude of message $m_i$ that the $i$-th receiver expects to receive and $||\mathbf{h_k c_k} m_k||^2$ represents the interference caused by messages sent to other receivers in the same concurrent transmission. The main idea of ZF-BF is to nullify the interference caused by other concurrently transmitted messages, and channel inversion [12] is proven to be the optimal precoding matrix to ensure zero interference, i.e., $\mathbf{C} = \mathbf{H}^\dagger = \mathbf{H}^*(\mathbf{H}^*\mathbf{H})^{\mathbf{-1}}$ and

$$\mathbf{y} = \mathbf{HH}^*(\mathbf{H}^*\mathbf{H})^{\mathbf{-1}}\mathbf{m} + \mathbf{n} = \mathbf{m} + \mathbf{n}, \tag{4}$$

indicating that the $i$-th receiver can receive its own message $m_i$ without any interference from other concurrently transmitted messages. A $2 \times 2$ example of ZF-BF (i.e., $N = 2$, $M = 2$) is shown in Fig. 2, where the precoded message of $m_1$ can be visualized as a vector being steered along the direction orthogonal to $\mathbf{h_2}$, thus causing no interference at $rx_2$, i.e., $\mathbf{h_2}(\mathbf{c_1} m_1) = 0$.

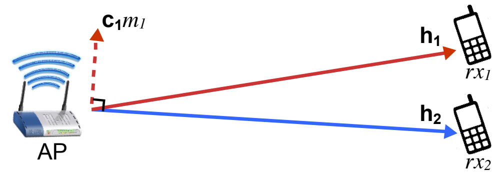

**Figure 2—Zero-forcing beamforming.** $rx_2$ receives zero interference from $m_1$ because the precoded $\mathbf{c_1}m_1$ is in the null space of $\mathbf{h_2}$. On the other hand, $m_1$ is decodable at $rx_1$ because $\mathbf{c_1}m_1$ is not orthogonal to $\mathbf{h_1}$

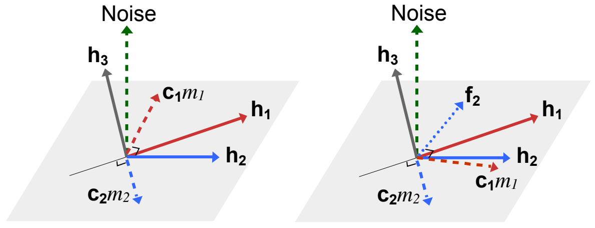

(a) Physical security in ZF-BF (b) When CSI is forged

**Figure 3—How forged CSI works.** Dashed lines represent the signals sent by transmit antennas and solid lines represent the CSI. When CSI is forged, the precoded signal $\mathbf{c_1}m_1$ is no longer orthogonal to $\mathbf{h_2}$.

## 2.2 Physical-Layer Security in MU-MIMO

As shown in Eq. (4), in ZF-BF, each receiver only receives the message sent to itself without knowledge of concurrently transmitted data to other receivers. Thus, it is proven to be not leaking information to other concurrent receivers in terms of physical security [12, 24]. To prevent other eavesdroppers from sniffing the concurrent transmission, artificial noise is introduced to mislead potential eavesdroppers [13]. Usually, the artificial noise is transmitted in the null space of $\mathbf{H}$ to avoid interference to concurrent receivers. It has been shown experimentally that artificial noise can ensure eavesdroppers to have 15dB lower SNR than the signal at the intended receiver if the eavesdroppers and the intended receiver are a half wireless wavelength apart [3]. A toy example of this physical security is shown in Fig. 3(a), where $\mathbf{h_1}$ and $\mathbf{h_2}$ are the CSI of $rx_1$ and $rx_2$ which are the intended receivers of messages $m_1$ and $m_2$, respectively, and $\mathbf{h_3}$ is the CSI of $rx_3$ which is not in the concurrent transmission. We assume $rx_2$ and $rx_3$ are both eavesdroppers who are trying to decode the message sent to $rx_1$ with physical security enabled. Since in the precoding phase, message $m_1$ is projected as $\mathbf{c_1}m_1$ in the space orthogonal to $h_2$, $rx_2$ cannot receive any information about $m_1$ because the projection of $\mathbf{c_1}m_1$ onto $\mathbf{h_2}$ is zero. Moreover, because an artificial noise is transmitted, $rx_3$ is unlikely to be able to decode $m_1$ since the projection of sent artificial noise onto $\mathbf{h_3}$ is not 0. Thus, this example demonstrates the capability of physical security to thwart eavesdroppers.

Physical security like this has been proposed to thwart eavesdroppers *without encryption*. The overhead of keeping data confidential this way is minimal in MU-MIMO because all necessary computations are done in wireless communication chips, and no decoding procedure is required at clients. Some researchers even claim that this protocol can provide better security due to the diversity of CSI [9]. However, all of these assume perfect CSI fed back by clients, and we will next introduce a threat model based on this assumption.

## 2.3 General Attack Model

The proposed attacks are demonstrated in a 2 × 2 MU-MIMO system as shown in Fig. 2, consisting of one 2-antenna AP and two 1-antenna clients for simplicity. We focus on the downlink transmission because downlink CSI is unknown to the transmitter and it needs to be fed back by the receivers. Here, we focus on explicit CSI feedback since implicit CSI feedback relying on channel reciprocity has been abandoned in 802.11ac. Moreover, we assume malicious clients are able to report forged CSI, and the packets precoded by ZF-BF are concurrently sent to two clients, $rx_1$ and $rx_2$. In the scenario of physical security, transmitters use additional antennas to send artificial noise, and there exists one additional eavesdropper, $rx_3$, not in the concurrent transmission as shown in Fig. 3. The setting can easily be extended to more complicated systems with more antennas. Under this setting, the 2 × 2 MU-MIMO system model can be represented as:

$$\begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \begin{pmatrix} h_{11} & h_{12} \\ h_{21} & h_{22} \end{pmatrix} \begin{pmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \end{pmatrix} \begin{pmatrix} \sqrt{p_1}m_1 \\ \sqrt{p_2}m_2 \end{pmatrix} + \begin{pmatrix} n_1 \\ n_2 \end{pmatrix}, \quad (5)$$

where $y_1$, $y_2$ are their received signals at receivers; $m_1$, $m_2$ are the messages sent for $rx_1$ and $rx_2$; and $h_{ij}$ represents the channel coefficient (CSI) from the $j$-th transmitter antenna to the $i$-th receiver. Since the transmitter precodes the sent message based on ZF-BF, the precoding matrix $\mathbf{C}$ is set to $\mathbf{H^{-1}}$, and thus the received message at $rx_i$ is $y_i = \sqrt{p_i}m_i + n_i$ where $p_i$ represents the power coefficient used to control the magnitude of sent message $m_i$ where messages are assumed as constant-modulus signals, i.e., $|m_i| = 1$. Since the AP's total power is fixed, the transmit power allocation must satisfy: $\sum_{i=1}^{N} \|\mathbf{c_i}\|^2 p_i < P$.

Without loss of generality, we assume the second client, $rx_2$, is malicious and reports the forged CSI, $\mathbf{f_2} = [f_{21}, f_{22}]$, instead of the genuine CSI, $\mathbf{h_2} = [h_{21}, h_{22}]$. In this case, the channel matrix perceived at transmitters will be $\mathbf{F} = \begin{pmatrix} h_{11} & h_{12} \\ f_{21} & f_{22} \end{pmatrix}$ instead of $\mathbf{H}$. Using ZF-BF, the received symbol will become:

$$\begin{aligned} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} &= \begin{pmatrix} h_{11} & h_{12} \\ h_{21} & h_{22} \end{pmatrix} \begin{pmatrix} h_{11} & h_{12} \\ f_{21} & f_{22} \end{pmatrix}^{-1} \begin{pmatrix} \sqrt{p_1}m_1 \\ \sqrt{p_2}m_2 \end{pmatrix} + \begin{pmatrix} n_1 \\ n_2 \end{pmatrix} \\ &= \begin{pmatrix} 1 & 0 \\ \frac{h_{21}f_{22} - h_{22}f_{21}}{h_{11}f_{22} - f_{21}h_{12}} & \frac{h_{11}h_{22} - h_{12}h_{21}}{h_{11}f_{22} - f_{21}h_{12}} \end{pmatrix} \begin{pmatrix} \sqrt{p_1}m_1 \\ \sqrt{p_2}m_2 \end{pmatrix} + \begin{pmatrix} n_1 \\ n_2 \end{pmatrix} \end{aligned} \quad (6)$$

The first insight in this equation is that no matter how $rx_2$ forges and reports the CSI, $rx_1$ always receives the signal $y_1 = \sqrt{p_1}m_1 + n_1$. That is, it is impossible for $rx_2$ to inject any payload into $y_1$ by forging CSI, i.e., misleading the decoded signals at $rx_1$. However, comparing Eqs. (4) and (6), one can easily see the received signals at $rx_2$ change when CSI is forged, making it possible to mount the proposed two types of attack. For example, the information of $m_1$ is leaked to $rx_2$ with the magnitude of $\frac{h_{21}f_{22} - h_{22}f_{21}}{h_{11}f_{22} - f_{21}h_{12}}$ because the CSI is falsely reported as $\mathbf{f_2}$, instead of $\mathbf{h_2}$. The result of reporting forged CSI can also be visualized as in Fig. 3(b), where the AP falsely precodes the direction of message $m_1$ to the direction that is not orthogonal to $\mathbf{h_2}$. We will introduce how the attackers forge well-designed CSI and report it to the AP for malicious purposes in the following sections.

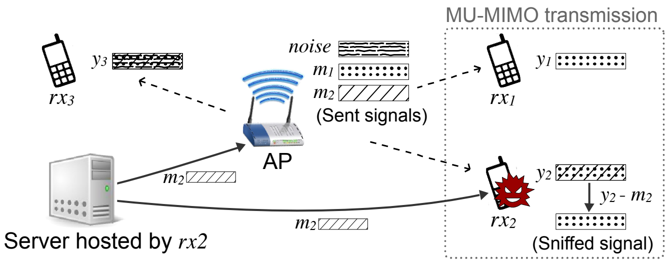

**Figure 4—Flow of the sniffing attack.** $rx_2$ receives the leaked information of $m_1$ by misleading the precoding process at the AP with forged CSI. The interference caused by $m_2$ is further removed because $m_2$ is known to $rx_2$ by connecting to a server hosted by himself.

## 3. Sniffing Attack

We now introduce the sniffing attack against physical security by reporting forged CSI. Different from the traditional eavesdropping where sniffers passively wait for a chance to see and collect victims' data, in the sniffing attack, eavesdroppers will actively forge and report CSI. We first introduce a way for eavesdroppers to decode others' packets in MU-MIMO networks with physical security enabled and then propose a heuristic algorithm that efficiently realizes this attack.

### 3.1 Decoding the Sniffed Packets

As shown in Eq. (6), when the forged CSI is reported, the received signal at the eavesdropper, $y_2$, contains a mixture of signals from $m_1$ and $m_2$, but both of them are not decodable at $rx_2$ because of their mutual interference. However, the purpose of the eavesdropper $rx_2$ is sniffing $m_1$ rather than receiving its own message $m_2$, and hence it is reasonable to assume $m_2$ is already known to $rx_2$ without loss of generality. When $m_2$ is known to the eavesdropper, $rx_2$ can decode $m_1$ via interference cancellation [15]. That is, $rx_2$ first removes the interference caused by $m_2$ from $y_2$, and then decodes $m_1$ from the remaining signals as:

$$\begin{aligned} m_1^{rx_2} &= \frac{(h_{11}f_{22} - f_{21}h_{12})}{\sqrt{p_1}(h_{21}f_{22} - h_{22}f_{21})}(y_2 - \frac{h_{11}h_{22} - h_{12}h_{21}}{h_{11}f_{22} - f_{21}h_{12}}\sqrt{p_2}m_2) \\ &= m_1 + \frac{(h_{11}f_{22} - f_{21}h_{12})}{\sqrt{p_1}(h_{21}f_{22} - h_{22}f_{21})}n_2 \end{aligned} \tag{7}$$

where all components, except $m_1$, in this equation are known to $rx_2$, and the sniffed SNR is controlled by the forged CSI, i.e., $f_{21}$ and $f_{22}$. One way for $rx_2$ to know $m_2$ for signal cancellation is to download the same message from a server maintained by themselves as shown in Fig. 4. We have validated this trick on several machines by socket programs connected through the AP.

### 3.2 Selection of Forged CSI

As shown in Eq. (7), an intuitive selection of forged CSI for the eavesdropper is $(f_{21}, f_{22})$ that maximizes the sniffed SNR, i.e., minimizing $\frac{h_{11}f_{22}-f_{21}h_{12}}{\sqrt{p_1}(h_{21}f_{22}-h_{22}f_{21})}$. However, this intuition is valid only if the interference from $m_2$ can be completely removed and there is no interference caused by artificial noise. In the process of removing $m_2$, higher $\frac{\sqrt{p_2}(h_{11}h_{22}-h_{12}h_{21})}{h_{11}f_{22}-f_{21}h_{12}}$ incurs a higher residual interference due to imperfect signal cancellation. Moreover, if the null space of the forged CSI is not the same as $\mathbf{h_1}$ and $\mathbf{h_2}$, the received signal will be interfered with by the artificial noise as shown in Fig. 3. Based on these two observations, we propose an efficient heuristic to select the forged CSI. That is, the forged CSI is selected as a weighted sum of genuine CSI as:

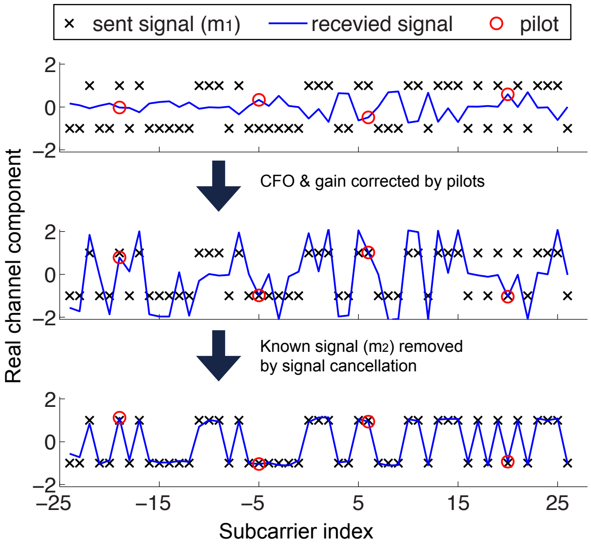

**Figure 5—Example of signal cancellation of sniffed signal.** Sniffed signal at $rx_2$ is first corrected by pilots, then the interference caused by messages sent to $rx_2$ is removed via signal cancellation.

$$[f_{12}, f_{22}] = [wh_{11} - h_{12}, wh_{21} - h_{22}], \tag{8}$$

where $w$ is a adjustable coefficient. Selecting forged CSI with this heuristic ensures the precoded message is not nullified by ZF-BF at $rx_2$ and the strength of leaked information at $rx_2$ is bounded. Applying this heuristic selection of CSI to Eq. (6), we can derive the received signal at the eavesdropper as:

$$y_2 = w\sqrt{p_1}m_1 + \sqrt{p_2}m_2 + n_2. \tag{9}$$

In this scenario, the AP is assumed to set power coefficients equally as $p_1 = p_2 = p$ and ensure the sent signals to meet the power constraint, i.e., $p(\|\mathbf{h_1}\|^2 + \|\mathbf{f_2}\|^2)/det(\mathbf{F})^2 \leq P$. As shown in this equation, if $m_2$ can be removed completely, then the sniffed SNR of message $m_1$ at $rx_2$ is proportional to $w\sqrt{p}$.

The main idea of this selection of forged CSI is that the null space of $w\mathbf{h_1} - \mathbf{h_2}$ is the same as that of $\mathbf{h_1}$ and $\mathbf{h_2}$, so the artificial noise sent in the null space still causes no interference to the received signals at $rx_2$. As shown in Fig. 4, $rx_2$ receives no interference from the artificial noise and can thus decode $m_1$ after removing the known message $m_2$. The effectiveness of this heuristic selection of forged CSI will be evaluated in Section 5. One thing to note in the implementation of this attack is that eavesdroppers should keep the reported CSI on pilot subcarriers intact because signals sent through those subcarriers are used to remove the central frequency offset (CFO) caused by imperfect clock synchronization between transmitters and receivers.

Fig. 5 illustrates an example of decoding the sniffed message at $rx_2$. As shown in the figure, the unprocessed signal at $rx_2$ seems unrelated to the target message $m_1$ due to the mixture of $m_1$ and $m_2$ plus lack of clock synchronization between transmitters and receivers. However, after correcting CFO and gain via intact pilot subcarriers and removing the component of $m_2$, the residual signal is found highly correlated to the target message $m_1$.

# 4. POWER ATTACK

Here we introduce another potential threat, called *power attack*, by exploiting forged CSI. As mentioned in Section 2, the total amount of power for an AP to transmit data via its antennas is fixed. Under this constraint, several power-allocation mechanisms have been proposed with different objectives. Of these, two most representative mechanisms are: (1) *equal power* (EP) allocation (i.e., $p_1 = p_2 = \ldots = p_N$) that maximizes fairness among concurrent receivers, and (2) *maximizing throughput* (MT) allocation (i.e., $argmax_{p_i} \sum_i^N log(1 + p_i/\sigma_i)$) that maximizes the aggregated capacity of concurrent receivers. Note that both mechanisms rely on CSI feedback from receivers. Malicious clients can thus unfairly boost their received SINR by misleading the precoding process with forged CSI. How to gain favorable SINR with forged CSI varies with the underlying power-allocation mechanism, which is actually implementation-specific in WiFi chips. Thus, without loss of generality, we will illustrate the concept of power attack against EP and MT allocations, which are most commonly adopted in MU-MIMO [4]. This attack can be extended to other power-allocation mechanisms as long as they rely on CSI feedback.

## 4.1 Decoded SINR at Malicious Clients

Suppose $rx_2$ is the malicious client, then according to Eq. (6), the decoded SINR of messages sent to $rx_2$ when forged CSI $\mathbf{f_2}$ is reported becomes:

$$SINR_{m_2} = log(\frac{p_1|\frac{h_{11}h_{22}-h_{12}h_{21}}{h_{11}f_{22}-f_{21}h_{12}}|^2}{p_2|\frac{h_{21}f_{22}-h_{22}f_{21}}{h_{11}f_{22}-f_{21}h_{12}}|^2 + \sigma_2^2}) \tag{10}$$

where $p_1|\frac{h_{11}h_{22}-h_{12}h_{21}}{h_{11}f_{22}-f_{21}h_{12}}|^2$ represents the received power of message $m_2$ and $p_2|\frac{h_{21}f_{22}-h_{22}f_{21}}{h_{11}f_{22}-f_{21}h_{12}}|^2$ is the power of interference from $m_1$. Note that the denominator of received power, i.e., $|h_{11}f_{22} - f_{21}h_{12}|^2$, is controlled by the forged CSI, $f_{21}$ and $f_{22}$. However, decreasing this dominator also increases the interference from $m_1$ because it has the same dominator. The optimal selection of forged CSI is to maximize $SINR(m_2)$ subject to the power constraint $P$, and this optimization problem is akin to that of finding the optimal regularization term in perturbed channel inversion [31]. Our simulation result of increased SINR in EP allocation due to the forged CSI is plotted in Fig. 6, where the forged CSI (i.e., $f_{21}$ and $f_{22}$) is selected as real numbers for easy visualization, and the original CSI at $rx_2$ is labeled with a cross in this figure. After several simulation runs with different parameter settings, we find the optimal selection of forged CSI lying in a similar direction as the original CSI, which is the direction causing no interference from concurrent transmitters. One possible explanation of this phenomena against the conclusion in [31] that perturbated precoding is optimal in terms of SINR is that the degrees of freedom in our optimization problem are not full. That is, unlike the optimization problem of perturbated precoding matrix which can modify any components in the precoding matrix, our optimization problem can only control the second row of precoding matrix, i.e., $f_{21}$ and $f_{22}$. Thus, based on this observation, instead of solving the optimization problem in real time, we propose an efficient heuristic to select proper CSI to gain a favorable power allocation to message $m_2$ for both EP and MT allocations.

## 4.2 Selection of Forged CSI in EP

To exploit the vulnerability of the AP's power allocation based on CSI feedback, one simple and effective heuristic is to report a scaled version of CSI which has the same direction as the genuine CSI. This idea is identical to ZF-BF which nullifies the interference caused by concurrent transmissions even though it has been proven in [31] not optimal in terms of received SINR. Using this approach, the forged CSI is selected as:

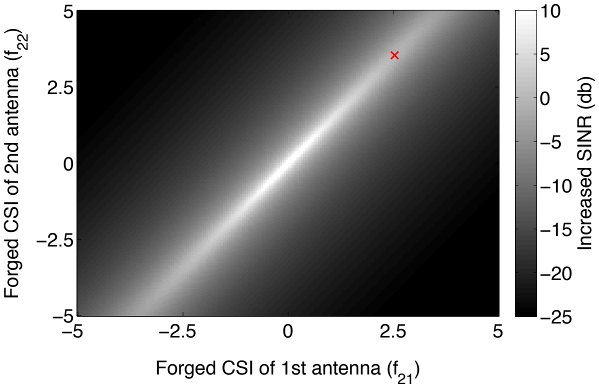

**Figure 6—Effectiveness of forged CSI selection.** The increased SINR (dB), i.e., SINR with forged CSI subtracted by SINR with the genuine CSI, is shown in gray scale, where the white color represents a 12dB increase and the black color represents a 27dB decrease.

$$[f_{21}, f_{22}] = w[h_{21}, h_{22}], \tag{11}$$

where $w$ is a constant within $[0, 1]$ in EP allocation to pretend $rx_2$ suffering from heavy channel fading and $w$ is set larger than 1 in MT allocation to pretend having high quality of $rx_2$'s channel.

Applying this forged CSI to Eq. (6), the received signal is simplified as:

$$y_2 = \frac{\sqrt{p_2}}{w} m_2 + n_2, \tag{12}$$

where $\frac{\sqrt{p_2}}{w} = \frac{h_{11}h_{22}-h_{12}h_{21}}{\sqrt{w^2(|h_{21}|^2+|h_{22}|^2)+|h_{11}|^2+|h_{12}|^2}}$ in the equal power (EP) allocation, because $p_1 = p_2 = det(\mathbf{F})^2 P/(\|\mathbf{h_1}\|^2 + \|\mathbf{f_2}\|^2)$.

This equation indicates the optimal selection of $w$ to be 0. However, this selection is not feasible because the AP might refuse to send data toward clients who claim their channel gain is 0. Moreover, higher received SINR of packets doesn't imply higher throughput because the throughput is related to the modulation and rate selection which are both controlled by transmitters. Rate adaption in MU-MIMO still remains to be an open question. Existing protocols either select modulation based on history or suggestions from receivers. In both of these schemes, forging CSI with different $w$ makes no difference because the two rate adaptations do not rely on the magnitude of reported CSI. However, there exists ongoing research on the estimation of SINR based on rate adaptation with reported CSI [35, 16]. In these schemes, a small $w$ will incur a lower rate assigned to transmit $m_2$ because the AP perceives the channel gain toward $rx_2$ to be small, i.e., $w < 1$. For example, forged CSI with $w = 0.3$ might increase the received SINR of $m_2$ from 10dB to 15dB, but it might decrease the assigned data rate from 11Mbps to 5.5Mbps, which eventually decreases nearly to a half of the throughput of malicious users. If these protocols are adopted in future MU-MIMO networks, one possible way to realize power attack is also to forge the reported noise estimation, because most of those protocols rely on reported noise estimation to calculate SINR at receivers. For example, when CSI is reported as $w\mathbf{h_2}$, $rx_2$ can also report its noise variance as $w\sigma_2$ to mislead the rate-adaptation protocol in the AP. The way to cheat on noise feedback for different rate-adaptation protocols is part of our future work, and in the following sections, the channel capacity defined as $log(1 + SINR)$ will be used to estimate the effectiveness of the proposed power

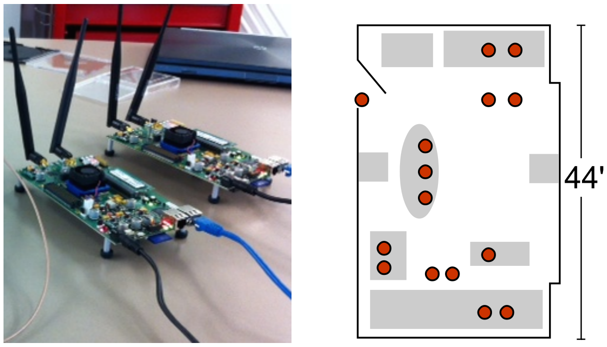

(a) WARP testbed (b) Testing locations

**Figure 7—Testing environment.** Experiments are conducted in an area representing a typical environment of WiFi transmission full of multi-paths. The red circles represent the position of WARP boards and the grey blocks represent the pillars, tables, racks, and other obstacles.

attack without loss of generality because it represents a theoretical upper-bound of transmission throughput in the real world.

### 4.3 Forged CSI Selection in MT

Analysis of forged CSI selection at MT allocation is more complex since the power allocation becomes a nonlinear optimization problem to maximize total capacity, i.e., $capacity(p_1,\dots,p_N) = \sum_{i=1}^{N} log(1+p_i/\sigma_{n_i})$, subject to the total power constraint.

This problem can be solved by a well-known waterfilling algorithm [46]:

$$p_i = \left(\frac{u}{\|\mathbf{c_i}\|^2} - \sigma_{n_i}\right)^+ \tag{13}$$

where $(x)^+$ denotes $max(x,0)$ and $u$ represents the water level chosen to satisfy the total power constraint, i.e., $\sum_{i=1}^{2}\|\mathbf{c_i}\|^2 p_i \le P$. Applying our heuristic $\mathbf{f_2} = w\mathbf{h_2}$ to this equation, the power coefficient allocated to $m_2$ becomes:

$$p_2 = \left(w^2\frac{P\|det(\mathbf{H})\|^2+\sigma_1\|\mathbf{h_2}\|^2}{2\|\mathbf{h_1}\|^2} - \frac{\sigma_2}{2}\right)^+. \tag{14}$$

Based on this derivation, in MT allocation, selection of $w > 1$ is desirable to gain favorable power allocation, especially when the first term, $\frac{P\|det(\mathbf{H})\|^2+\sigma_1\|\mathbf{h_2}\|^2}{2\|\mathbf{h_1}\|^2}$, is close to the second term $\frac{\sigma_2}{2}$. With a large enough $w$, it is even possible to gain power from that allocated to message $m_2$, while the power originally allocated to $m_2$ is 0. The effectiveness of this CSI selection will be discussed in Section 5.

## 5. Evaluation

We now evaluate the effect of the proposed attacks based on forged CSI while focusing on the sniffed SNR and the capacity increase with forged CSI. We implemented the proposed attacks on a testbed built with WARP boards [20], each of which is equipped with two antennas as shown in Fig. 7(a). In the case of physical security, an additional WARP board is connected to the transmitter with CM-MMCX modules to transmit artificial noises. An 802.11-like MU-MIMO is implemented in this testbed where the long training sequences following the 802.11n/ac standard are used to estimate the CSI. 30-symbol payloads are transmitted in a 2.4GHz band with 64 subcarriers. Only 52 of these 64 subcarriers are used to transmit data and 4 of them are used as pilots to correct the central frequency offset (CFO). There are three single-antenna receivers, two of which are intended users, $rx_1$ and $rx_2$, while the remaining receiver, $rx_3$, is the sniffer not in the concurrent transmission with others. The receiver $rx_2$ is designed to be the malicious user who reports forged CSI in concurrent transmission. We placed the nodes in the marked locations in Fig. 7(b) and each result at a single location is an average of 5 experimental runs. During the experiment, the locations of the transmitter and receivers are interchangeable, so there are more than 300 transmissions for each forged CSI setting. In each transmission, a single user transmission without beamforming is used first to estimate the channel gain for each receiver. The trace including received SNR lower than 5dB is discarded because it is not in the operational range of WiFi. The received SNR distribution in our experiments ranges between 5dB and 28dB which is representative of current wireless systems.

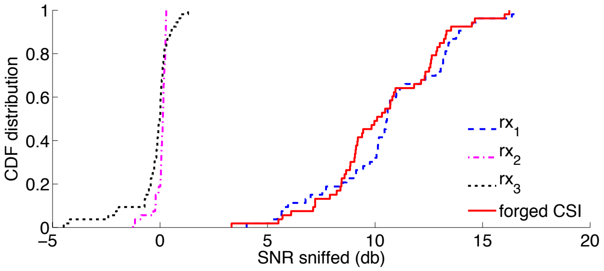

**Figure 8—SNR of sniffed message $m_1$.**

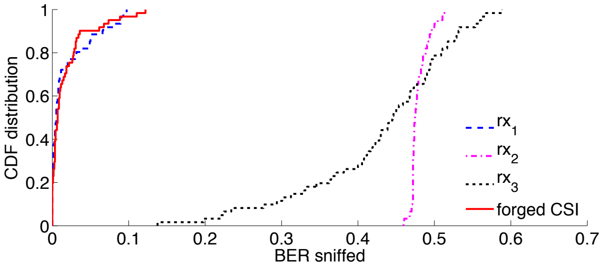

**Figure 9—BER of sniffed message $m_1$.**

### 5.1 Sniffing Attack

Figs. 8 and 9 show the decoded SNR and bit-error-rate (BER) of message $m_1$ received by different users $rx_1$, $rx_2$ and $rx_3$. Both $rx_2$ and $rx_3$ want to sniff $rx_1$'s message. In this experiment, 1/4 of the AP's power is used to send artificial noise, and signals are modulated with BPSK for easy comparison. We first investigate the performance of physical security with artificial noise and ZF-BF. Due to the artificial noise, the sniffer ($rx_3$) not in concurrent transmission experiences about 9dB SNR degradation in decoding $m_1$, compared to the decoded SNR at $rx_1$. On the other hand, the receiver$rx_2$ in the concurrent transmission receives almost nothing because the sent signal of $\mathbf{c_1}m_1$ is nullified in the direction of $\mathbf{h_2}$. This result is consistent with the proof in [3, 12]. One thing to note, though, is that $rx_3$ still has 4% probability to decode $m_1$ with less than 0.2 BER because the sent artificial noise is not strong enough to change the received signal in some traces where the direction of $\mathbf{h_3}$ is near the plane spanned by $\mathbf{h_1}$ and $\mathbf{h_2}$. This problem can be addressed further with higher modulation schemes like QPSK/16-QAM or by sending stronger artificial noise. The enhancement of secrecy by using stronger noise power is plotted in Fig. 11, where BER of sniffed data at $rx_3$ is increased to 0.45 when 65% of the AP's power is used to transmit artificial noise, but this also increases the BER of received data at $rx_1$ due to the decease of power used to transmit

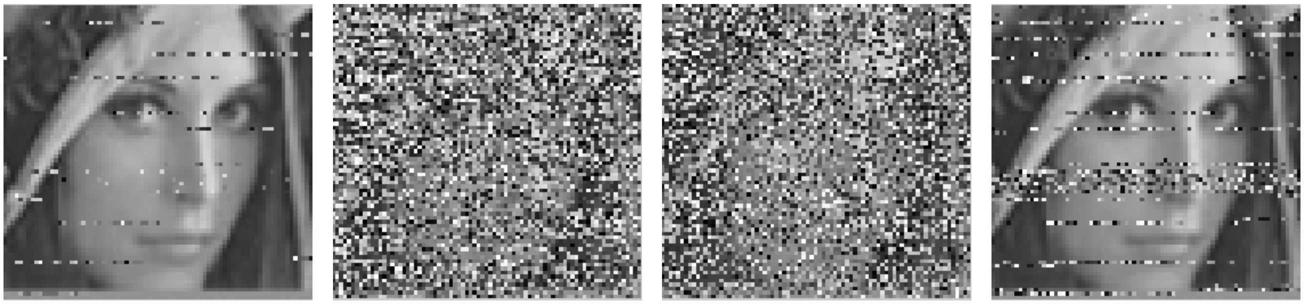

(a) Image at $rx_1$ (b) Image at $rx_2$ w/o cheating (c) Image at $rx_3$ (d) Image at $rx_2$ w/ cheating

**Figure 10—A bitmap image is transmitted to $rx_1$.** The information is leaked to $rx_2$ when CSI is forged even when physical security is enabled.

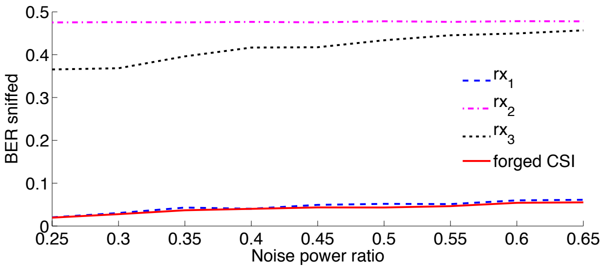

**Figure 11—BER of sniffed message $m_1$ at different noise levels.**

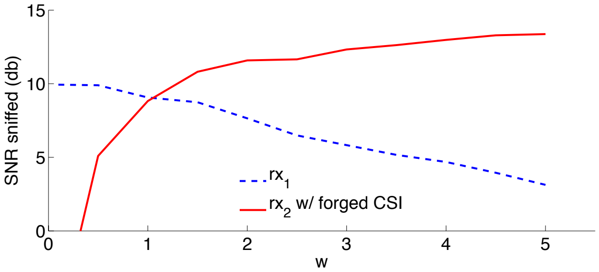

**Figure 12—SNR of sniffed message $m_1$ with different $w$.**

$m_1$. How to select a proper modulation or magnitude of artificial noise is beyond the scope of this paper since we will demonstrate an easier way for $rx_2$ to sniff message $m_1$ with the same decoding capability as the intended receiver $rx_1$ regardless of the strength of artificial noise.

As discussed previously, if the genuine CSI is reported, $rx_2$ will always receive nothing about message $m_1$, making the decoded BER of $rx_2$ close to the theoretical bound of 0.5. However, as shown in Figs. 8 and 9, when the CSI is forged to be $w\mathbf{h_1} - \mathbf{h_2}$ with parameter $w = 1$, the sniffed SNR and BER of $m_1$ at $rx_2$ are almost the same as decoded at $rx_1$. This implies that $rx_2$ can always sniff packets sent to $rx_1$ if it is decodable at $rx_1$ even when physical security is enabled. An illustrative example of sending and sniffing the transmission of a grayscale bitmap image via WARPs is provided in Fig. 10, demonstrating the effectiveness of physical security in thwarting eavesdroppers and how the forged CSI undermines the effectiveness of physical security.

These results are also consistent with the derivation of Eq. 9 in that $rx_2$ receives the same magnitude of $\sqrt{p}m_1$ as at $rx_1$ if $w$ is set to 1, but the decoded SNR at $rx_2$ is slightly lower than that at $rx_1$ because of imperfect cancellation of the interference caused by $\sqrt{p}m_2$. The larger the $w$, the larger the magnitude of sniffed message $m_1$. The average decoded SNR of $m_1$ with different $w$ is plotted in Fig. 12. One takeaway from this figure is that selecting a large $w$ helps decode the sniffed message, but this trend is not prominent after $w$ gets larger than 2 because reporting a larger $w$ incurs less total power to be used for transmitting data to $rx_2$ as shown in the power attack. Moreover, reporting a forged CSI as $w\mathbf{h_1} - \mathbf{h_2}$ with a large $w$ is unwise because the AP might not place $rx_2$ in the same concurrent transmission group with $rx_1$ if the reported CSI indicates a direction that degrades the performance of MU-MIMO. For example, reporting a forged CSI with $w = 5$ decreases the decoded SNR of $m_1$ at $rx_1$ from 11dB to 4dB, and hence it is unlikely for the AP to place those two clients in the same concurrent transmission group. These observations led to the selection of $w = 1.5$ for the purpose of eavesdropping, which ensures $rx_2$ to have 3dB higher capability in decoding $m_1$ than $rx_1$ with less effect on the received SNR at $rx_1$. However, the optimal selection of $w$ also depends on other conditions, such as the users scheduling algorithm adopted by the AP and the channel condition of eavesdroppers, which are beyond the scope of this paper. One thing worth noting in this sniffing attack is that malicious clients should be equipped with at least the same number of antennas as other concurrent receivers to sniff packets. It is straightforward to extend this attack to the system with multiple antennas.

## 5.2 Power Attack

To measure the performance of the proposed power attack, a metric called *capacity increase ratio* is used to estimate the resources unfairly gained with forged CSI, which is defined as the ratio of capacity increase with forged CSI to the capacity without forged CSI.

Based on the observation in Section 4, we let $rx_2$ report forged CSI as $w\mathbf{h_2}$ with $w = 0.3$ and $w = 4.0$ in EP and MT, respectively. The overall performance of these settings is summarized in Figs. 13 and 14. First, we find that in the EP power allocation, reporting forged CSI instead of genuine CSI can, on average, gain an unfair 20% additional capacity. This unfair power allocation also causes $rx_1$ a 40% decrease of capacity. The capacity increase with forged CSI actually depends on the channel conditions of concurrent clients, $rx_1$ and $rx_2$. To study this, we group the metric by different SNR gaps which are defined as the original SNR of $rx_2$ minus the original SNR of $rx_1$ if MU-MIMO is disabled. As shown in Fig. 15, if the SNR gap between $rx_1$ and $rx_2$ is between $-5 \sim -10$dB, i.e., the channel condition of $rx_1$ is far better than that of $rx_2$, pretending to have a “weaker" channel in the EP scheme

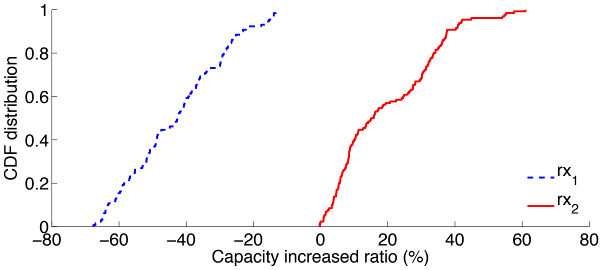

**Figure 13—Capacity increase distribution in EP.**

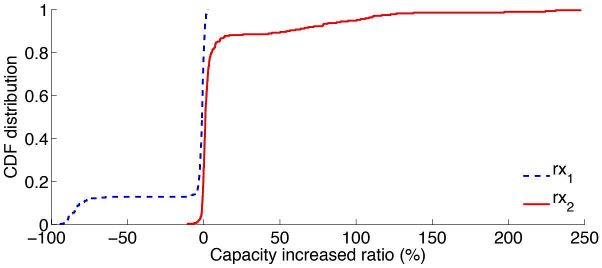

**Figure 14—Capacity increase distribution in MT.**

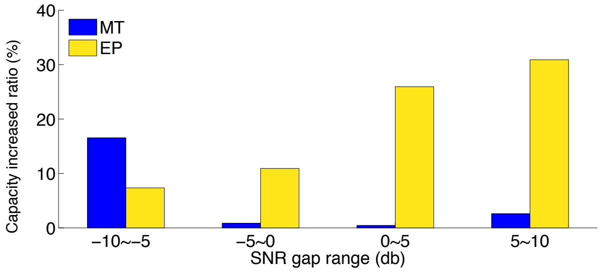

**Figure 15—Capacity increase at $rx_2$ with forged CSI.**

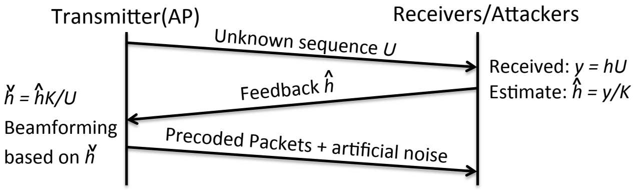

**Figure 16—Proposed CSIsec feedback scheme.**

doesn't help much because the AP is already spending more power for $rx_2$. In contrast, if the SNR gap ranges from 5 to 10dB, implying that the channel condition of $rx_2$ is much better and the AP intentionally reserves more power for $rx_1$. In this situation, pretending a weak channel with $w < 1$ helps $rx_2$ regain the shared power allocation, resulting in an unfairly high capacity increase ratio as shown in Fig. 15.

An interesting observation from our experiments is that the behavior of reporting forged CSI in MT is totally different from that in EP. As shown in Fig. 14, the overall performance of reporting forged CSI in MT is less effective for malicious clients than in EP. On average, only 13% additional capacity is gained with forged CSI in MT. Moreover, the distribution of capacity increase ratio is very skewed. 80% of transmissions increase capacity only by less than 6% by reporting forged CSI. This low effectiveness mainly comes from the forged parameter $w > 1$. When $w$ is set larger than 1, the precoding process of ZF-BF will inherently impose a weight $\frac{1}{w}$ on the message sent to $rx_2$ as shown in Eq. (12). Even in MT where the AP is misled to make a higher power allocation $p_2$ to $rx_2$ as shown in Eq. (14), the loss in the precoding matrix offsets the power unfairly acquired from the AP. We should also note that with 10% probability, $rx_2$ can gain more than 50% additional capacity. In such cases, $rx_2$ is actually located in the SNR region that the AP will remove power allocated to $rx_2$ due to its bad channel condition, compared to $rx_1$. In this situation, "balloning" CSI with $w > 1$ helps mislead the AP to share its power allocation with the malicious clients, boosting their received capacity. Based on this idea, we find that if the original SNR of $rx_2$ is 10 to 15dB less than $rx_1$, reporting forged CSI has potential for a large capacity boost as shown in Fig. 15, which is opposite to the EP power allocation.

Reporting forged CSI for the selfish purpose in MT is also less attractive than in EP. Note that the capacity increases in both schemes come at the expense of performance loss of $rx_1$. Moreover, in both schemes, reporting forged CSI with proper selection of $w$ incurs no penalty for receivers, and hence malicious clients have incentives to cheat on CSI in both schemes.

## 6. COUNTERMEASURES

We categorize the proposed forged CSI attacks (1) based on the genuine CSI and (2) not based on the genuine CSI. For example, the sniffing attack belongs to the first category while the power attack belongs to the second category.

One necessary assumption for the attacks in the first category is that malicious clients should know their own and other clients' CSI. This assumption is valid in existing MU-MIMO systems since CSI is estimated by receivers using a known sequence, $K$, in downlink transmission and fed back by receivers to the APs [38]. One obvious way to prevent CSI leakage during the feedback process is to introduce cryptosystems into the feedback process. For example, if the estimated CSI at a client is encrypted by the AP's public key before sending it to the AP, other concurrent clients cannot decode the CSI and forge CSI based on that information. However, this method requires the modification of CSI feedback protocol in both clients and the AP. Moreover, encrypting the feedback CSI incurs additional overheads of key exchange and encryption/decryption, which have been shown to be avoidable by using physical security. Thus, to prevent receivers from reporting forged CSI with limited overhead, we propose a novel CSI feedback system, called CSIsec protocol, as shown in Fig. 16.

In the CSIsec protocol, transmitters are assumed benign and transmit an unknown sequence $U$ – instead of the original known sequence $K$ – for clients to estimate CSI. $U$ is a random unknown sequence that varies whenever transmitters calls for the CSI estimation process (i.e., by changing the HT−LTF field), and only the transmitters know the secret value of $U$. This way, the estimated CSI at a receiver is $\hat{h} = (Uh + n)/K$, which is no longer an unbiased estimation of CSI, $h$, because the estimation process at receivers is misled by the unknown sequence $U$. After the transmitter receives the feedback from the receiver, i.e., $\hat{h}$, CSI is re-estimated by $\check{h} = K\hat{h}/U = h + n/U$, and this re-estimated CSI is used to calculate the precoding matrix. All the necessary modifications of CSIsec protocol is at the transmitter side, and there is no need to modify the receiver side of existing protocols.

### 6.1 Secure Analysis of `CSIsec` Protocol

To prove the security guarantee by `CSIsec`, the following questions are answered sequentially:

- Can the sniffing attack work under the protection of `CSIsec`?
- Can the unknown sequence be recovered by attackers?
- Can CSI be leaked via other side channels?

To mount the sniffing attack with forged CSI as introduced in Section 3, the attacker must report the forged CSI that is orthogonal to the same null space of the original CSI. Otherwise, the sniffed data is not decodable due to the sent artificial noise as shown in Section 5. For example, suppose $rx_2$ wants to sniff the packets sent to $rx_1$, then $rx_2$ should first sniff the CSI $\mathbf{h_1}$ of $rx_1$ when $rx_1$ feeds back the CSI to the AP. However, in the `CSIsec` protocol, the sniffed CSI $rx_1$ feeds back is actually $\hat{\mathbf{h_1}} = \frac{U_1}{K}\mathbf{h_1} + \mathbf{n}$ instead of $\mathbf{h_1}$, and even the estimated $\mathbf{h_2}$ at $rx_2$ is $\hat{\mathbf{h_2}} = \frac{U_2}{K}\mathbf{h_2}$ instead of $\mathbf{h_2}$. Recall that the way malicious clients can fool transmitters and sniff the messages to $rx_1$ is to make transmitters believe the CSI at $rx_2$ is $\check{\mathbf{h_2}} = \frac{K}{U_2}\hat{\mathbf{h_2}} = w\mathbf{h_1} - \mathbf{h_2}$ that has the same null space as the genuine CSI. This implies the feedback from $rx_2$ should be $\hat{\mathbf{h_2}} = \frac{U_2}{K}(w\mathbf{h_1} - \mathbf{h_2}) = \frac{wU_2}{U_1}\hat{\mathbf{h_1}} - \frac{U_2}{K}\mathbf{h_2}$, which can only be estimated when $U_2$ and $U_1$ are known to the attackers. Thus, `CSIsec` protocol prevents malicious clients from reporting forged CSI based on the genuine CSI of concurrent receivers. Our evaluation shows that failure to estimate CSI correctly incurs a 9dB drop of sniffed SNR which is considered secure in case of physical security [3].

The above analysis is valid under the assumption that $U_1$ and $U_2$ are secret known to transmitters only. This is a reasonable assumption because the received data at clients/attackers are $y = hU + n$, where $n$ is and additive Gaussian noise. When the CSI of clients is assumed independent of each other and drawn from a Gaussian distribution, it is easy to find that the distribution of received/sniffed data at clients, i.e., $y$, is also a Gaussian distribution, which is considered impossible to be decomposed if $h$, $U$ and $n$ are all unknown to clients. Thus, even when there are multiple attackers collude with each other, none of them can recover $U$ from the CSI estimation process, ensuring the security guaranteed by `CSIsec`. Suppose attackers can exhaustively search the space of all possible values of $U$ and know the answer once the right sequence is tried. In existing 802.11n/ac systems with 52 subcarriers, an unknown sequence, $U$, modulated with QPSK, i.e., $U = \{\pm 1, \pm i\}^{52}$, requires $2^{208}$ operations for malicious clients to find the CSI using a brute-force search, and this whole process must be completed before the timeout of feedback process.

Let us consider the possibility of existence of a side channel (outside the CSI estimation process) that malicious clients can use to estimate CSI. For example, CSI, denoted by $h$, can be estimated through other fields if the sent signal is known, i.e., $h = y/x : \ x \in$ other fields. As the structure defined in Fig. 17, 802.11n/ac packets can be divided into data and preamble fields. In MU-MIMO with physical security, the data field is precoded by ZF-BF and transmitted with artificial noise as discussed in Section 2, so it does not leak CSI to eavesdroppers. In contrast, preambles such as `L-STF`, `L-LTF`, `L-SIG`, and `HT-SIG` are transmitted without precoding for backward compatibility, thus enabling all receivers to sense these control fields. In the preamble fields, `HT-LTF1` to `HT-LTFN` represents the known sequence sent via each individual antenna for receivers to estimate their CSI. These fields do not leak CSI either, because `CSIsec` sends an unknown sequence $U$ in those fields instead of the original known sequence $K$. Thus, the only remaining fields that might leak CSI to eavesdroppers are `L-STF`, `L-LTF`, `L-SIG`, and `HT-SIG`. However, even if these fields are modulated without spatial precoding, malicious clients cannot estimate CSI through those non-beamformed fields because APs in 802.11n/ac use all antennas with cyclic shifts to send those fields. Thus, the estimated CSI through those fields at receiver $rx_i$ is $\gamma_1 h_{i2} + \gamma_2 h_{i2} + \ldots + \gamma_N h_{iN}$, not CSI for each antenna, where $\gamma_j$ represents the predefined cyclic shift. When CSI is assumed to be an independent Gaussian distribution, the received data in those fields also follow the Gaussian distribution, so it too is unlikely to decompose CSI components, $h_{i1}$, $h_{i2}$,..., $h_{iN}$ by those fields.

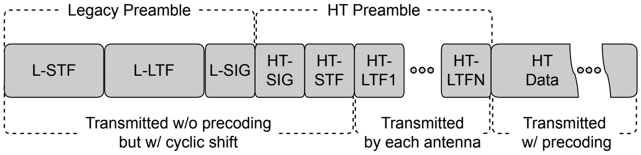

**Figure 17—Packet structure in 802.11n/ac.**

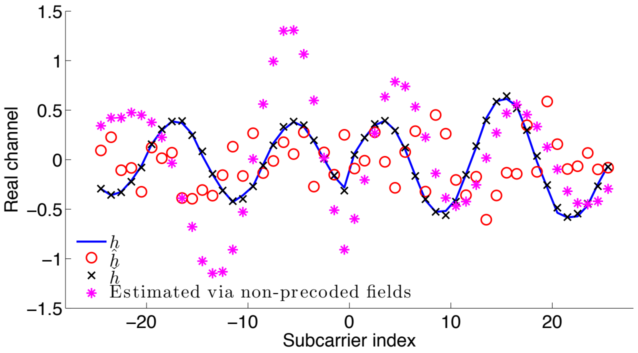

**Figure 18—Estimated CSI with different methods.**

An example of estimated CSI by different methods is shown in Fig. 18, where estimated CSI via `HT-LTF` at clients, $\hat{h}$, is divergent from the original CSI, $h$, due to the unknown sequence $U$. Moreover, the CSI estimated via `L-LTF`, i.e., the fields transmitted without spatial coding, is also different from the original CSI because of cyclic prefix. On the other hand, The AP is shown to be able to recover the original CSI with the misguided feedback from clients. In our experiments, the transmission precoded with CSI estimated by an unknown sequence $U$ incurs only an average of 1dB SNR loss due to its non-optimality of peak-to-average-power-ratio (PAPR).

We have shown above that malicious clients can only attack the precoding process based on the genuine CSI and knowledge of $U$. In reality, it is impossible to attack the precoding since CSI and $U$ are both unknown to the attackers. We also (i) showed the leakage of CSI under `CSIsec` is minimal and (ii) provided experimental evaluation results. In conclusion, `CSIsec` can prevent clients from getting an accurate estimate of CSI and also thwart any forged CSI attack based on the genuine estimation of CSI.

### 6.2 Disincentivized Power Allocation

In `CSIsec`, even though malicious clients are unable to know their own CSI, the second category of forging CSI attacks without knowing the CSI, such as the proposed power attack is still possible. For example, instead of reporting $\hat{h} = \frac{U}{K}h$ back to the AP, a malicious client can choose to report $w\hat{h}$ even though he doesn't know the real $h$. This feedback has the same effect as power attack

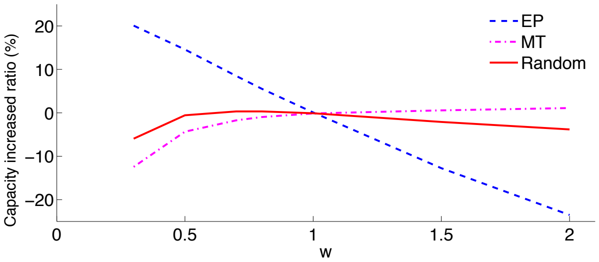

**Figure 19—Capacity increase in randomized settings.**

presented in Section 4 to mislead transmitters that the attacker's CSI is $\check{h} = \frac{K}{U}\hat{h} = wh$, instead of $h$.

It is easy to see that there is no way to prevent or even detect this behavior via feedback from clients when only the scale of forged CSI is modified. This problem is analogous to a malicious speedometer that always reports $w \times$ *speed* to drivers instead of the real *speed*. It is impossible for the drivers to determine if the speedometer is lying by querying the malicious speedometer (cheating clients). The only way to detect a malicious speedometer is to know the ground truth via other side channels, such as the estimated time driven from home to work. In wireless networks, there exist side channels of CSI at clients, e.g., spatial dependency in CSI among clients [17]. However, the existing work that explores the dependency among clients requires additional devices, help from (malicious) clients, and other strict assumptions on fading channels. It is also proven that CSI is quickly de-correlated in the real world when the distance between clients is greater than a half of the wavelength [4].

Thus, instead of finding other side channels to validate the correctness of reported CSI, future MU-MIMO should incorporate randomness into its protocols to discourage malicious clients from falsely reporting the scale of their CSI. For example, the AP can randomly switch between EP and MT power allocations to confuse malicious clients in choosing the right forged parameters $w$. In this way, the expected capacity to be gained by forging the scale of CSI is smaller than 0 because malicious clients are unable to determine which forged CSI to report. For example, as shown in Fig. 19, when the AP has 80% and 20% probabilities to allocate power based on MT and EP, respectively, malicious clients gain zero or negative reward regardless how they forge CSI.

## 7. Related Work

We first discuss the related work on CSI estimation in wireless networks and then present the marriage between MIMO networks and physical security. Finally, we discuss several state-of-art attacks on CSI feedback.

### 7.1 CSI in Wireless Communications

CSI is an important metric in wireless systems [1, 2, 5]. Imperfect estimation of CSI is known to cause serious problems in wireless systems [4, 47]. Therefore, many researchers proposed ways of improving the accuracy of CSI estimation and reducing the overhead of CSI transmission. In [21], the authors proposed a CSI estimation scheme based on the idea similar to compressive sensing and the authors of [43] investigated adaptive CSI feedback which ensures the accuracy and decreases the feedback overhead. In [19], the authors studied the tradeoff in channel correlation, user diversity and MU-MIMO efficiency, and then proposed an optimal CSI feedback scheme for downlink MU-MIMO systems by manipulating the time and frequency intervals in CSI feedback. Meanwhile, a large amount of effort has been made to overcome the problems caused by imperfect CSI [11, 28, 39]. However, none of these have considered the vulnerability of CSI feedback. We are the first to consider *intentional* forging and reporting of CSI to attack the precoding process for malicious purposes.

### 7.2 Physical Security

Physical security has been widely studied in wireless systems [27]. For example, the authors of [9] proved that, with an enough number of antennas, confidential messages protected with physical security are harder to break by brute-force approaches than traditional cryptosystems, and the authors of [22, 24] derived a theoretical bound of secret capacity of transmitting confidential messages in MU-MIMO. Physical security can be utilized in different forms in real-world systems. For example, proper artificial noise can be added without affecting intended receivers while keeping transmitted messages confidential [13]. In STROBE [3], this idea was experimentally validated and proved that the received SNR at sniffers is 15dB lower than the legitimate receivers. In MU-MIMO, ZF-BF was used to keep data confidential to concurrent receivers [12]. On the other hand, uplink authentication based on physical security was also proposed to differentiate attackers from legitimate transmitters through signatures like uplink CSI [30, 44]. A detailed survey of how physical security can be applied in multi-user wireless environments can be found in [27, 37]. However, all of these assume the reported CSI is genuine, but we prove that forged CSI can easily undermine the effectiveness of wireless network mechanisms. A recent effort [34] also shows the vulnerability of physical security based on artificial noise. There the precoding matrix with artificial noise is estimated by an adaptive filter under the assumption that the sent message is known to attackers. Multiple attackers and CSI of channels to them are necessary to realize their attack. The main difference between this and ours is that we don't need the known-plaintext assumption, making our proposed attacks more practical and general. Moreover, our attacks actively control the precoding matrix at transmitters with forged CSI, providing high efficiency in decoding sniffed data. For example, in our sniffing attack, a single malicious client can decode the same level of SNR as legitimate receivers with only one antenna. Note that the proposed `CSIsec` is also capable of mitigating the eavesdroppers shown in [34] because CSI estimation at eavesdroppers is necessary for prefiltering.

### 7.3 Attack Patterns Against CSI

Attacks against CSI feedback have been studied extensively. For example, the authors of [26] selectively jammed the CSI feedback to subvert MIMO network's performance, and a corresponding countermeasure to mitigate this problem was proposed in [29]. However, this paper is the first that discusses the vulnerability of the precoding process by reporting forged CSI. There also exists work [14, 23, 32] that uses falsified metrics to break systems, but none of them addresses the same questions as in this paper. For example, the scheme in [14] falsely replies a page message for DDoS attack in GSM, and [32] dishonestly increases the reported CQI to gains more transmission opportunity in LTE. Our work differs from them in that it exploits the precoding process in MU-MIMO. The closest to ours is the *mimicry* attack [23], in which attackers spoof uplink signals properly based on uplink CSI of legitimate transmitters against CSI-based authentication [30, 44]. In the *mimicry* attack, the uplink CSI is unknown to attackers, which will incur additional complexity to learn the uplink CSI by using methods like those in [17, 40], making this attack less practical. Instead of *mimicry* targeting the uplink, we focus on downlink security which is more vulnerable because downlink CSI must be fed back from receivers (attackers). Instead of estimating CSI explicitly, there exist techniques called *blind CSI estimation* [36, 42], which adaptively learns CSI based on the subspace structure of channel fading. A similar procedure is also used to create discriminatory channel estimation that eavesdroppers are unable to learn CSI [7, 18]. The concept of sending unknown sequences in CSIsec is also adopted in this work, but all of the existing protocols require the modification of clients while CSIsec only requires the modification of transmitters.

## 8. DISCUSSION

Since we target low-level design of wireless systems, the details of attacks need to be adapted for the different mechanisms under consideration, but it is still valid for other protocols. For example, if proportional fairness is adopted for power allocation, i.e., allocating power proportional to the clients' reported quality of links, the same technique to forge CSI with a higher magnitude can also be used by the attacker. On the other hand, if proportional fairness is adopted in the same way as in LTE for user scheduling [32], monotonically increased CSI can also gain slots unfairly allocated to that attacker. Actually, there are numerous protocols and variants of power allocation, link adaptation, and user scheduling in wireless systems. To illustrate and validate our proposed model with a real-world implementation, we present most representative mechanisms such as maximizing fairness or maximizing sum-capacity. The case of proportional fairness actually lies between these extremes, and different parameter settings can make it either more aggressive (maximizing sum-capacity) or conservative (maximizing fairness). Some proportional fairness schemes such as [25] also choose QoS as an indicator to allocate resource which can also be compromised because QoS is reported by clients, but how to manipulate QoS reports to gain an unfair share of resource is beyond the scope of this paper.

## 9. CONCLUSION

This paper studies the security of CSI feedback in MU-MIMO downlink and is the first to explore potential threats in the precoding process by forging CSI. We proposed two possible attacks, sniffing attack and power attack and validated their possibility in real-world MU-MIMO systems. Note, however, that these two are not the only possible attacks by forging CSI. We believe that any existing protocol relying on CSI feedback might be vulnerable to some extent. For example, the authors of [46] proposed an optimal user scheduling algorithm in ZF-BF that greedily adds clients with the CSI most orthogonal to other clients into concurrent transmission. In such a case, malicious clients can exploit the same trick used in *sniffing attack* to gain more for concurrent transmissions. Likewise, if the AP chooses a user scheduling algorithm which maximizes the total throughput, forging CSI as discussed in the *power attack* can be adopted to gain in concurrent transmissions.

To mitigate the problem caused by attacks on CSI feedback, we proposed CSIsec for existing MU-MIMO systems which only needs modification in the AP at minimal cost. Even though the attacks proposed in this paper have not yet been exploited in the wild, these threats are likely to become real as MU-MIMO like 802.11ac and physical security become prevalent. It is thus necessary for the research and industry communities to account for this vulnerability in the design of next-generation wireless communication systems.

## ACKNOWLEDGEMENTS

The authors would like to thank Alex Halderman, Safavi-Naini, Kassem Fawaz, Yuanchao Shu, Kate Ching-Ju Lin, Wei-Liang Shen, Yaohui Li, and anonymous reviewers for constructive comments on the earlier versions of this paper.

The work reported in this paper was supported in part by the NSF under Grants CNS-0905143 and CNS-1114837, by the AFOSR under Grant FA9550-10-1-0393, and by the ARO under Grant W811NF-12-1-0530.

## 10. REFERENCES

[1] Ieee standard: 802.11tm:wireless lan medium access control (mac) and physical layer (phy) specifications. 2007.

[2] Wireless lan medium access control (mac) and physical layer (phy) specification, ieee std. 802.11ac draft 2.0. 2012.

[3] N. Anand, S.-J. Lee, and E. Knightly. Strobe: Actively securing wireless communications using zero-forcing beamforming. In *Proceedings of IEEE INFOCOM '12*, pages 720–728.

[4] E. Aryafar, N. Anand, T. Salonidis, and E. W. Knightly. Design and experimental evaluation of multi-user beamforming in wireless lans. In *Proceedings of ACM MOBICOM '10*, pages 197–208.

[5] H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire. Achieving high data rates in a distributed mimo system. In *Proceedings of ACM MOBICOM '12*, pages 41–52.

[6] O. Bejarano, E. Knightly, and M. Park. Ieee 802.11ac: from channelization to multi-user mimo. *IEEE Communications Magazine*, 2013.

[7] T.-H. Chang, W.-C. Chiang, Y. Hong, and C.-Y. Chi. Training sequence design for discriminatory channel estimation in wireless mimo systems. *IEEE Transactions on Signal Processing*, 58(12):6223–6237, Dec 2010.

[8] M. H. M. Costa. Writing on dirty paper. *IEEE Transactions on Information Theory*, 29(3):439–441, May 1983.

[9] T. Dean and A. Goldsmith. Physical-layer cryptography through massive mimo. In *IEEE Information Theory Workshop (ITW) '13*, pages 1–5.

[10] W. Diffie and M. Hellman. New directions in cryptography. *Information Theory, IEEE Transactions on*, 22(6):644–654, 1976.

[11] M. Ding and S. Blostein. Mimo minimum total mse transceiver design with imperfect csi at both ends. *IEEE Transactions on Signal Processing*, 57(3):1141–1150, March 2009.

[12] G. Geraci, M. Egan, J. Yuan, A. Razi, and I. Collings. Secrecy sum-rates for multi-user mimo regularized channel inversion precoding. *IEEE Transactions on Communications*, 60(11):3472–3482, 2012.

[13] S. Goel and R. Negi. Guaranteeing secrecy using artificial noise. *IEEE Transactions on Wireless Communications*, 7(6):2180–2189, 2008.

[14] N. Golde, K. Redon, and J.-P. Seifert. Let me answer that for you: Exploiting broadcast information in cellular networks. In *Proceedings of USENIX SEC '13*, pages 33–48.

[15] S. Gollakota, S. D. Perli, and D. Katabi. Interference alignment and cancellation. In *Proceedings of the ACM SIGCOMM '09*, pages 159–170.

[16] D. Halperin, W. Hu, A. Sheth, and D. Wetherall. Predictable 802.11 packet delivery from wireless channel measurements. In *Proceedings of the ACM SIGCOMM'10*, pages 159–170.

[17] X. He, H. Dai, W. Shen, and P. Ning. Is link signature dependable for wireless security? In *Proceedings of INFOCOM '13*, pages 200–204.

[18] C.-W. Huang, T.-H. Chang, X. Zhou, and Y.-W. Hong. Two-way training for discriminatory channel estimation in wireless mimo systems. *IEEE Transactions on Signal Processing*, 61(10):2724–2738, May 2013.

[19] N. Jindal and S. Ramprashad. Optimizing csi feedback for mu-mimo: Tradeoffs in channel correlation, user diversity and mu-mimo efficiency. In *Proceedings of IEEE Vehicular Technology Conference (VTC Spring) '11*, pages 1–5.

[20] A. Khattab, J. Camp, C. Hunter, P. Murphy, A. Sabharwal, and E. W. Knightly. Warp: A flexible platform for clean-slate wireless medium

access protocol design. *SIGMOBILE Mob. Comput. Commun. Rev.*, 12:56–58, 2008.

[21] T.-H. Lin and H. Kung. Concurrent channel access and estimation for scalable multiuser mimo networking. In *Proceedings of IEEE INFOCOM '13*, pages 140–144.

[22] R. Liu, T. Liu, H. Poor, and S. Shamai. Multiple-input multiple-output gaussian broadcast channels with confidential messages. *IEEE Transactions on Information Theory*, 56:4215–4227, 2010.

[23] Y. Liu and P. Ning. Poster: Mimicry attacks against wireless link signature. In *Proceedings of the 18th ACM Conference on Computer and Communications Security*, pages 801–804.

[24] H. Ly, T. Liu, and Y. Liang. Multiple-input multiple-output gaussian broadcast channels with common and confidential messages. *IEEE Transactions on Information Theory*, 56:5477–5487, 2010.

[25] A. Mahmoud, A. Al-Rayyah, and T. Sheltami. Adaptive power allocation algorithm to support absolute proportional rates constraint for scalable ofdm systems. In *Proceedings of IEEE Vehicular Technology Conference*, pages 1–4, 2010.

[26] R. Miller and W. Trappe. Subverting mimo wireless systems by jamming the channel estimation procedure. In *Proceedings of ACM WiSec '10*, pages 19–24.

[27] A. Mukherjee, S. A. A. Fakoorian, J. Huang, and A. L. Swindlehurst. Principles of physical layer security in multiuser wireless networks: A survey. *CoRR*, 2010.

[28] A. Mukherjee and A. Swindlehurst. Robust beamforming for security in mimo wiretap channels with imperfect csi. *IEEE Transactions on Signal Processing*, 59(1):351–361, Jan 2011.

[29] G. Noubir, R. Rajaraman, B. Sheng, and B. Thapa. On the robustness of ieee 802.11 rate adaptation algorithms against smart jamming. In *Proceedings of ACM WiSec '11*, pages 97–108.

[30] N. Patwari and S. K. Kasera. Robust location distinction using temporal link signatures. In *Proceedings of ACM MobiCom '07*, pages 111–122.

[31] C. Peel, B. Hochwald, and A. Swindlehurst. A vector-perturbation technique for near-capacity multiantenna multiuser communication-part i: channel inversion and regularization. *IEEE Transactions on Communications*, 53(1):195–202, 2005.

[32] R. Racic, D. Ma, H. Chen, and X. Liu. Exploiting and defending opportunistic scheduling in cellular data networks. *Mobile Computing, IEEE Transactions on*, 9(5):609–620, 2010.

[33] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of ACM Symposium on Theory of Computing*, STOC '05, pages 84–93.

[34] M. Schulz, A. Loch, and M. Hollick. Practical known-plaintext attacks against physical layer security in wireless mimo systems. In *Proceedings of NDSS 2014*.

[35] W.-L. Shen, Y.-C. Tung, K.-C. Lee, K. C.-J. Lin, S. Gollakota, D. Katabi, and M.-S. Chen. Rate adaptation for 802.11 multiuser mimo networks. In *Proceedings of ACM MOBICOM '12*, pages 29–40.

[36] C. Shin, R. Heath, and E. Powers. Blind channel estimation for mimo-ofdm systems. *IEEE Transactions on Vehicular Technology*, 56(2):670–685, 2007.

[37] Y.-S. Shiu, S.-Y. Chang, H.-C. Wu, S.-H. Huang, and H.-H. Chen. Physical layer security in wireless networks: a tutorial. *IEEE Wireless Communications*, 18(2):66–74, 2011.

[38] D. Tse and P. Viswanath. *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.

[39] P. Ubaidulla and A. Chockalingam. Relay precoder optimization in mimo-relay networks with imperfect csi. *IEEE Transactions on Signal Processing*, 59(11):5473–5484, Nov 2011.

[40] T. Wang and Y. Yang. Analysis on perfect location spoofing attacks using beamforming. In *Proceedings of IEEE INFOCOM '13*, pages 2778–2786.

[41] H. Weingarten, Y. Steinberg, and S. Shamai. The capacity region of the gaussian multiple-input multiple-output broadcast channel. *IEEE Transactions on Information Theory*, 52:3936–3964, Sept 2006.

[42] L. Withers, R. Taylor, and D. Warme. Echo-mimo: A two-way channel training method for matched cooperative beamforming. *IEEE Transactions on Signal Processing*, 56(9):4419–4432, 2008.

[43] X. Xie, X. Zhang, and K. Sundaresan. Adaptive feedback compression for mimo networks. In *Proceedings of ACM MOBICOM '13*, pages 477–488.

[44] J. Xiong and K. Jamieson. Securearray: Improving wifi security with fine-grained physical-layer information. In *Proceedings of ACM MobiCom '13*, pages 441–452.

[45] A. Yarali and B. Ahsant. 802.11n: The new wave in wlan technology. In *Proceedings of Mobility '07*, pages 310–316.

[46] T. Yoo and A. Goldsmith. On the optimality of multiantenna broadcast scheduling using zero-forcing beamforming. *IEEE Journal on Selected Areas in Communications,*, 24(3):528–541, 2006.

[47] J. Zhang, M. Kountouris, J. Andrews, and R. Heath. Achievable throughput of multi-mode multiuser mimo with imperfect csi constraints. In *Proceedings of IEEE ISIT '09.*, pages 2659–2663.

# APPENDIX

Based on the finding in [33], the learning with error (LWE) problem is conjectured as a hard problem in that even quantum computation needs an exponential time to solve. Similar to the proof in [9], we choose to reduce our CSIsec protocol to the LWE problem under the Gaussian channel assumption. Given the system dimension $n$, some prime integer $p \leq poly(n)$, and an arbitrary number of equations with error, the traditional LWE problem is:

$$\begin{aligned} \langle s, a_1 \rangle &\approx_\chi b_1 (\bmod\, p) \\ \langle s, a_1 \rangle &\approx_\chi b_1 (\bmod\, p) \\ &\vdots \end{aligned} \tag{15}$$

where $s \in \mathbb{Z}_p^n$ is the secret to recover, $a_i$ is chosen independently from $\mathbb{Z}_p^n$, and $b_i \in \mathbb{Z}_p$ is the result of inner-product of $s$ and $a_i$ with additive error following a $\chi$ distribution. Based on Theorem 1.1 in [33], this LWE problem can be reduced to the known shortest independent vectors problem (SIVP) which is conjectured hard to solve in linear time. Thus, reducing our problem to the LWE problem also proves that CSIsec is hard to break.

Unlike the proof in [9], which tries to protect secure messages via CSI, our target is to secure the unknown sequence $U$ because it is the crucial information that guarantees CSIsec to work as shown in Section 6. For this purpose, we assume the unknown sequence is drawn uniformly from discrete periodic constellation such as M-PAM, and CSI is independent among clients which follow a Gaussian distribution with zero mean. This assumption is commonly used in wireless networks, especially in indoor environments where rich multipath communications exist. By mapping the unknown sequence to $s$ and CSI to $a_i$, and by assuming the additive noise follows the distribution of $\chi$, our MU-MIMO problem can be represented as an instance of the known LWE problem. Under this setting, if any attacker, or multiple attackers, have an oracle to recover the unknown sequence, $s$, from the received signal $b_i$, then this oracle can also be used to solve the known LWE problems as shown in Eq. (15). Details of this reduction are omitted because of space limit. This reduction shows that, under the assumption that breaking CSIsec is equivalent to finding $U$ (as argued in Section 6), breaking CSIsec is as hard as solving the LWE problem, which is conjectured hard to solve.