Fewer Demands, More Chances: Active Eavesdropping in MU-MIMO Systems

Xingya Zhao
The Ohio State University
Columbus, Ohio, USA
zhao.2053@osu.edu

Anwesha Roy
The Ohio State University
Columbus, Ohio, USA
roy.464@osu.edu

Avishek Banerjee*
The Ohio State University
Columbus, Ohio, USA
banerjee.152@osu.edu

Kannan Srinivasan
The Ohio State University
Columbus, Ohio, USA
kannan@cse.ohio-state.edu

ABSTRACT

As the demand for high-speed and reliable wireless networks continues to increase, multi-user multiple-input multiple-output (MU-MIMO) technology has become a popular choice for wireless communication systems. However, this technology also brings new security challenges, one of which is the vulnerability during the channel sounding process. In this paper, we propose an active eavesdropping attack targeting MU-MIMO systems. The attack consists of two phases. First, the attacker sends a forged pilot packet to the victims. After that, the access point transmits streams intended for victims to the attacker, who operates in full-duplex mode and relays the streams to the victims. Compared to existing eavesdropping attacks targeting MU-MIMO systems, our proposed attack requires less prior knowledge and coordination from attackers and maximizes eavesdropping opportunities. We evaluate the proposed attack in various settings and prove its effectiveness with multiple victims and partial channel knowledge. Additionally, we explore the use of physical-layer features to detect our proposed attack.

CCS CONCEPTS

• **Networks** → **Mobile and wireless security**.

KEYWORDS

physical layer security, wireless network

ACM Reference Format:
Xingya Zhao, Anwesha Roy, Avishek Banerjee, and Kannan Srinivasan. 2024. *Fewer Demands, More Chances*: Active Eavesdropping in MU-MIMO Systems. In *Proceedings of the 17th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '24), May 27–30, 2024, Seoul, Republic of Korea.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3643833.3656136

*Current affiliation is Nokia Bell Labs. This work was conducted during a previous affiliation.

WiSec '24, May 27–30, 2024, Seoul, Republic of Korea

ACM ISBN 979-8-4007-0582-3/24/05…$15.00
https://doi.org/10.1145/3643833.3656136

1 INTRODUCTION

Wireless communication has become an essential part of modern society with a growing demand for high-speed and reliable wireless networks. In response to this demand, multiple-input multiple-output (MIMO) technology has been widely adopted in wireless communication systems due to its ability to improve spectral efficiency and enhance the quality of service [1, 8, 10, 22]. Multi-user MIMO (MU-MIMO) further extends MIMO technology. It allows multiple users to communicate with a multi-antenna access point (AP) simultaneously at the same frequency by spatial multiplexing. MU-MIMO technology has been incorporated into the latest wireless communication standards, such as IEEE 802.11ac [9], IEEE 802.11ax [11] and 5G [26, 34]. The proliferation of wireless devices and the exponential growth of data traffic have also made MU-MIMO increasingly popular in both academic research and industrial applications in recent years [16, 27, 32, 37, 50].

While MU-MIMO technology offers significant benefits to wireless communication systems, it also introduces new security challenges. One of them arises from the channel sounding process [39]. To perform MU-MIMO, the AP needs to measure accurate channel state information (CSI) between the clients and itself, which is completed through the exchange of control packets. To ensure that clients at different locations can all participate in the channel sounding and later MU-MIMO communications, the AP broadcasts the control packets omnidirectionally. Additionally, to reduce the overhead of the channel sounding, the control packets are all transmitted in plaintext. The broadcasted plaintext packets make it possible for a potential attacker to passively eavesdrop on CSIs of clients or even launch active attacks.

Several studies in the literature have investigated vulnerabilities in the channel sounding process, leading to various eavesdropping attacks on MU-MIMO systems. Tung et al. [39] and Mao et al. [29] propose active eavesdropping attacks for MU-MIMO systems with explicit or implicit channel feedback. The malicious party executes the attacks by joining the network as a malicious client and sending forged CSI feedback or pilots to the AP to corrupt its channel measurements. The polluted channel measurements allow the attacker to receive signals containing the information intended for the victim client and itself. When signals intended for the attacker are known, the attacker can cancel them from the received signals and decode the messages meant for the victim from the remaining

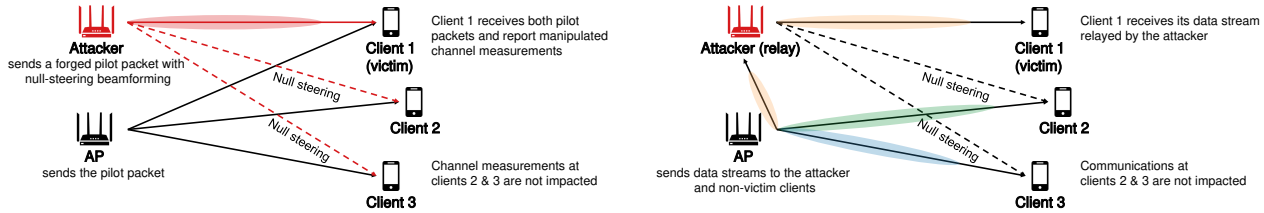

(a) Channel measurement manipulation

(b) Data stream relaying

Figure 1: Attack model. In the first phase, the attacker transmits a forged pilot packet to cancel out the AP-victim channel and inject the AP-attacker channel. In the second phase, the AP transmits the victim's data stream to the attacker, which then relays the received data stream to the victim.

signals. Wang et al. [47] extend this attack model to attack multiple victims with more attacker devices as malicious clients.

While active eavesdropping attacks have been proven effective in compromising MU-MIMO systems, they place specific requirements on the attacker devices: (1) **Participation in targeted transmissions as client(s)**: If there are more clients than the maximum data streams an AP can support in one transmission, only a subset of clients is selected for each transmission based on channel conditions, user fairness, and system capacity [37, 52]. This client selection process can reduce the opportunities for successful eavesdropping attacks [45], especially in multi-victim scenarios where the number of attackers participating in the transmissions must be equal to or greater than the number of victims [47]. (2) **Prior knowledge of packets for malicious client(s)**: Attacker devices need to know contents of the packets intended for them as input for signal cancellation, which typically requires cooperating servers to transmit predefined data. (3) **Shared eavesdropped signals among attacker devices**: In multi-victim scenarios, the multiple attacker devices must collaborate and share eavesdropped signals to decode messages intended for the victims.

In this work, we propose a novel eavesdropping attack in MU-MIMO systems where the attacker can use a multi-antenna full-duplex device to eavesdrop on one or multiple victims. Our proposed attack consists of two phases, as illustrated in Figure 1 with a one-victim example. In the first phase, the attacker sends a forged pilot packet with null-steering beamforming to the victims while the AP sends the legitimate pilot packet to the clients. The pilot in the forged packet is manipulated so that the channels measured from this packet will cancel the AP-victim channel and inject the AP-attacker channel. In the second phase, the AP transmits the streams intended for the victims to the selected antennas at the attacker, who operates in MIMO full-duplex mode and relays the received streams to the victims.

Compared to existing eavesdropping attacks targeting MU-MIMO systems, the proposed attack offers several significant advantages. First, *it demands less prior knowledge and coordination from the attacker*. To execute this attack, the attacker only requires control over a single multi-antenna full-duplex device. This malicious device does not need to join the network as a client together with the victims, and our attack does not rely on external servers to transmit any known data packets. What's more, *it maximizes eavesdropping opportunities* by operating independently of user selection results and can be performed whenever the targeted victims are selected.

We summarize our contributions as follows:

- We propose a novel active eavesdropping attack on MU-MIMO systems. Compared to existing eavesdropping attacks targeting MU-MIMO systems, our proposed attack requires less prior knowledge and coordination from attackers and maximizes eavesdropping opportunities.
- We prove the effectiveness of our proposed attack in various settings, including cases with multiple victims and partial channel knowledge. The secrecy capacity[1] of the victims can be brought down to zero.
- We evaluate the performance of using physical-layer features, such as angle of arrival (AoA) and carrier frequency offset (CFO), to detect the proposed attack.

2 BACKGROUND

2.1 MU-MIMO Systems

MU-MIMO is a space division multiplexing technology for wireless communication systems. By creating multiple independent spatial streams, it allows a multi-antenna AP to communicate with multiple users simultaneously in one frequency band and thus significantly improves the overall network efficiency. MU-MIMO has been introduced as a mandatory feature to Wi-Fi protocols since 802.11ac [9] and supported by numerous commercial devices [16].

To generate independent spatial streams, the AP needs to measure channels between the clients and itself during the channel sounding process. In MU-MIMO systems with explicit channel feedback, the AP first broadcasts a pilot packet. Upon receiving the pilot packet, each client measures the channels from the AP's antennas to itself based on the known pilot and sends the channel measurements back to the AP in the form of a feedback packet. Based on the received feedback packets, the AP calculates the appropriate weights to apply to each data stream to transmit at its antennas to reduce interference among clients. The matrix formed by these weights is called the *precoding matrix*.

Precoding methods can be classified as linear and non-linear. Although the achievable capacity of linear precoding methods is slightly lower than some more complicated non-linear methods such as dirty paper coding, the linear precoding methods are widely

[1] Secrecy capacity measures the maximum rate of the confidential information under the threat of eavesdroppers. It is calculated as $C_S = \max\{0, C - C_E\}$ where C denotes the legitimate channel capacity and C_E denotes the eavesdropper capacity. A secrecy capacity of zero means that the victims are completely compromised.

preferred for their lower computation overheads [19, 48]. A representative example of linear precoding methods is zero-forcing beamforming [5, 38, 52]. Consider a case of an M-antenna AP and N single-antenna clients ($N < M$) and let $\mathbf{H}$ denote the N-by-M channel matrix between the AP and clients, where the entry in the i-th column and j-th row represents the channel value from the AP's i-th antenna observed at the j-th client. With zero-forcing beamforming, the precoding matrix $\mathbf{C}$ is calculated as:

$$\mathbf{C} = \mathbf{H}^{+} = \mathbf{H}^{*}(\mathbf{H}\mathbf{H}^{*})^{-1} \tag{1}$$

where $\mathbf{H}^{*}$ represents the conjugate transpose of $\mathbf{H}$, $(\cdot)^{+}$ represents Moore-Penrose inverse, and $(\cdot)^{-1}$ represents inverse.

Let $\mathbf{x}$ denote the N-by-1 data vector to be transmitted to the N clients, and $\mathbf{P}$ denote the diagonal N-by-N power allocation matrix $\mathrm{diag}(p_1, \cdots, p_N)$, where p_j represents the power allocated to the j-th client during transmission. The precoded vector to be sent at M antennas is $\mathbf{C}\sqrt{\mathbf{P}}\mathbf{x}$ and the received signal at receivers will be:

$$\mathbf{y} = \mathbf{H}\mathbf{C}\sqrt{\mathbf{P}}\mathbf{x} + \mathbf{n} = \mathbf{H}\mathbf{H}^{*}(\mathbf{H}\mathbf{H}^{*})^{-1}\sqrt{\mathbf{P}}\mathbf{x} + \mathbf{n} = \sqrt{\mathbf{P}}\mathbf{x} + \mathbf{n} \tag{2}$$

where $\mathbf{n}$ denotes the noise vector observed at receivers. With precoding, the received signal at each receiver will have negligible interference from other clients, and each client can decode the signal independently without any knowledge about the other clients.

The power allocation matrix $\mathbf{P}$ needs to satisfy the constraint $\|\mathbf{C}\sqrt{\mathbf{P}}\mathbf{x}\|^2 \leq P$, where P is the total transmit power. The values of each entry in $\mathbf{P}$ can be decided by the specific power allocation strategy. The two most representative strategies are equal power allocation and maximal throughput power allocation. The equal power allocation maximizes fairness among concurrent receivers with $p_1 = \cdots = p_N$, and the maximal throughput power allocation maximizes the aggregated capacity of concurrent receivers with $\mathrm{argmax}_{p_j} \sum_{j=1}^{N} \log_2(1 + p_j/|n_j|^2)$, where n_j represents the noise observed at the j-th client and $|n_j|^2$ is the noise power.

2.2 Full-Duplex Implementations

Full-duplex devices are designed to transmit and receive signals simultaneously at the same frequency bands. The main challenge to implement full-duplex devices is canceling the strong self-interference caused by the transmitted signal at the receiving side. Various techniques have been proposed to tackle this challenge, such as combining antenna, analog and digital cancellations [18], using Balun transformers to improve the analog cancellation [25], and canceling non-linear distortions [15]. In [6, 14], the authors extend the full-duplex implementations to MIMO scenarios and propose methods to cancel the interference across the multiple antennas.

Full-duplex devices can be used as real-time relays when they retransmit the signals that are just received. Full-duplex relays have been implemented and used to improve wireless communication system performance [13, 17] and sensing system security [33].

3 RELATED WORK

3.1 Eavesdropping in Wireless Networks

While traffic through wireless communication systems is often protected by end-to-end encryption, the security built upon cryptographic protocols can be time-limited, as demonstrated by historical vulnerabilities in standards such as wired equivalent privacy (WEP) [21], Wi-Fi protected access (WPA) [42], WPA2 [43, 44], and WPA3 [41, 44]. The ever-present risk of current protocols being compromised necessitates a comprehensive understanding of security problems in lower network layers such as eavesdropping.

Eavesdropping attacks include passive and active eavesdropping attacks. Active eavesdropping attacks involve an attacker participating in wireless transmissions, such as transmitting jamming signals [53] or performing man-in-the-middle attacks [40] for more advantage. On the other hand, passive eavesdropping attacks involve an attacker intercepting wireless transmissions proactively. They are typically carried out with a wireless receiver that can capture the transmissions between legitimate transmitters and receivers, such as eavesdropping on access control tokens in RFID systems [23] or mmWave communications [12].

The widespread use of multi-antenna systems presents more challenges for eavesdropping. In multiple-input single-output (MISO) systems, beamforming enables the transmitter to send a directional beam to the receiver, limiting passive eavesdropping effectiveness unless the attacker is located at specific positions [51]. In MIMO systems, a passive eavesdropper receives mixed signals of all data streams. Additionally, if the data streams are precoded, precoding matrices will introduce more unknowns, making passive eavesdropping impractical with a reasonable number of antennas at the attacker device. In [36], an eavesdropping attack with known plaintext is proposed to counter the effect of precoding matrices.

The above-mentioned works concentrate on eavesdropping in the physical layer. Many existing attacks in wireless networks, such man-in-the-middle attacks [4, 20, 30, 54], focus on compromising protocols in higher layers and assume raw signal access. We believe that the contributions of physical layer eavesdropping attacks, including our proposed attack, are orthogonal to these attacks. And the eavesdropping attacks can help realize the signal access assumptions in the attacks focusing on higher network layers.

3.2 Attacks in MU-MIMO Systems

In MU-MIMO systems, each transmission contains multiple data streams and precoding is mandatory for users to decode their packets independently. This presents similar eavesdropping challenges encountered in MIMO systems. Additionally, due to the independence of data streams of multiple users, assuming prior knowledge of plaintext for all packets within the same transmission becomes more difficult, hindering the extension of known plaintext attacks as demonstrated in [36].

One related line of research to our proposed attack involves active eavesdropping attacks on MU-MIMO networks using malicious clients. In [39], the authors propose letting the attacker join MU-MIMO communications as a malicious client and send forged CSI feedback during channel sounding, enabling the reception of mixed information for both the victim and the attacker. By canceling the known signals intended for the attacker sent by a cooperating server, the attacker can eavesdrop on messages received by the other client. Similarly, in [28, 29], the authors propose a comparable attack targeting networks with implicit channel feedback by sending forged pilots to the base station to corrupt its channel measurements. Generalizations of this attack to multiple victim client scenarios are explored in [47], where multiple attacker devices

forge CSI feedback as a polynomial function of the CSI of victims and attackers. Considering that to perform this attack the number of attackers must be no less than the victims in one MU-MIMO communication, in [45], the authors study how to optimize the opportunity of having the malicious clients being selected with the victim clients in the same transmissions.

4 THREAT MODEL AND METHODOLOGY

We make the following assumptions about the attacker:

(i) The attacker controls a multi-antenna full-duplex device whose antenna count is greater than or equal to the number of targeted victim clients. This attacker device always has sufficient transmit power.
(ii) The attacker device is within the communication ranges of the AP, victim clients, and optionally non-victim clients.
(iii) The attacker has some basic knowledge of the communication system, such as packet format and pilot for channel sounding.
(iv) The attacker device can anonymously query the channels from the victim clients, and optionally the non-victim clients to itself.

Assumption (iii) is based on the fact that pilots used for channel measurements are usually defined in corresponding standards and are thus commonly known by devices [29]. Combined with assumption (ii), the attacker is able to measure channels from the AP, victim clients, and optionally non-victim clients to itself from regular transmissions in the system, such as beacons, channel sounding packets for MU-MIMO user selection updates, and previous data transmissions. If some parties have not participated regular transmissions for a long time, the attacker can leverage assumption (iv) to query the channels of interest. Assumption (iv) has been proved feasible in real-world Wi-Fi networks, where an AP will always respond clear-to-send frames to fake request-to-send frames [46], or acknowledgment frames to fake data frames [2] even if the client is unauthorized. To query the channels from clients, the attacker can send fake beacons and get the clients' responses [3].

With these assumptions, we propose an active eavesdropping attack on MU-MIMO systems with explicit channel feedback, outlined in two phases as shown in Figure 1. In the first phase, during the AP's channel measurement, the attacker simultaneously sends a forged pilot packet with null-steering beamforming to victims. The pilot in the forged packet manipulates channels measured from this packet to cancel the AP-victim channel and inject the AP-attacker channel. In the second phase, the AP precodes data streams with measured channels and transmits the stream intended for victims to selected attacker antennas. To ensure the communications for the victims are not interrupted, the attacker operates in the MIMO full-duplex mode and relays the received streams to the victims.

In the remainder of this section, we will first introduce the two phases of this attack with an example case of one victim client and the attacker having prior knowledge of channels from all clients to itself. We will then extend this attack to multi-victim scenarios and discuss strategies when the channels of non-victims are not accessible to the attacker.

4.1 Channel Measurement Manipulation

MU-MIMO systems rely on channel measurements for effective beamforming. In MU-MIMO systems using explicit channel feedback, the beamformer initiates channel measurement by transmitting a pilot packet omnidirectionally to all potential beamformees. Upon receipt of the pilot packet, each beamformee estimates the channel between itself and the beamformer and reports the results. Figure 2 demonstrates the channel feedback process in 802.11ac, where the pilot packet is referred to as the null data packet (NDP). To avoid feedback collapses, the AP sends beamforming report poll packets to notify one specific client to send its report each time. The short interframe space (SIFS) is the minimum separation time between high-priority frames, such as these control frames used for channel sounding.

To manipulate the victim's channel measurements, the attacker transmits a forged pilot packet at the same time as the AP, as illustrated in Figure 2. The simultaneous transmission can be achieved by letting the attacker prepare the forged pilot packet in advance and send it one SIFS after the NDP announcement transmission.

The forged pilot is designed to contain the information of a channel that can cancel the AP-victim channel and inject the AP-attacker channel. When the victim client receives both the original pilot packet and this forged packet, its measurement result will be the channel between the AP and the attacker, rather than the channel between the AP and itself. To formulate this process, we consider the case of an M-antenna AP (the transmitter), N clients (receivers), and a K-antenna attacker. Let $h_{t_i r_j}$ denote the channel from the AP's i-th antenna to the j-th client, $h_{t_i a_k}$ denote the channel from the AP's i-th antenna to the attacker's k-th antenna, and $x_{p,i}$ denote the original pilot value sent from the AP's i-th antenna. Assume that the first client is chosen as the victim, and the attacker wants to inject the channel of its first antenna $h_{t_i a_1}$ with a scaling factor α. Then the attacker needs to modify the forged pilot so that the victim can receive it as $(\alpha h_{t_i a_1} - h_{t_i r_1}) x_{p,i}$, where $h_{t_i r_1}$ can be heard from the victim's broadcasted beamforming report in the last round of MU-MIMO channel measurement (based on assumption (ii)), and $h_{t_i a_1}$ can be queried directly from the AP (based on assumption (iv)). Together with the original pilot packet $h_{t_i r_1} x_{p,i}$ received from the AP, the victim client will consider

$$y = (\alpha h_{t_i a_1} - h_{t_i r_1}) x_{p,i} + h_{t_i r_1} x_{p,i} + n = \alpha h_{t_i a_1} x_{p,i} + n \quad (3)$$

as the received pilot value, where n represents the noise. It will report a channel value close to $\alpha h_{t_i a_1}$ if the noise power is significantly smaller than the signal power.

While manipulating the channel measurements at the victim, the impact of forged pilots on non-victim clients should be minimized to avoid interference with their communications with the AP. To address this issue, the attacker utilizes zero-forcing beamforming on victim and non-victim clients when transmitting the forged pilot packet. In this transmission, we let the data intended for the victim be as derived above, and the data for non-victims be null, i.e.,

$$\mathbf{x_{A,i}} = [(\alpha h_{t_i a_1} - h_{t_i r_1}) x_{p,i} \quad 0 \quad \cdots \quad 0]^T \quad (4)$$

where $(\cdot)^T$ denotes matrix transpose. Let $\mathbf{P_A}$ represent the power allocation matrix used by the attacker, where the attacker sets $p_{A,1} = 1$ with its sufficient transmit power. According to Equation 2, if the channel stays stable and the noise has significantly lower

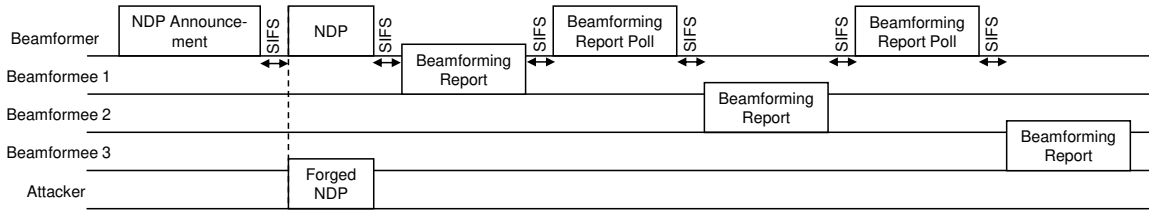

Figure 2: MU-MIMO channel measurement process in 802.11ac and the attacker channel injection

power than signals, the signal vector $\mathbf{y_{A,i}}$ received from the attacker is supposed to be very close to $\sqrt{\mathbf{P_A}}\mathbf{x}$ in zero-forcing beamforming transmissions. Thus the victim will receive the forged pilot from the attacker while all other clients will receive zero, i.e.,

$$\mathbf{y_{A,i}} \approx [(\alpha h_{t_i a_1} - h_{t_i r_1})x_{p,i} \quad 0 \quad \cdots \quad 0]^T \tag{5}$$

Meanwhile, the AP is also broadcasting the original pilot packet to all users. For the pilot value $x_{p,i}$, it arrives at clients as its original value multiplied by corresponding channels between the i-th antenna of the AP to clients. When noise is significantly weaker than signals, the signal vector $\mathbf{y_{T,i}}$ received from the AP will be

$$\mathbf{y_{T,i}} \approx [h_{t_i r_1} x_{p,i} \quad h_{t_i r_2} x_{p,i} \quad \cdots \quad h_{t_i r_N} x_{p,i}]^T \tag{6}$$

and the sum signal vector will be

$$\mathbf{y_i} = \mathbf{y_{A,i}} + \mathbf{y_{T,i}} \approx [\alpha h_{t_i a_1} x_{p,i} \quad h_{t_i r_2} x_{p,i} \quad \cdots \quad h_{t_i r_N} x_{p,i}]^T \tag{7}$$

This approach ensures that the victim client measures its channel as $\alpha h_{t_i a_1}$ while non-victim clients are less impacted. To control the power of the injected channels, we introduce the scaling factor α. The power of $h_{t_i a_1}$ can differ significantly from $h_{t_i r_1}$ due to variations in transmit power and locations between the AP and the attacker, which could affect the power allocation or even user selection results in MU-MIMO networks. The impact of this scaling factor on the attack efficiency will be evaluated in Section 5.2.

4.2 Data Stream Relaying

After successfully injecting pilot signals, the AP will regard the AP-attacker channel as the channel to the victim, and transmit the victim's data stream to the attacker. To avoid interrupting the communication between the AP and the victim, we let the attacker device function as a multi-antenna full-duplex relay during data transmissions. Similar to the pilot injection phase, the attacker performs null-steering zero-forcing beamforming while relaying the signal. Initially focusing on a single frequency band, let $x_{d,j}$ denote the data intended for the j-th client. Assuming the first client is the victim, according to Equation 2, the attacker receives

$$y_{d,1} \approx \sqrt{p_1} x_{d,1} \tag{8}$$

during the data transmission when the noise power is neglectable. To relay the signal with null-steering zero-forcing beamforming, the attacker prepares the data vector to relay as

$$\mathbf{r_d} = [\beta y_{d,1} \quad 0 \quad \cdots \quad 0]^T \tag{9}$$

where data stream for the victim is a scaled version of what the attacker receives about the victim's data, and the data streams for non-victims are null. We use β to denote the scaling factor used in data stream relaying. In this way, the victim client can get the information intended for it from the attacker, while other non-victim clients are less impacted by the relayed signals.

Many communication protocols use frequency division multiplexing (FDM) methods that involve multiple subcarriers, such as the orthogonal frequency division multiplexing (OFDM) used in LTE [1] and Wi-Fi standards since 802.11a [7]. A common practice to perform zero-forcing beamforming with multiple subcarriers is to first multiply the modulated symbols with the precoding matrix at each subcarrier in the frequency domain as in Section 2.1. The transmitter then converts the precoded symbols of all subcarriers to the time domain, and adds the cyclic prefixes (CP) to complete the OFDM symbol. While we adopt a similar method to obtain the precoded pilots during the pilot injection phase, this becomes infeasible during data transmission. This is because the attacker has sufficient time to prepare a forged pilot packet before the channel measurement process with the known pilot and channels during pilot injection, but the data packets during data transmission are unpredictable. Decoding the packet and performing frequency domain precoding is neither feasible given the intolerable delay.

To facilitate the real-time beamforming, we transform the precoding matrices in the frequency domain into precoding filters in the time domain. In [13], the authors implement a MIMO full-duplex relay with a construct-and-forward filter to make relayed signals constructively combine with the direct signals from the source. Our precoding filters can be implemented in the same way without introducing additional delay time.

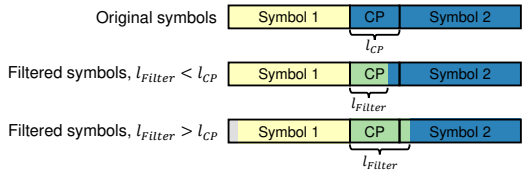

Figure 3: Impact of filter lengths on symbols. Green sections represent samples carrying information from both symbols.

After being converted to the time domain, the initial length of precoding filters will match the number of subcarriers. However, with a large number of subcarriers, the precoding filter length may exceed the maximum possible length permitted by the relay implementation. Moreover, if the precoding filter length is greater than the cyclic prefix length, applying the filters will increase the inter-symbol interference, as shown in Figure 3. This increased interference can adversely affect data transmission.

To constrain the length of precoding filters, we leverage the empirical observation that *the power of precoding filters typically concentrates on a small number of consecutive samples.* Figure 4

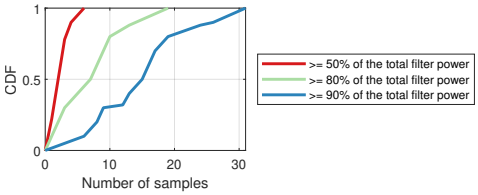

Figure 4: Distribution of the sample numbers taken to reach certain percentages of the total power of precoding filters

shows the cumulative distribution function (CDF) of the minimum number of consecutive samples required to reach specific power levels. We consider 50 traces of an MU-MIMO network serving three clients. We observe that choosing just 6 consecutive samples from the filters covers over 50% of total filter power, while selecting up to 19 consecutive samples covers over 80% of total filter power.

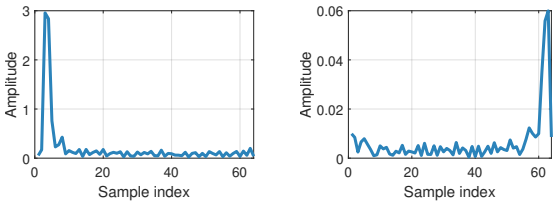

(a) Example channel **(b) Precoding filter**

Figure 5: An example channel in the time domain and its corresponding precoding filter at full length. The precoding filter looks similar to a time-reversed version of the channel, with its power concentrated in the last few samples.

We believe that this observation will hold in most cases. For each subcarrier, while calculating the precoding matrix $\mathbf{C}$ according to Equation 1, we have observed that for every row of the matrix $\mathbf{HH}^*$, the magnitude of the diagonal entry is usually larger than or equal to the sum of the magnitudes of non-diagonal entries in this row, i.e., $\mathbf{HH}^*$ is usually a diagonally dominant matrix. This is because in $\mathbf{HH}^*$ the diagonal entries represent the channel powers of clients, and the non-diagonal entries represent the interference of channels between client pairs. In the inverse of a strict diagonally dominant matrix, such as $(\mathbf{HH}^*)^{-1}$ in many cases, the largest entry in each column is on the diagonal [31]. Thus in $\mathbf{C} = \mathbf{H}^*(\mathbf{HH}^*)^{-1}$, the values from a scaled version of $\mathbf{H}^*$ can take a large part. When considering the precoding values of multiple subcarriers, the dominating conjugate of channel values in the frequency domain (values from $\mathbf{C}$'s of these subcarriers) will lead to a conjugate reverse of channel values in the time domain. Since in time-domain channel values, the power will usually concentrate on the first few samples, we can conclude that in the precoding filters, the filter power will usually concentrate on the last few samples, as shown in Figure 5.

By selecting these consecutive samples with dominant power, we can obtain shorter filters without significantly compromising beamforming performance.

4.3 Scaling to Multiple Victims

The proposed attack can extend to multi-victim scenarios by utilizing multiple antennas of the attacker device. In the one-victim scenario, the attacker performs null-steering zero-forcing beamforming and has one non-null data stream for the forged pilot or relayed data packets. In the multi-victim scenario, the attacker can similarly create multiple non-null data streams, one for each victim. Theoretically, a K-antenna attacker node can create up to K data streams, enabling it to attack up to K clients in one transmission. If the attacker has equal or more antennas than the AP ($K \geq M$), it can attack all clients served in one transmission. In this section, we assume that the attacker aims to attack V out of N clients. For ease of explanation, we assume that the first V clients are victims, although the attack can be applied to any subset of size V.

To initiate the attack, the attacker selects V antennas to receive data streams intended for the victims. As channels from the AP to selected antennas will later be injected into the AP's channel measurement, these antennas should have the least correlated channels from the AP to minimize interference in the AP's channel measurement. For simplicity, we assume the first V antennas are chosen, with each victim corresponding to a respective antenna.

In the pilot injection phase, similar to Equation 4, the attacker prepares a data vector

$$\mathbf{x_{A,i}} = [(\alpha h_{t_i a_1} - h_{t_i r_1})x_{p,i} \quad \cdots \quad (\alpha h_{t_i a_V} - h_{t_i r_V})x_{p,i} \quad 0 \quad \cdots \quad 0]^T \tag{10}$$

where the first non-zero V values are the forged pilot value for the victims, and the following $N - V$ zeros are the null data streams for non-victim clients. With zero-forcing beamforming, the forged pilot will be received by clients as

$$\mathbf{y_{A,i}} \approx [(\alpha h_{t_i a_1} - h_{t_i r_1})x_{p,i} \quad \cdots \quad (\alpha h_{t_i a_V} - h_{t_i r_V})x_{p,i} \quad 0 \quad \cdots \quad 0]^T \tag{11}$$

together with the the signal $\mathbf{y_{T,i}}$ received from the AP

$$\mathbf{y_{T,i}} \approx [h_{t_i r_1} x_{p,i} \quad \cdots \quad h_{t_i r_V} x_{p,i} \quad h_{t_i r_{V+1}} x_{p,i} \quad \cdots \quad h_{t_i r_N} x_{p,i}]^T \tag{12}$$

the sum signal vector $\mathbf{y_i} = \mathbf{y_{A,i}} + \mathbf{y_{T,i}}$ will be

$$\mathbf{y_i} \approx [\alpha h_{t_i a_1} x_{p,i} \quad \cdots \quad \alpha h_{t_i a_V} x_{p,i} \quad h_{t_i r_{V+1}} x_{p,i} \quad \cdots \quad h_{t_i r_N} x_{p,i}]^T \tag{13}$$

In this way, the v-th victim client will measure its channel as $\alpha h_{t_i a_v}$, while measurements at non-victim clients are not impacted.

In the data relaying phase, the attacker behaves similarly as in Section 4.2, except that there will be V antennas receiving signals for the eavesdropping purpose, and now there are V data streams to relay to the clients with zero-forcing beamforming. Let $x_{d,j}$ denote the data intended for the j-th client, and $y_{d,j} \approx \sqrt{p_j} x_{d,j}$ denote the signal received by the attacker about the j-th victim's data. The attacker prepares the data vector to relay as

$$\mathbf{r_d} = [\beta y_{d,1} \quad \cdots \quad \beta y_{d,V} \quad 0 \quad \cdots \quad 0]^T \tag{14}$$

The precoding filters can be shortened in the same way as mentioned in Section 4.2.

4.4 Strategy with Partial Channel Knowledge

Assumptions (ii) and (iv) in Section 4 take into account scenarios where the attacker lacks channel information on some or all non-victim clients. This can occur if the attacker is too distant from certain clients to detect their signals. Another case is that the attacker has fewer antennas than the number clients in a transmission ($K < N$). In the second case, the attacker cannot generate data

streams for all clients and has to rely on partial channel feedback due to this constraint on data stream generation.

We suggest that the attacker can prioritize non-victim clients with stronger received signal strengths (RSS) and disregard those with weaker RSS. If the attacker's antenna count equals or exceeds the total number of victims and known non-victims, it can proceed with the attack as usual. Otherwise, it can ignore channels with the lowest RSS values. For non-victim clients not known or ignored by the attacker, their communication with the AP may be affected by the attacker's signals. However, due to their weaker RSS, this interference will have a lesser impact compared to other non-victim clients. Thus, neglecting them optimizes overall performance when the attacker has limited antennas.

5 EVALUATION

5.1 Data Collection

Certain key information for the attack evaluation, such as the raw signal and SINRs, is not accessible in commercial devices. To overcome this limitation, we use WARP v3 software-defined radios to collect channels in a typical indoor office environment. The full-duplex device parameters are set as in [13]. We emulate the full-duplex relay scheme by first letting the AP transmit and the relay receive. After that, the AP remains silent, and the relay retransmits its received signal. Both received signals are later combined to form a single received signal during the attack.

We generate packets following the 802.11ac physical layer standard and use band 11 with a 20 MHz bandwidth at 2.4 GHz for the experiments. Each channel measurement contains values of 64 subcarriers, 52 of which are data subcarriers. The AP and the attack are both equipped with four VERT2450 antennas, and the AP serves three single-antenna clients unless otherwise specified.

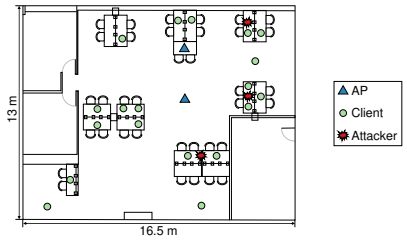

Figure 6: Layout of the office space and device locations

We consider a total of 30 settings in the typical office environment, and collect 5 channels with each setting. Each setting has a unique combination of AP/clients/attacker locations. The AP-client distances vary between 1 m and 10 m, the AP-attacker distances vary between 0.5 m and 9 m. The data collection spans two months and includes line-of-sight (LoS) and non-line-of-sight (NLoS) settings. In NLoS settings, we introduce common office obstacles such as cubicle panels, chairs, and books. Figure 6 illustrates the office layout and example device locations.

5.2 Impact of Key Parameters on Eavesdropping Efficiency

5.2.1 Precoding filter lengths. To evaluate the impact of the precoding filter lengths on the attacker's data relaying performance, we select 50 traces collected at 5 locations with 3 clients and 1 victim. By varying the precoding filter length, we evaluate its effect using two metrics: signal-to-noise ratio (SINR) at the victim of the received relayed data, and the leakage at the non-victims resulting from transmissions. We define the leakage as the sum of received signal power at the non-victim clients from the attacker. Lower leakage indicates that the attacker causes less interference to the non-victims' communications with the AP.

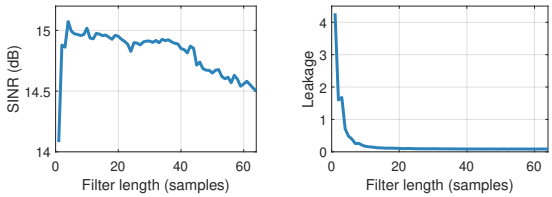

(a) SINR at the victim **(b) Leakage at non-victims**

Figure 7: SINR at the victim and leakage at the non-victims with varying precoding filter length

From Figure 7 we can see that initially the SINR at victims increases with the filter length. This is because very short filters cannot fully perform the beamforming. However, beyond a filter length of 20, SINR decreases due to increased inter-symbol interference, as discussed in Section 4.2. The average leakage at non-victims decreases with filter length, since the null data streams for non-victims do not suffer from inter-symbol interference and benefit from better beamforming performance. We use precoding filters of length 16 for balanced performance in the following experiments.

5.2.2 Scaling factors for the channel injection and data stream relaying. The scaling factors in the pilot injection and data relaying phases, α and β, play a significant role in the eavesdropping efficiency and MU-MIMO communication performance. We notice that the varying RSS of AP-attacker channels has a great impact on the eavesdropping efficiency. To eliminate its impact, we redefine the scaling factors with the RSS at non-victims as references, i.e.,

$$\alpha' = \frac{\sum_{k\in\mathbf{V}} \alpha \|\mathbf{H_{ta,k}}\|/|\mathbf{V}|}{\sum_{j\notin\mathbf{V}} \|\mathbf{H_{tr,j}}\|/(N-|\mathbf{V}|)}, \quad \beta' = \frac{\sum_{k\in\mathbf{V}} \beta \|\mathbf{H_{ta,k}}\|/|\mathbf{V}|}{\sum_{j\notin\mathbf{V}} \|\mathbf{H_{tr,j}}\|/(N-|\mathbf{V}|)} \tag{15}$$

where $\mathbf{V}$ denotes the victim set, $\mathbf{H_{ta,k}}$ denotes the channel matrix from the AP to the attacker's k-th antenna, $\mathbf{H_{tr,j}}$ denotes the channel matrix from the AP to the j-th client, and N is the number of clients. Let C denote the subcarrier count and A_t denote the antenna count at the AP. Both $\mathbf{H_{ta,k}}$ and $\mathbf{H_{tr,j}}$ are of size A_t-by-C.

To select appropriate scaling factors α' and β', we aim to fulfill the following requirements: (1) The estimated signal-to-interference-plus-noise ratios (SINR) of all clients should closely match their actual SINRs during data transmissions; (2) The SINR of the attacker should be close to or higher than that of the victims; (3) The victims should achieve as high SINRs as possible. The first requirement is to accommodate rate adaptation, where the AP selects the optimal transmission rate based on channel conditions. Setting the rate too high can cause packet loss and retransmissions, while setting it too low will reduce the network throughput. The second requirement aims to maximize eavesdropping efficiency, while the third minimizes the attacker's impact on victims' communications.

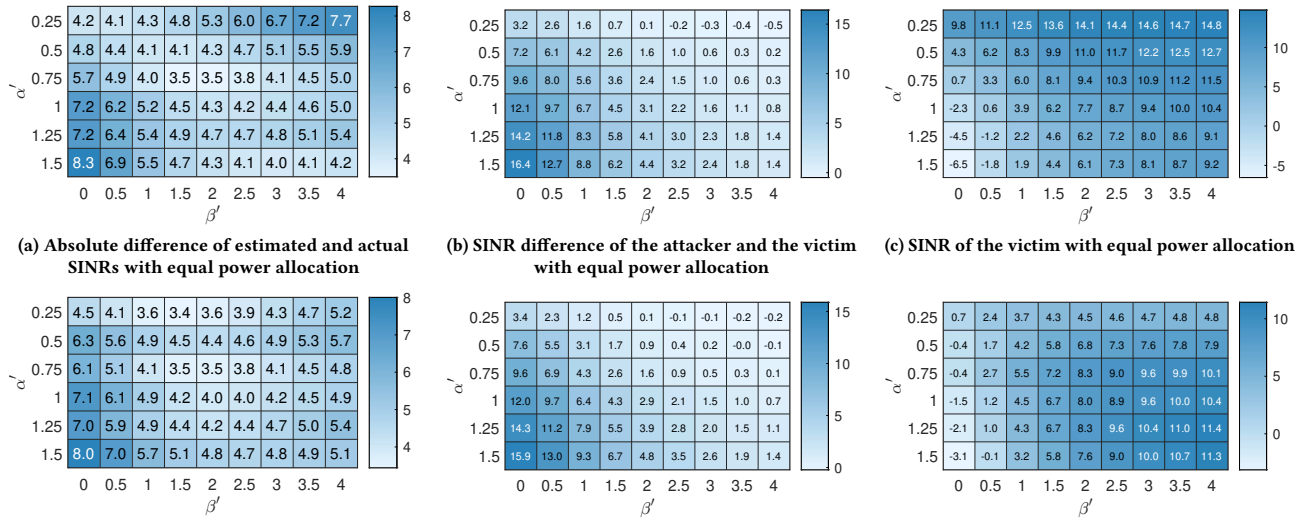

(a) Absolute difference of estimated and actual SINRs with equal power allocation

(b) SINR difference of the attacker and the victim with equal power allocation

(c) SINR of the victim with equal power allocation

(d) Absolute difference of estimated and actual SINRs with maximal throughput power allocation

(e) SINR difference of the attacker and the victim with maximal throughput power allocation

(f) SINR of the victim with maximal throughput power allocation

Figure 8: Metrics with varying scaling factors α' and β'

To evaluate the effectiveness of different scaling factors in meeting three critical requirements, we consider an experimental scenario where an AP serves three clients, of which one is selected as the victim. We consider 10 settings in the office environment and collect 5 traces in each setting. The metrics corresponding to the three requirements are shown in Figure 8. The difference between estimated and actual SINRs typically ranges from 2.5 to 3.5 dB in regular MU-MIMO transmissions without an attacker. Figures 8a and 8d indicate that selecting $\alpha' = 0.75$ and $\beta' = 1.5$ or 2 maintains this range for both power allocation strategies. Figures 8b and 8e show that in most cases the attacker's SINR at the attacker exceeds that at the victim, reducing secrecy capacity to 0. This is because the victim receives both the attacker's relayed signal and interference from the AP's beamforming. Despite this degradation brought by the nature of our attack model, with proper scaling factor selection considering the first metric, the AP will take the victims as clients with inherently weaker channels and adjust the transmission rates to accommodate them. Finally, Figures 8c and 8f show that with a fixed α' value, the victim's SINR increases with β', which indicates greater amplification is applied during data relay.

In the following evaluations, we will use $\alpha' = 0.75$ and $\beta' = 2$. The corresponding α and β values are calculated with Equation 15.

5.3 Overall Eavesdropping Efficiency

To investigate the overall eavesdropping efficiency, we collect 150 traces with 30 settings with varying AP/clients/attacker locations in the office environment. We consider a case of one AP serving three clients, and the AP adopts the equal power or maximal throughput power allocation strategies. To establish a baseline, while collecting each trace, we disable the attack once and monitor the signals received by the attacker. This baseline represents a receiver colocated with the attacker when the proposed attack is not executed. We refer to the baseline as a *passive eavesdropper*. The passive eavesdropper targets the same victim as the attacker in each transmission.

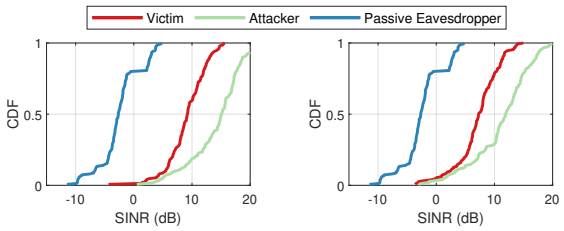

(a) Equal power **(b) Maximal throughput**

Figure 9: Distributions of average SINRs at the victims, the attacker, and a passive eavesdropper

From Figure 9 we can see that in almost all cases, the attacker gets higher SINRs than the victim and reduces the victim's secrecy capacity to zero. Compared to the passive eavesdropper, an attacker performing our proposed active eavesdropping attack has an SINR gain of around 18 dB with equal power allocation, and around 14 dB with maximal throughput power allocation.

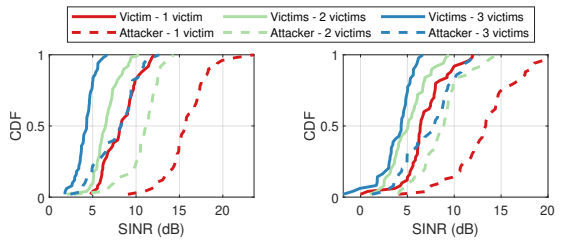

(a) Equal power **(b) Maximal throughput**

Figure 10: Distributions of average SINRs at the victims and the attacker with varying victim counts

5.4 Eavesdropping Efficiency with Multiple Victims

To evaluate how the eavesdropping efficiency varies with victim counts, we collect 50 traces with 10 settings in the office environment with varying AP/clients/attacker locations. We consider a case of one AP serving three clients, and the AP adopts the equal power or maximal throughput power allocation.

From Figure 10 we can see that with both power allocation strategies, SINRs decrease for victims and the attacker as the victim count increases due to increased channel correlation among attacker antennas compared to clients. In our test settings, the average correlation among channels from the AP to attacker antennas is 0.623, while to different clients it is 0.496. As more clients become victims, the AP utilizes more channels from attacker antennas, leading to increased channel correlation and interference. The signals are then relayed to victims and cause SINR drops.

5.5 Eavesdropping Efficiency with Partial Channel Knowledge

To evaluate the eavesdropping efficiency with partial channel knowledge of non-victim clients, we collect 50 traces with 10 settings in the office environment with varying AP/clients/attacker locations. We assume one of the three clients is selected as the victim, and the attacker is aware of channels of 0-2 non-victim clients. For the cases of 1 known non-victim, we assume the non-victim with higher RSS at the attacker is known and the other one is unknown.

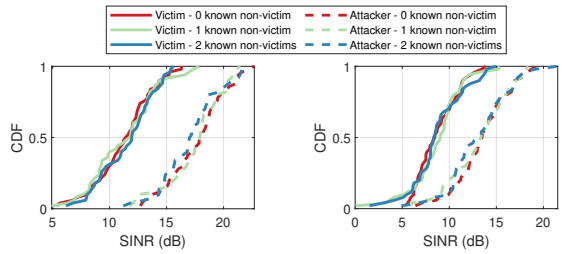

(a) Equal power **(b) Maximal throughput**

Figure 11: Distributions of SINRs at the victim client and the attacker with partial channel feedbacks

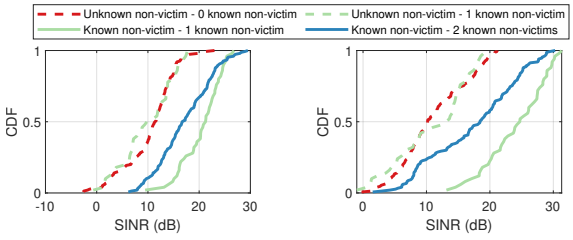

(a) Equal power **(b) Maximal throughput**

Figure 12: Distributions of SINRs at non-victim clients with partial channel feedbacks

Figures 11 and 12 depict the SINRs of the attacker, the victim, and non-victim clients with varying numbers of non-victim channels known at the attacker. Figures 11a and 11b indicate that the number of known non-victim clients has a negligible effect on the attacker and the victim for both power allocation strategies. Conversely, SINRs of unknown non-victim clients notably decrease compared to known counterparts, as shown in Figures 12a and 12b. We attribute this to the lack of null streams generated by the attacker for unknown non-victim clients, which makes them suffer from the interference of the relayed signals. Another observation is that the SINRs of known non-victim clients decrease as more non-victims are known at the attacker. This is because more clients involved in generating the precoding matrices make the precoding values for the same client across subcarriers less correlated, which yields received signals with lower powers. Since we assume that the attacker adjusts its transmit power to maintain the RSS at the victim at a constant level, the attacker needs to allocate higher transmit power per client. Consequently, the increased transmit power leads to higher leakage and lower SINRs at non-victim clients.

5.6 Comparative Analysis with the Malicious Client Eavesdropping Attack

We compare our proposed attack's eavesdropping efficiency with a representative MU-MIMO attack [39]. In [39], the attacker joins MU-MIMO communications as a client, exploiting forged channels to receive both intended victim messages and its own. With the message intended for itself as prior knowledge, the attacker cancels its signals to decode the victim's message. Following the system setting in [39], we consider an AP serving two clients. We collect 50 traces with 10 different settings in the office environment. We set the adjustable coefficient $w = 1$ for the attack in [39]. We refer to the attack in [39] as the *malicious client* method and our proposed attack as the *malicious relay* method in this section.

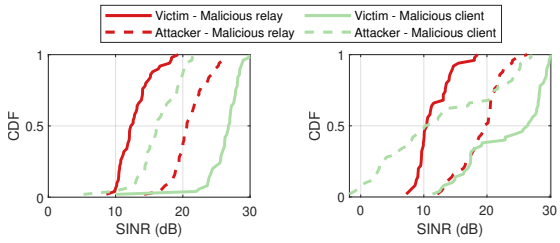

(a) Equal power **(b) Maximal throughput**

Figure 13: Distributions of average SINRs at the victims and the attacker with different attack methods

Figure 13 illustrates the SINRs at the attacker and the victim with two attack methods. In our method, the attacker's SINR is typically 8-10 dB higher than the victim's, whereas with the malicious client method, the attacker's SINR is generally 10-15 dB lower than the victim's. This difference stems from the attack methods' nature. In our proposed attack, the attacker's SINR is higher than the victim's because the signals received by the attacker are only intended for the victim, while at the victim they are a mix of the signals relayed by the attacker and the interference from the AP. Conversely, the malicious client method requires the attacker to estimate and cancel its own signals from received signals, introducing the unavoidable cancellation errors that decrease victim SINR. The victim's communication with the AP remains unaffected by the attacker, leading to a higher SINR for the victim compared to the attacker.

6 COUNTERMEASURES

The authors of prior research on eavesdropping attacks in MU-MIMO systems have proposed various countermeasures. In [39], the authors propose to use secret pilot values for channel sounding. In [29], a two-phase pilot commitment process is proposed to prevent unauthorized access to CSI. While these methods can defend against our attack, they require altering the existing communication protocols and can introduce extra control signal exchange overhead in MU-MIMO, which already has noticeable delays. Therefore, we evaluate the effectiveness of two representative features, AoA and CFO, used in physical-layer source authentication in detecting our attacks. These source authentication methods are fully compatible with existing protocols. They utilize metrics calculated during decoding and introduce minimal overhead.

6.1 Detection with Angle of Arrival

AoA describes signal arrival direction at the receiver, which can be estimated at multi-antenna devices with the multiple signal classification (MUSIC) algorithm [35]. Recent research has applied AoA profiles in detecting malicious activity in wireless networks [45, 49]. To detect our proposed attack, the AP can employ the MUSIC algorithm with CSIs from the feedback packets as input and monitor changes in AoA profiles for each client. Sudden deviations in a client's AoAs may indicate the attack initiation. This is because the channel measurements from feedback packets represent AP-attacker antenna channels, their AoA profiles may differ from AP-victim channels before the attack.

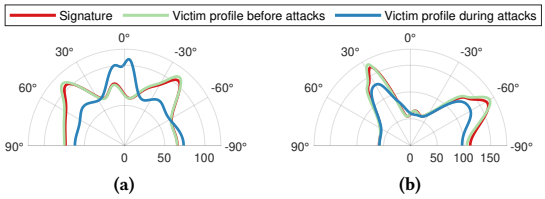

Figure 14: Example AoA spectra for detection

We evaluate the effectiveness of AoA-based detection with traces collected in Section 5.6. For each setting, we use 3 out of 5 traces to extract the AoA signature of the victim. The victim's channels before and after the attack in the remaining traces serve as input. We employ a simplified version of the method from [49], i.e., getting AoA spectra with the MUSIC algorithm and extracting local maximum angles as features. From Figure 14 we can see that in both examples, the AoA spectra closely match the victim's signature in the absence of an attack. During an attack, there are noticeable differences for cases such as Figure 14a. However, Figure 14b presents a challenging scenario missed by the detection method due to the close proximity of the attacker and victim. This AoA-based method achieves an accuracy of 90%, with an 80% true positive rate (TPR) and a 100% true negative rate (TNR).

In [45], the authors propose to assess a correlation metric between AoAs of feedback packets and angles of departure (AoD) of the reported channels to detect their proposed attack. In our proposed attack, since the reported channels will be AP-attacker antenna channels, the AoDs of reported channels and AoAs of attacked packets will be very close. Thus, we believe performances of the defense in [45] and our evaluation should be comparable.

6.2 Detection with Carrier Frequency Offset

CFO represents the carrier signal frequency difference between two devices. It is a ubiquitous phenomenon in wireless communication systems and is usually caused by oscillator drifts or Doppler shifts. In [24], CFO is employed as a signature for device authentication based on transmitter-receiver oscillator biases. To detect the proposed attack with CFO signatures, clients need to monitor CFO changes between the AP and themselves over time. Since CFO values are already estimated with pilots for successful decoding, reusing CFO as authentication signatures introduces minimal overhead. Sudden deviations in CFO values observed by a client can indicate the attack initiation and its victim status, as victim clients receive mixed signals from the attacker and AP during the attack, leading to combined CFO values due to oscillator drifts.

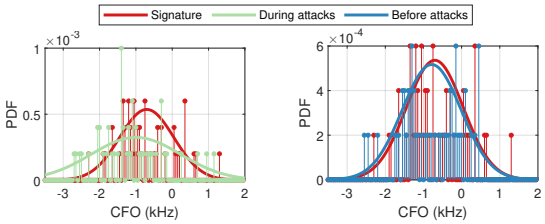

Figure 15: Example CFO distributions for detection

We extract CFO distributions from 50 traces for the same AP-client pair as the signature, comparing them to CFO distributions before and during attacks from the experiment in Section 5.6. Figure 15 displays probability density functions (PDF) and Gaussian approximations, which shows clear distinctions between clean signatures and observations during attacks. To evaluate the detection accuracy with CFO, we use 4 of the 5 traces per setting to create distribution profiles with and without the attack and the remaining one for testing. We determine the CFO observation result by comparing the likelihood of PDF functions at that CFO value. The CFO-based method reports a 65% accuracy, with a 40% TPR and a 90% TNR. The lower accuracy of the CFO-based method is due to significant overlap in CFO distributions.

7 CONCLUSION

In this paper, we introduce an active eavesdropping attack on MU-MIMO systems using a multi-antenna full-duplex device. Our attack comprises two phases: beamformed channel measurement manipulation and data stream relaying. We perform extensive experiments to evaluate the effectiveness of our attack under various settings, demonstrating its capability to successfully eavesdrop on AP-victim communications and bring the victims' secrecy capacity down to zero. We also investigate the feasibility of using physical-layer features to detect the proposed attack.

ACKNOWLEDGMENTS

We would like to thank our anonymous shepherd and reviewers for their insightful comments and constructive suggestions. This work is supported by NSF ECCS-2128567 and CNS-2007581 awards.

REFERENCES

[1] 3GPP. 2022. TS36.201 E-UTRA; LTE physical layer; General description. 3GPP.

[2] Ali Abedi and Omid Abari. 2020. Wi-Fi says "hi!" back to strangers!. In *Proceedings of the 19th ACM Workshop on Hot Topics in Networks*. 132–138.

[3] Ali Abedi and Deepak Vasisht. 2022. Non-cooperative Wi-Fi localization & its privacy implications. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking*. 570–582.

[4] Mayank Agarwal, Santosh Biswas, and Sukumar Nandi. 2015. Advanced stealth man-in-the-middle attack in WPA2 encrypted Wi-Fi networks. *IEEE communications letters* 19, 4 (2015), 581–584.

[5] Ehsan Aryafar, Narendra Anand, Theodoros Salonidis, and Edward W Knightly. 2010. Design and experimental evaluation of multi-user beamforming in wireless LANs. In *Proceedings of the sixteenth annual international conference on Mobile computing and networking*. 197–208.

[6] Ehsan Aryafar, Mohammad Amir Khojastepour, Karthikeyan Sundaresan, Sampath Rangarajan, and Mung Chiang. 2012. MIDU: Enabling MIMO full duplex. In *Proceedings of the 18th annual international conference on Mobile computing and networking*. 257–268.

[7] IEEE Standards Association. 1999. IEEE 802.11 a-1999 IEEE Standard for Telecommunications and Information Exchange Between Systems—LAN/MAN Specific Requirements—Part 11: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications: High Speed Physical Layer in the 5 GHz Band. IEEE.

[8] IEEE Standards Association. 2009. IEEE 802.11n-2009 IEEE Standard for Information technology– Local and metropolitan area networks– Specific requirements– Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 5: Enhancements for Higher Throughput. IEEE.

[9] IEEE Standards Association. 2013. IEEE 802.11ac-2013 IEEE Standard for Information technology–Telecommunications and information exchange between systems–Local and metropolitan area networks–Specific requirements–Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications–Amendment 4: Enhancements for Very High Throughput for Operation in Bands below 6 GHz. IEEE.

[10] IEEE Standards Association. 2017. IEEE 802.16-2017 IEEE Standard for Air Interface for Broadband Wireless Access Systems. IEEE.

[11] IEEE Standards Association. 2021. IEEE 802.11ax-2021 IEEE Standard for Information Technology–Telecommunications and Information Exchange between Systems Local and Metropolitan Area Networks–Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 1: Enhancements for High-Efficiency WLAN. IEEE.

[12] Sarankumar Balakrishnan, Pu Wang, Arupjyoti Bhuyan, and Zhi Sun. 2019. Modeling and analysis of eavesdropping attack in 802.11 ad mmWave wireless networks. *IEEE Access* 7 (2019), 70355–70370.

[13] Dinesh Bharadia and Sachin Katti. 2014. Fastforward: Fast and constructive full duplex relays. *ACM SIGCOMM Computer Communication Review* 44, 4 (2014), 199–210.

[14] Dinesh Bharadia and Sachin Katti. 2014. Full duplex MIMO radios. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. 359–372.

[15] Dinesh Bharadia, Emily McMilin, and Sachin Katti. 2013. Full duplex radios. In *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*. 375–386.

[16] Wireless CAT. 2022. List of MU-MIMO supported devices. (2022).

[17] Bo Chen, Vivek Yenamandra, and Kannan Srinivasan. 2015. FlexRadio: Fully flexible radios and networks. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*. 205–218.

[18] Jung Il Choi, Mayank Jain, Kannan Srinivasan, Phil Levis, and Sachin Katti. 2010. Achieving single channel, full duplex wireless communication. In *Proceedings of the sixteenth annual international conference on Mobile computing and networking*. 1–12.

[19] Max Costa. 1983. Writing on dirty paper (corresp.). *IEEE transactions on information theory* 29, 3 (1983), 439–441.

[20] Xuewei Feng, Qi Li, Kun Sun, Yuxiang Yang, and Ke Xu. 2023. Man-in-the-middle attacks without rogue AP: When WPAs meet ICMP redirects. In *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, 3162–3177.

[21] Scott Fluhrer, Itsik Mantin, and Adi Shamir. 2001. Weaknesses in the key scheduling algorithm of RC4. In *Selected Areas in Cryptography: 8th Annual International Workshop, SAC 2001 Toronto, Ontario, Canada, August 16–17, 2001 Revised Papers 8*. Springer, 1–24.

[22] GSMA. 2012. MIMO in HSPA: the Real-World Impact.

[23] Gerhard P Hancke. 2011. Practical eavesdropping and skimming attacks on high-frequency RFID tokens. *Journal of Computer Security* 19, 2 (2011), 259–288.

[24] Weikun Hou, Xianbin Wang, Jean-Yves Chouinard, and Ahmed Refaey. 2014. Physical layer authentication for mobile systems with time-varying carrier frequency offsets. *IEEE Transactions on Communications* 62, 5 (2014), 1658–1667.

[25] Mayank Jain, Jung Il Choi, Taemin Kim, Dinesh Bharadia, Siddharth Seth, Kannan Srinivasan, Philip Levis, Sachin Katti, and Prasun Sinha. 2011. Practical, real-time, full duplex wireless. In *Proceedings of the 17th annual international conference on Mobile computing and networking*. 301–312.

[26] Xingqin Lin, Jingya Li, Robert Baldemair, Jung-Fu Thomas Cheng, Stefan Parkvall, Daniel Chen Larsson, Havish Koorapaty, Mattias Frenne, Sorour Falahati, Asbjorn Grovlen, et al. 2019. 5G new radio: Unveiling the essentials of the next generation wireless access technology. *IEEE Communications Standards Magazine* 3, 3 (2019), 30–37.

[27] Lingjia Liu, Runhua Chen, Stefan Geirhofer, Krishna Sayana, Zhihua Shi, and Yongxing Zhou. 2012. Downlink MIMO in LTE-advanced: SU-MIMO vs. MU-MIMO. *IEEE Communications Magazine* 50, 2 (2012), 140–147.

[28] Yunlong Mao, Ying He, Yuan Zhang, Jingyu Hua, and Sheng Zhong. 2019. Secure TDD MIMO networks against training sequence based eavesdropping attack. *IEEE Transactions on Mobile Computing* 19, 12 (2019), 2916–2932.

[29] Yunlong Mao, Yuan Zhang, and Sheng Zhong. 2016. Stemming downlink leakage from training sequences in multi-user MIMO networks. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. 1580–1590.

[30] Ulrike Meyer and Susanne Wetzel. 2004. A man-in-the-middle attack on UMTS. In *Proceedings of the 3rd ACM workshop on Wireless security*. 90–97.

[31] Alexander M Ostrowski. 1952. Note on bounds for determinants with dominant principal diagonal. *Proc. Amer. Math. Soc.* 3, 1 (1952), 26–30.

[32] Hannaneh Barahouei Pasandi and Tamer Nadeem. 2021. LATTE: online MU-MIMO grouping for video streaming over commodity WiFi. In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*. 491–492.

[33] Yue Qiao, Ouyang Zhang, Wenjie Zhou, Kannan Srinivasan, and Anish Arora. 2016. PhyCloak: Obfuscating sensing from communication signals. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*. 685–699.

[34] Qualcomm. 2016. Exploring 5G new radio: Use cases, capabilities & timeline. (2016). https://www.qualcomm.com/content/dam/qcomm-martech/dm-assets/documents/heavyreading_whitepaper_-_exploring_5g_new_radio_use_cases_capabilities_and_timeine.pdf

[35] Ralph Schmidt. 1986. Multiple emitter location and signal parameter estimation. *IEEE transactions on antennas and propagation* 34, 3 (1986), 276–280.

[36] Matthias Schulz, Adrian Loch, and Matthias Hollick. 2014. Practical known-plaintext attacks against physical layer security in wireless MIMO systems. In *Network and Distributed System Security (NDSS) Symposium*.

[37] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. 2016. Practical MU-MIMO user selection on 802.11 ac commodity networks. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*. 122–134.

[38] David Tse and Pramod Viswanath. 2005. *Fundamentals of wireless communication*. Cambridge university press.

[39] Yu-Chih Tung, Sihui Han, Dongyao Chen, and Kang G Shin. 2014. Vulnerability and protection of channel state information in multiuser MIMO networks. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. 775–786.

[40] Yu-Chih Tung, Kang G Shin, and Kyu-Han Kim. 2016. Analog man-in-the-middle attack against link-based packet source identification. In *Proceedings of the 17th ACM International Symposium on Mobile Ad Hoc Networking and Computing*. 331–340.

[41] Mathy Vanhoef. 2021. Fragment and Forge: Breaking Wi-F Through Frame Aggregation and Fragmentation. In *30th USENIX Security Symposium (USENIX Security 21)*. 161–178.

[42] Mathy Vanhoef and Frank Piessens. 2013. Practical verification of WPA-TKIP vulnerabilities. In *Proceedings of the 8th ACM SIGSAC symposium on Information, computer and communications security*. 427–436.

[43] Mathy Vanhoef and Frank Piessens. 2017. Key reinstallation attacks: Forcing nonce reuse in WPA2. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 1313–1328.

[44] Mathy Vanhoef and Eyal Ronen. 2020. Dragonblood: Analyzing the Dragonfly Handshake of WPA3 and EAP-pwd. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 517–533.

[45] Sulei Wang, Zhe Chen, Yuedong Xu, Qiben Yan, Chongbin Xu, and Xin Wang. 2019. On user selective eavesdropping attacks in MU-MIMO: CSI forgery and countermeasure. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 1963–1971.

[46] Wei Wang, Raj Joshi, Aditya Kulkarni, Wai Kay Leong, and Ben Leong. 2013. Feasibility study of mobile phone Wi-Fi detection in aerial search and rescue operations. In *Proceedings of the 4th Asia-Pacific workshop on systems*. 1–6.

[47] Xiaoshan Wang, Yao Liu, Xiang Lu, Shichao Lv, Zhiqiang Shi, and Limin Sun. 2017. On eavesdropping attacks and countermeasures for MU-MIMO systems. In *MILCOM 2017-2017 IEEE Military Communications Conference (MILCOM)*. IEEE, 40–45.

[48] Hanan Weingarten, Yossef Steinberg, and Shlomo Shitz Shamai. 2006. The capacity region of the Gaussian multiple-input multiple-output broadcast channel. *IEEE transactions on information theory* 52, 9 (2006), 3936–3964.

[49] Jie Xiong and Kyle Jamieson. 2013. Securearray: Improving WiFi security with fine-grained physical-layer information. In *Proceedings of the 19th annual international conference on Mobile computing & networking*. 441–452.

[50] Qing Yang, Xiaoxiao Li, Hongyi Yao, Ji Fang, Kun Tan, Wenjun Hu, Jiansong Zhang, and Yongguang Zhang. 2013. BigStation: Enabling scalable real-time

signal processingin large MU-MIMO systems. *ACM SIGCOMM Computer Communication Review* 43, 4 (2013), 399–410.

[51] Chia-Yi Yeh and Edward W Knightly. 2021. Eavesdropping in massive MIMO: New vulnerabilities and countermeasures. *IEEE Transactions on Wireless Communications* 20, 10 (2021), 6536–6550.

[52] Taesang Yoo and Andrea Goldsmith. 2006. On the optimality of multiantenna broadcast scheduling using zero-forcing beamforming. *IEEE Journal on selected areas in communications* 24, 3 (2006), 528–541.

[53] Yong Zeng and Rui Zhang. 2016. Active eavesdropping via spoofing relay attack. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2159–2163.

[54] Lizhuo Zhang, Weijia Jia, Sheng Wen, and Di Yao. 2010. A man-in-the-middle attack on 3G-WLAN interworking. In *2010 International Conference on Communications and Mobile Computing*, Vol. 1. IEEE, 121–125.