CSIsnoop: Attacker Inference of Channel State Information in Multi-User WLANs

Xu Zhang
Rice University
Houston, Texas, USA
Xu.Zhang@rice.edu

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

ABSTRACT

Channel State Information (CSI) has been proposed to enhance physical layer security between a transmitter and a receiver because it decorrelates over half wavelength distances in rich scattering environments. Consequently, CSI was employed to generate passwords, to authenticate the source of packets, and to inject artificial noise to thwart eavesdroppers. However, in this paper, we present *CSIsnoop*, and show that an attacker can infer CSI in a multi-user WLAN, even if both channel sounding sequences from the access point and CSI measurement feedback from the clients are encrypted. The insights of *CSIsnoop* are that the CSI of clients can be computed based on transmit beamforming weights at the access point, and that the transmit beamforming weights can be estimated from downlink multi-user transmission. We implement *CSIsnoop* on a software defined radio and conduct experiments in various indoor environments. Our results show that on average *CSIsnoop* can infer CSI of the target client with an absolute normalized correlation of over 0.99, thereby urging reconsideration of the use of CSI as a tool to enhance physical layer security in multi-user WLANs.

CCS CONCEPTS

• **Networks** → **Mobile and wireless security**; **Network privacy and anonymity**; **Network experimentation**; *Wireless local area networks*;

KEYWORDS

CSIsnoop, Multi-User MIMO WLAN, Channel State Information Inference, Known-Transmitted-Symbol Attack

ACM Reference format:
Xu Zhang and Edward W. Knightly. 2017. *CSIsnoop*: Attacker Inference of Channel State Information in Multi-User WLANs. In *Proceedings of Mobihoc '17, Chennai, India, July 10-14, 2017,* 10 pages.
https://doi.org/http://dx.doi.org/10.1145/3084041.3084048

1 INTRODUCTION

Channel State Information (CSI) plays a key role in multi-user beamforming systems, because it enables an Access Point (AP) to increase throughput by concurrently sending multiple data streams to multiple clients. According to IEEE 802.11ac/af [7, 8], a known channel sounding sequence is broadcasted by the AP, from which clients compute and feedback their measurement results of the channel's effect on the known sequence. Based on the collected CSI, the AP can compute transmit beamforming weights and, for example, zero-force the signals of one client at other clients in order to eliminate inter-client interference [2, 24]. Moreover, CSI can be also used to increase network throughput by grouping clients with orthogonal channels [24, 25].

Besides data transmission, CSI has also been proposed to enhance network security, because it decorrelates over half a wavelength (several centimeters in 2.4/5 GHz WiFi) in rich scattering environments. Therefore, wireless devices can employ CSI for secret key establishment [12], which appears especially promising when there are limited resources or lacking key management infrastructure. In addition, CSI can be used as a signature to authenticate the source of packets, as each client in the network will have a unique CSI signature [9, 22]. Finally, CSI can be used to inject artificial noise orthogonal to the intended recipient to degrade an eavesdropper's channel [1, 5]. Such artificial noise is nulled at the clients by the AP so that the signal SINR at the clients will not be reduced.

Because of the importance of CSI, it has been proposed to encrypt CSI during the standard-defined explicit channel sounding process: either by encrypting the measurement feedback from the clients, or by encrypting the channel sounding sequence from the AP [19]. Therefore, a malicious node within range of the network cannot learn clients' CSI by overhearing their measurement feedback.

However, we discover that the above methods cannot ensure the confidentiality of CSI in multi-user MIMO WLANs. In particular, in this paper we describe *CSIsnoop*, a framework by which a *passive* attacker can infer the CSI of clients by overhearing their downlink beamforming transmission. The first step of *CSIsnoop* is to employ the knowledge of part of the transmitted symbols (e.g., the MAC header) and trains an adaptive filter to separate the different data streams at the multiple-antenna malicious node (Eve), which we term as a known-transmitted-symbol attack (a PHY analogy of the known-plaintext attack). The malicious node subsequently estimates her channel from the AP and combines it with the adaptive filter to compute the transmit beamforming weights that the AP must have used. Finally, *CSIsnoop* uses the estimated beamforming weights to compute the CSI of clients. We analyze *CSIsnoop* with various number of clients within the network, and also show that even if the AP encrypts the channel sounding sequence [19], it is still possible for the multiple-antenna malicious node to estimate CSI for both herself and the target client. Moreover, we discuss how

Mobihoc '17, July 10-14, 2017, Chennai, India

ACM ISBN 978-1-4503-4912-3/17/07…$15.00
https://doi.org/http://dx.doi.org/10.1145/3084041.3084048

an active adversary can accelerate the computing process by using a variant of *CSIsnoop* which we name *CSIsnoop+a*.

Our results reveal a fundamental conflict between using CSI to optimize PHY design and hiding CSI from malicious nodes, which urges reconsideration of the use of CSI to enhance physical layer security in multi-user WLANs. In particular, with *CSIsnoop*, we demonstrate that a malicious node can now break the following security schemes in multi-user WLANs:

- The malicious node can compute CSI of clients even if the AP encrypts the channel sounding sequence and the clients encrypt their CSI measurement feedback [19].
- With the computed CSI, the malicious node can further estimate the CSI-based password [12].
- With the computed CSI, the malicious node can fake CSI-based signatures [9, 22].
- With the computed CSI, the malicious node can remove most of the artificial noise and decode the overheard packets [1, 5].

In addition, we show that *CSIsnoop* can be also employed to degrade downlink and uplink throughput in the network. Specifically, we identify a new threat, which we term selective jamming, to the uplink multi-user transmission in next generation wireless standards [4, 6, 18]: once the malicious node obtains the CSI of a target client, she can selectively jam the client's data stream in uplink multi-user transmission, while not interfering with the data streams from all other clients. This is fundamentally different from current jamming techniques [14, 23], which treat all the concurrent data streams identically.

Finally, we implement *CSIsnoop* on WARP v3 [20], deploy a testbed, and conduct experiments in various indoor environments. Specifically, we consider that the AP has 2, 3, or 4 antennas and collect measurements from over 100,000 over-the-air transmissions. Our main experimental results can be summarized as follows:

- With the adversary's average signal SNR being 30 dB and the average condition number of the channel matrix between the AP and the adversary being 5, *CSIsnoop* can infer the target client's CSI with an absolute normalized correlation of over 0.99.
- The accuracy of *CSIsnoop* is related to the attacker's channel: if Eve is able to move to a location with high signal strength but small condition number of her CSI matrix from the AP, she can perform better by observing just a single frame than if she were able to observe multiple frames but in a less favorable location. This also provides hints of how attacks based on *CSIsnoop* can be mitigated.
- *CSIsnoop* enables the malicious node to compute over 85% of the CSI-based password.
- For selective jamming, *CSIsnoop* creates a 20 dB average increase in interference to the uplink data stream of the target client compared to other clients.

The rest of the paper is organized as follows. Sec. 2 introduces our threat model. Sec. 3 and Sec. 4 describe the principles of *CSIsnoop* and its variant *CSIsnoop+a*. We discuss implementation in Sec. 5 and experimental evaluations in Sec. 6. Several attack applications of *CSIsnoop* are explored in Sec. 7. Sec. 8 discusses related work and Sec. 9 concludes the paper.

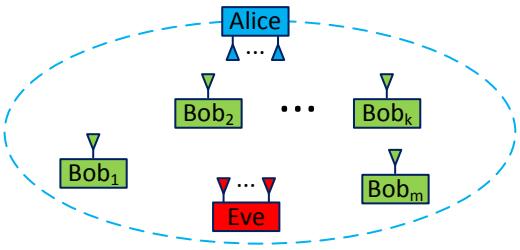

Figure 1: Threat model of *CSIsnoop*.

2 THREAT MODEL

In this paper, we consider the threat model illustrated in Fig. 1. Alice is a multiple-antenna multi-user AP and each Bob is a legitimate single-antenna client. The network has multiple clients, i.e., multiple "Bobs". Within the scope of this paper, we consider that explicit channel sounding is employed (as is the case with IEEE 802.11ac/af [7, 8]). That is Alice transmits beacons such that the Bobs can measure CSI and send the results back to Alice. However, each Bob can encrypt his CSI feedback in order to prevent malicious nodes from directly overhearing his CSI measurements. We also consider the case that Alice encrypts her broadcasted channel sounding sequence [19] (i.e., the beacon sequence is only known to Alice within the network).

After acquiring Bobs' CSI H_{AB}, Alice uses zero-forcing beamforming (ZF-BF) to compute her transmit beamforming weights, which is $W_A = H_{AB}^{\dagger} = (H_{AB}^H H_{AB})^{-1} H_{AB}$. ZF-BF has been widely used because it can asymptotically achieve network capacity with relatively low computational complexity [2, 11, 19, 24]. Nonetheless, the *CSIsnoop* framework can be extended to other different beamforming algorithms, e.g., conjugate beamforming. Moreover, we assume that Alice uses all of her antennas to transmit no matter what the number of data streams is. In other words, Alice will fully utilize her antenna resources to boost the downlink network throughput.

There is a malicious node, Eve, within range of Alice. We consider a rich scattering environment typical of an indoor WLAN so that Eve can overhear signals of Alice's downlink transmission to all Bobs. And Eve has the same number of antennas as Alice (Alice's antenna number can be known from the network control signals).

We further assume that Eve knows a subsequence of the symbols that are transmitted by Alice for each Bob's downlink data packet. These known symbols can be portions of the MAC packet header that follow a pre-defined format in the standards [7, 8]. Finally, we assume that Eve knows which Bobs are included in a specific downlink transmission. This can be done by Eve overhearing the control signals broadcasted by Alice before the data transmission, or the ACK packets sent from the Bobs after the data transmission.

3 CSISNOOP

In this section, we describe *CSIsnoop*, a technique by which a *passive* adversary can infer the CSI of different clients by overhearing downlink multi-user transmission. We begin by analyzing the base case where the number of data streams from Alice to the Bobs equals the number of antennas at Alice. Then we generalize to scenarios where the number of data streams is smaller. Finally, we

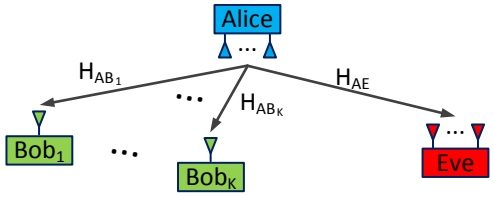

Figure 2: Base case analysis of *CSIsnoop*, where Bobs' data stream number equals the number of antennas at Alice in one downlink transmission.

show that Eve is able to estimate her channel (an important step of *CSIsnoop*) even if the channel sounding sequences are encrypted by Alice.

3.1 Base Case Analysis of $K \times K$

To begin with, we first consider the base case of *CSIsnoop* as shown in Fig. 2, where Alice has K antennas and beamforms K data streams to K Bobs in one downlink transmission. In this case, the pseudo-inverse computation of ZF-BF is simplified to matrix inversion. Thus H_{AB} and W_A uniquely determine each other.

We denote the channel between Alice and Bob_j to be H_{AB_j}, and $H_{AB} = [H_{AB_1}^T, H_{AB_2}^T, \dots]^T$. Eve also has K antennas, and the channel between Alice and Eve is H_{AE}. X_j is a $1 \times L$ vector that contains L symbols of Bob_j's data stream. Therefore, for all K Bobs, $X = [X_1^T, \dots, X_K^T]^T$ represents the $K \times L$ transmitted symbols from Alice. Similarly, we denote Y_j as the L symbols received by the j^{th} antenna of Eve, and $Y = [Y_1^T, \dots, Y_K^T]^T$ as the $K \times L$ symbols overheard by Eve. Therefore, Y can be represented as

$$Y = H_{AE} W_A P X + N, \tag{1}$$

where W_A is Alice's transmit beamforming weights, P is the transmit power scaling matrix and $P = diag(\{\sqrt{p_1}, \dots, \sqrt{p_K}\})$, and N is random noise.

In order to compute H_{AB} between Alice and the Bobs, Eve needs to first estimate W_A used by Alice. We identify the known-transmitted-symbol attack that can be employed by *CSIsnoop* to accomplish this. After that, Eve computes H_{AB} based on the relationship between H_{AB} and W_A. We divide the whole process into 4 steps and discuss each of them as follows.

(1) Known-Transmitted-Symbol Attack. *CSIsnoop* employs known-transmitted-symbol attack to compute an adaptive filter based on only X and Y, and uses this filter to further compute $H_{AE} W_A P$. Specifically, if Eve knows the $1 \times L$ vector X_j of Bob_j (as discussed in Sec. 2, these known symbols can be part of the MAC header that follow a pre-defined format in the standards [7, 8]), she can compute a $1 \times K$ receive beamforming vector $W_{E,j}^{(1)}$ such that

$$E\{\|e_j^{(1)}\|\} = E\{\|X_j - W_{E,j}^{(1)} Y\|\} \tag{2}$$

is minimized. By taking the derivative of $E\{\|e_j^{(1)}\|\}$ over $W_{E,j}^{(1)}$ and setting it to zero, Eve obtains

$$W_{E,j}^{(1)} = X_j Y^\dagger, \tag{3}$$

where $Y^\dagger$ is the Moore-Penrose pseudo-inverse of Y.

Therefore, if Eve launches known-transmitted-symbol attack targeting Bobs' K data streams, she can compute

$$W_E^{(1)} = X Y^\dagger, \tag{4}$$

where $W_E^{(1)} = [{W_{E,1}^{(1)}}^T, \dots, {W_{E,K}^{(1)}}^T]^T$. By ignoring the random noise N, Eve can estimate

$$H_{AE} W_A P = {W_E^{(1)}}^{-1}. \tag{5}$$

Here $W_E^{(1)}$ is a $K \times K$ square matrix because of the K data streams. So Eve can directly compute its inverse.

However, the above process will amplify the random noise N when calculating $Y^\dagger$, which degrades the accuracy of the estimation of $H_{AE} W_A P$. Therefore, instead of Eq. (2), *CSIsnoop* minimizes

$$E\{\|e^{(2)}\|\} = E\{\|Y - W_E^{(2)} X\|\}, \tag{6}$$

which leads to

$$W_E^{(2)} = Y X^\dagger, \tag{7}$$

and

$$H_{AE} W_A P = W_E^{(2)}. \tag{8}$$

In the following, we use W_E to represent $W_E^{(2)}$.

(2) Estimation of H_{AE}. It can be observed from Eq. (8) that in order to compute $W_A P$, Eve needs to estimate H_{AE} first. As explicit channel sounding is employed within the network, if the channel sounding sequences broadasted from Alice are not encrypted, Eve can use them to estimate H_{AE} directly (even though these sequences are designed for the Bobs to measure their CSI). The computation of H_{AE} at Eve under encrypted channel sounding sequences is discussed in Sec. 3.3.

(3) Computation of $W_A P$. With W_E and H_{AE}, Eve can compute

$$W_A P = H_{AE}^{-1} W_E. \tag{9}$$

This is because Eve has the same number of antennas as Alice and H_{AE} is a $K \times K$ square matrix. Thus its inverse can be directly computed.

(4) Computation of H_{AB}. ZF-BF computes the transmit beamforming weights as the pseudo-inverse of the CSI matrix. For the base case, because the number of data streams equals the number of antennas at Alice, the pseudo-inverse computation is matrix inversion. As a result, Eve can compute

$$P^{-1} H_{AB} = P^{-1} W_A^{-1} = W_E^{-1} H_{AE}. \tag{10}$$

It should be noted that the j^{th} row of $P^{-1} H_{AB}$ (which is $\frac{1}{\sqrt{p_j}} H_{AB_j}$) and the j^{th} row of H_{AB} are in the same sub-space. Therefore, even though Eve does not know P, she can still locate the signal sub-space from Alice to Bob_j. This already provides enough information for Eve to learn the relative relationship among the CSI of different Bobs, and to further break various CSI-based security mechanisms and decrease network throughput as discussed in Sec. 7.

3.2 Generalization to $M \times K$ with $M < K$

In the above base case analysis, we assume that the number of data streams from Alice to the Bobs in one downlink transmission equals the number of antennas at Alice. Therefore, Eve can directly compute the inverse of $W_A P$ in Eq. (10). However, in practice, it is

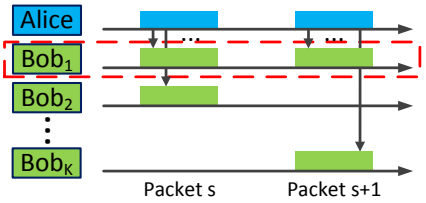

Figure 3: *CSIsnoop* combines the computation of multiple downlink transmissions to compute the CSI of a target Bob.

possible that the number of data streams in one downlink transmission is smaller, e.g., not all Bobs are backlogged. Under such circumstances, while Eve can still compute Alice's transmit beamforming weights W_A, H_{AB} is no longer a square matrix and the sub-space spanned by each of its rows can be no longer uniquely determined by its pseudo-inverse. In the following, we discuss how *CSIsnoop* addresses this problem by overhearing multiple downlink transmissions.

We first use an example in Fig. 3 to illustrate Eve's computing process. Suppose that there are M_s and M_{s+1} Bobs involved in transmission s and $s+1$, respectively. And both M_s and M_{s+1} are smaller than K. Because Eve has the same number of antennas as Alice, she can always compute $W_A P$ from Eq. (9). And we denote $W_{A,s}P_s = [V_{s,1}, \ldots, V_{s,M_s}]$ and $W_{A,s+1}P_{s+1} = [V_{s+1,1}, \ldots, V_{s+1,M_{s+1}}]$ to be the multiplication of transmit beamforming weights and transmit power scaling matrix at Alice for downlink multi-user frame s and $s+1$, respectively.

If Bob_1 is included in both transmissions, because ZF-BF transmits the data of all other Bobs into the null space of H_{AB_1} of Bob_1, Eve can obtain

$$H_{AB_1} V_{s,i} = H_{AB_1} V_{s+1,i} = 0, \ \forall i \geq 2. \tag{11}$$

Consequently, when the total number of different Bobs that are included in these 2 transmissions is no smaller than K, Eve can have at least $K-1$ uncorrelated vectors V from Eq. (11) such that $H_{AB_1} V = 0$, which enables Eve to locate the sub-space spanned by the $1 \times K$ vector H_{AB_1}.

Therefore, when the number of data streams in one downlink transmission is smaller than K, for Eve to compute H_{AB_j}, she needs to overhear multiple downlink multi-user transmissions that (1) each includes Bob_j, and that (2) the total number of different Bobs that are involved should be no smaller than K. However, these multiple transmissions need not be consecutive as long as H_{AB_j} remains stable. In comparison, $H_{AB_{i\neq j}}$ of other Bobs can even vary during the transmissions. In fact, as can be observed from Eq. (11), a changing $H_{AB_{i\neq j}}$ will only increase the number of different V's that Eve can have and thereby enable Eve to locate the sub-space of H_{AB_j} more quickly. H_{AE} can also change as long as its variation can be detected by Eve so that Eq. (9) remains accurate.

3.3 Estimating H_{AE} with Encrypted Channel Sounding Sequences

For explicit channel sounding, in order to prevent Eve from knowing Bobs' CSI by overhearing their measurement feedback, we can either (1) ask each Bob to encrypt his sounding feedback, or (2) ask Alice to encrypt her channel sounding sequences. However, as

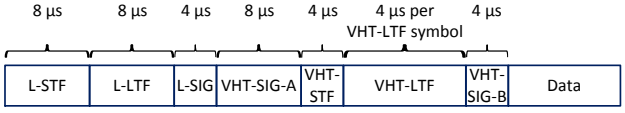

Figure 4: VHT PPDU format of IEEE 802.11ac [7] and IEEE 802.11af [8] (with different length of each field). The Null Data Packet for explicit channel sounding has the same format except that there is no data field.

discussed in Sec. 3.1 and Sec. 3.2, encrypting Bob's measurement feedback cannot stop Eve from computing Bob's CSI, because Eve can use the known channel sounding sequence to estimate H_{AE} and then compute H_{AB} based on *CSIsnoop*. In the following, we show that even if Alice encrypts her channel sounding sequences, Eve can still estimate H_{AE} and thereby compute H_{AB}.

We consider that Alice uses the *CSIsec* scheme [19] to encrypt her channel sounding sequences. Particularly, in each sub-carrier, Alice broadcasts a random symbol R instead of the known symbol D. Therefore, Bob_j and Eve measure their channel as $(R/D)H_{AB_j}$ and $(R/D)H_{AE}$, respectively. Because Alice knows R/D, she can remove it from Bob_j's measurement feedback. In contrast, Eve does not know R/D and thereby cannot obtain H_{AE}.

The key reason that Eve can estimate H_{AE} in *CSIsnoop* is because she now has multiple antennas instead of only a single antenna as in the discussion of [19]. In particular, *CSIsnoop* combines the measured $(R/D)H_{AE}$ with the L-LTF field defined in IEEE 802.11ac/af to compute H_{AE}. Fig. 4 shows the standard-defined packet format. The L-LTF field contains the long training sequence of IEEE 802.11a/b/g, and is designed to make IEEE 802.11ac compatible with IEEE 802.11a/b/g. However, to avoid un-intentional beamforming, the L-LTF field is sent with dynamic cyclic shift, i.e., a different phase shift is added to the signals sent by each of Alice's antennas. Such phase shifts are pre-defined in the standard and publicly known. Moreover, the L-LTF field is used by other clients to decode the following L-SIG field and thereby cannot be encrypted by Alice.

We still consider that both Alice and Eve have K antennas. In addition, we denote the known channel sounding symbol and its encrypted version from the i^{th} antenna of Alice to be D_i and R_i, respectively. Therefore, from the channel sounding phase between Alice and the Bobs, Eve can compute

$$G_{AE} = H_{AE}\Gamma, \tag{12}$$

where $\Gamma = diag(\{R_1/D_1, \ldots, R_K/D_K\})$. Suppose that the dynamic cyclic shift added to the i^{th} antenna of Alice is β_i. Thus from the L-LTF field of the *same* channel sounding packet, Eve can estimate

$$F_{AE} = H_{AE}B = H_{AE}[\beta_1, \ldots, \beta_K]^T. \tag{13}$$

Combining Eq. (12) and Eq. (13), Eve can obtain

$$G_{AE}^{-1}F_{AE} = \Gamma^{-1}B = [\frac{D_1}{R_1}\beta_1, \ldots, \frac{D_K}{R_K}\beta_K]^T. \tag{14}$$

Since Eve knows $B = [\beta_1, \ldots, \beta_K]^T$, she can solve Γ^{-1} from Eq. (14) and thereby compute $H_{AE} = G_{AE}\Gamma^{-1}$.

Therefore, whether or not Alice encrypts her channel sounding sequences, Eve can always estimate H_{AE} and further compute H_{AB}.

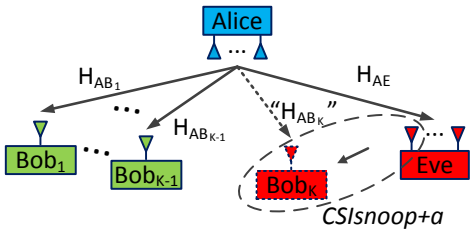

Figure 5: For *CSIsnoop+a*, Eve becomes active and joins in the downlink transmission, which enables Eve to compute H_{AB} more quickly in certain scenarios.

4 CSISNOOP+A

In the discussion in Sec. 3, Eve is completely passive and computes H_{AB} by only overhearing the downlink transmission from Alice to the Bobs. In this section, we describe a variant of *CSIsnoop* which we name *CSIsnoop+a*. For *CSIsnoop+a*, Eve becomes active and joins in the downlink multi-user transmission, which enables Eve to compute H_{AB} more quickly in certain scenarios.

One of the examples is shown in Fig. 5, where there are $K-1$ single-antenna Bobs within the network. As a result, the maximum number of different Bobs that are included in multiple transmissions can only be $K-1$. If Eve uses *CSIsnoop* to compute H_{AB_j} of Bob_j, at least one $H_{AB_{i\neq j}}$ of Bob_i need to change during the overheard multiple transmissions, so that Eve still has at least $K-1$ uncorrelated vectors V for Eq. (11).

In comparison, if *CSIsnoop+a* is employed, Eve can locate H_{AB_j} much more quickly. As shown in Fig. 5, Eve now fakes the legitimate client Bob_K by using her first antenna and participates in the downlink multi-user transmission, i.e., she sends channel sounding measurements back to Alice and asks Alice to beamform downlink data to the faked Bob_K (e.g., Eve can setup a remote server and ask the server to send data to the faked Bob_K). After that, Eve still uses her K antennas to overhear the K data streams, except that now she only needs to launch $K-1$ known-transmitted-symbol attacks, because the first row of W_E in Eq. (7) can be directly set to $[0,\ldots,0,1/\sqrt{pK}]$. Then Eve can obtain $P^{-1}H_{AB}$ by following step 2 to step 4 in Sec. 3.1.

In general, if there are N Bobs in the network and $N<K$, Eve can use her $K-N$ antennas to pretend to be $K-N$ clients and join in the downlink multi-user transmission. However, Eve still only needs to have K antennas in total. What is more, Eve can pretend to be legitimate clients using a subset of her antennas while at the same time overhear multiple downlink transmissions. Therefore, Eve does not need to have precise timing control of when her downlink data of the faked clients should arrive at the AP.

5 IMPLEMENTATION

Access Point (Alice) and Legitimate Clients (Bobs). We implement the functions of beamforming transmission of Alice and the Bobs in the software defined radio WARP v3 [20]. In particular, Alice broadcasts the channel sounding sequence defined in IEEE 802.11ac [7] from each of her antennas, which is used by the Bobs to measure their channel. After Alice receives the measurement feedback from the Bobs, she uses ZF-BF to calculate her transmit beamforming weights and beamforms Bobs' downlink data packet.

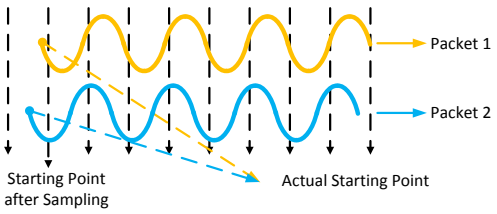

Figure 6: Fractional timing offset due to ADC sampling.

We also implement the *CSIsec* scheme [19], with which Alice encrypts the channel sounding sequence by multiplying it with a random complex number in each sub-carrier and for each of her antennas.

The transmission has a bandwidth of 20 MHz and includes 64 OFDM sub-carriers (based on IEEE 802.11ac). Bobs' data are modulated by BPSK or 4-QAM. And we use different WARP v3 boards to implement Alice and the Bobs so they are not clock-synchronized.

***CSIsnoop* and *CSIsnoop+a*.** We also implement all steps and calculations of *CSIsnoop* and *CSIsnoop+a* as described above in WARP v3. Specifically, when Eve overhears the downlink transmission, she first uses the normal method [15] to correct the timing and carrier frequency offset (due to the different oscillator frequencies between devices) based on the L-STF and the L-LTF field of each packet. After that, even though Eve has not computed the transmit beamforming weights W_A at Alice yet, she can still track and correct the residue carrier frequency offset with the aid of the pilot sub-carriers (these are reserved sub-carriers with pre-defined pilots throughout the entire packet [7, 8]). Suppose that the pilot is α. According to the standards, it is the same for all data streams. Therefore, what Eve overhears can be represented as

$$[Y_1,\ldots,Y_k]^T = H_{AE}W_AP\left([\alpha,\ldots,\alpha]_{1\times data\ stream\ \#}\right)^T. \quad (15)$$

Even if $H_{AE}W_AP$ is not known by Eve, for her j^{th} antenna, Y_j/α should have a fixed phase over time if there is no residue carrier frequency offset (but it can have different values across Eve's antennas). Therefore, by tracking the phase of every Y_j/α, Eve can detect and correct the residue carrier frequency offset.

Eve can subsequently begin to compute $H_{AE}W_AP$ by known-transmitted-symbol attack (Eq. (7)). In our implementation, we directly minimize the mean square error for this computation. While this has computational complexity of $O(n^3)$ (n is the size of the matrix X), it is optimal. Alternatively, we can use the iterative Least Mean Square algorithm to reduce computational complexity [16].

Finally, after Eve overhears multiple packets within channel coherence time of H_{AB_j} (during which H_{AB_j} stays stable), she can combine her estimates together to increase accuracy. However, Eve cannot simply compute the average value, because different packets will have different (1) transmit power scaling factors and (2) fractional timing offsets. The impact of scaling factors has already been discussed in Sec. 3.1. In the following, we analyze the impact of fractional timing offset.

Fractional timing offset is mainly due to the ADC sampling at the receiver. Particularly, in wireless networks the correlation computation of the training sequences in the L-STF and the L-LTF field is used to determine the start of a packet among a series of sampling points [15]. Because this computation occurs in the

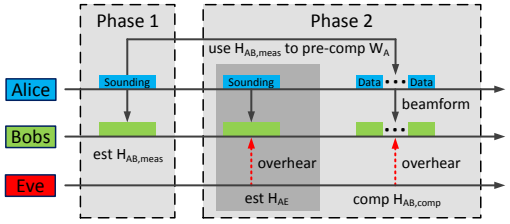

Figure 7: Experiment timeline.

digital domain, the actual start of the packet may deviate from the computed point with a maximum error of $\frac{\Delta}{2}$, where Δ is the ADC sampling interval (as shown in Fig. 6). Therefore, even if the physical channel stays unchanged, the actual channel that the packets go through will be different. According to the Fourier transform, delay in the time domain translates into rotation over sub-carriers in the frequency domain. As a result, each estimated H_{AB} at Eve will also have an additional unknown phase rotation.

To address the different transmit power scaling factors and fractional timing offsets, when Eve has a set of observations $\mathscr{H}_{AB_j} = [H_{AB_j}^{(1)\ T}, H_{AB_j}^{(2)\ T}, \dots]^T$, she searches $\hat{H}_{AB_j}$ that maximizes

$$\sum_{\forall i} \|H_{AB_j}^{(i)} \hat{H}_{AB_j}^H\| \tag{16}$$

with constraint $\|\hat{H}_{AB_j}\| = 1$. The evaluation in Sec. 6.4 shows that this will lead to a significantly more accurate result compared to simply calculating the average.

6 EXPERIMENTAL EVALUATION

In this section, we conduct experiments in various indoor environments to evaluate the performance of *CSIsnoop* and *CSIsnoop+a*. We first describe our experimental setup. Then we discuss the accuracy of the estimated H_{AB} as well as the impact of various factors beginning with the base case, followed by the generalized scenarios. Finally, we study how accurately Eve can estimate H_{AE} when Alice encrypts her channel sounding sequence.

6.1 Experimental Setup

We use our WARP v3 implementation to conduct experiments in typical lab, office, and apartment environments. Alice is configured to have 2, 3, or 4 antennas and beamform data to up to 4 single-antenna Bobs. Eve has the same number of antennas as Alice and stays within Alice's transmitting range. Therefore, Eve can overhear Bobs' downlink packets and use *CSIsnoop* to compute H_{AB}. During the experiments, we collect more than 100,000 over-the-air packets.

Due to the extra delay of the WARPLab framework [21], in order to ensure that the experiment is finished within channel coherence time (during which the channel stays stable), we divide it into 2 phases. In the first phase, we continuously measure $H_{AB,meas}$ for over 100 ms. In the second phase, we first ask Alice to broadcast the channel sounding sequence (to emulate the explicit channel sounding process), by which Eve can estimate H_{AE}. After that, Alice sends the pre-computed downlink beamforming data packets to the Bobs, which are based on the previously collected $H_{AB,meas}$. Eve launches known-transmitted-symbol attack targeting these downlink packets and uses *CSIsnoop* to compute $H_{AB,comp}$. The timeline of the experiment is shown in Fig. 7.

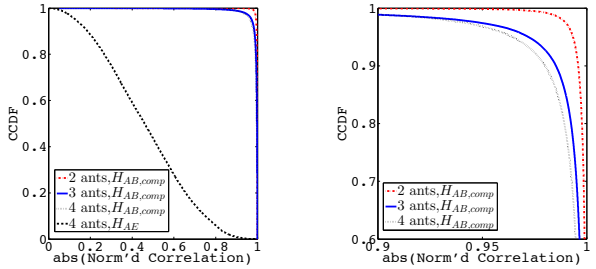

Figure 8: CCDF of the absolute value of normalized correlation between $H_{AB,meas}$ and $H_{AB,comp}$, and $H_{AB,meas}$ and H_{AE}. (b) is the zoomed-in upper-right portion of (a).

6.2 Accuracy of *CSIsnoop*

We first define a metric for estimation accuracy. As discussed in Sec. 3.1, Eve only needs to compute the sub-space spanned by H_{AB_j} instead of its exact value. Therefore, it is natural to use the degree of correlation as the evaluating metric. In particular, we use the absolute value of the normalized correlation, which is defined as

$$C = \frac{|H_{AB_j,comp} \cdot H_{AB_j,meas}^H|}{\|H_{AB_j,comp}\| \cdot \|H_{AB_j,meas}\|} \tag{17}$$

for H_{AB_j} of Bob_j. C ranges from 0 to 1. $C = 0$ indicates that $H_{AB_j,comp}$ and $H_{AB_j,meas}$ are orthogonal, while $C = 1$ indicates that $H_{AB_j,comp}$ and $H_{AB_j,meas}$ are perfectly correlated. Therefore, the closer to 1 is C, the more accurate the estimation is by Eve.

We assume that Eve uses 20 known symbols for the known-transmitted-symbol attack according to [16]. And Fig. 8(a) displays the complementary cumulative distribution function (CCDF) of C when Eve's average signal SNR is 30 dB and when the average condition number of H_{AE} is 5 (the influence of signal SNR and condition number are discussed later in Sec. 6.3). As a baseline for comparison, we also plot the CCDF of correlation between $H_{AB_j,meas}$ and H_{AE}: if *CSIsnoop* is not employed, we suppose that Eve uses H_{AE} to infer H_{AB_j}. The results indicate that *CSIsnoop* enables Eve to locate the sub-space spanned by H_{AB_j} with very high accuracy. For all 2-, 3-, and 4-antenna Alice, the average value of C largely increases from 0.46 to over 0.99.

We further zoom in the upper-right portion of Fig. 8(a) and show the details in Fig. 8(b). For 2-antenna Alice, it can be seen that over 99% of Eve's CSI inferences yield C larger than 0.99. However, the estimation accuracy decreases with the number of antennas at Alice. The main reason is that more noise will be included into the computation when Alice has more antennas.

6.3 Impact of H_{AE}

In the following, we study the influence of Eve's relative position to Alice via H_{AE} on *CSIsnoop*'s accuracy. H_{AE} is important because (1) if Eve is mobile or there are multiple Eves, *CSIsnoop* can search for places with favorable H_{AE} to reduce the estimation error; and (2) Alice can also use the connection between H_{AE} and *CSIsnoop*'s accuracy to design schemes to prevent H_{AB} from being computed.

Strength of H_{AE}. On average, the strength of H_{AE} determines the signal SNR at Eve (while for each overheard packet, the SNR

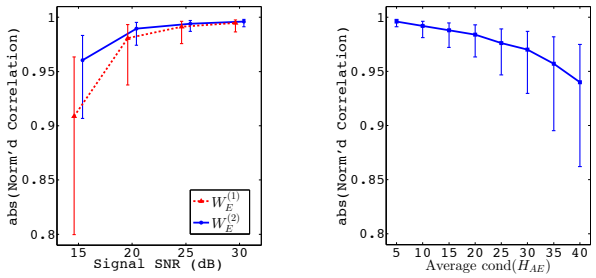

Figure 9: The variation of C (median with 25th/75th percentile) over (a) the average signal SNR at Eve and (b) the average condition number of H_{AE}. Alice has 4 antennas.

is also related to Alice's transmit beamforming weights). To examine how SNR impacts Eve's CSI inference accuracy, we consider measurements only from scenarios in which the average condition number of H_{AE} is 5, and plot the median of C (with 25th and 75th percentile) under different SNR in Fig. 9(a). We do not use the mean and standard deviation here because the distribution of C is highly non-normal (as shown in Fig. 8). It can be observed from the blue solid curve (when $W_E = W_E^{(2)}$ is employed by *CSIsnoop*) that, the median of C decreases to below 0.96 when the SNR becomes smaller than 15 dB. On the other side, when the SNR is over 25 dB, increasing the SNR further does not lead to a large increase of C.

Moreover, as discussed in Sec. 3.1, there are 2 ways for Eve to compute the adaptive filter based on known-transmitted-symbol attack. Therefore, Fig. 9(a) also depicts the median of C when $W_E = W_E^{(1)}$ is employed. It can be seen that $W_E^{(2)}$ consistently outperforms $W_E^{(1)}$ in accuracy. Moreover, the difference between $W_E^{(2)}$ and $W_E^{(1)}$ increases when the signal SNR becomes smaller: from 30 dB to 15 dB, the difference between their median increases from 0.001 to 0.05; and at 15 dB, the 25th percentile of C further reduces to below 0.8 when $W_E^{(1)}$ is used. The main reason is that during the computation of $W_E^{(1)}$, the random noise in the signals are amplified. Thus the reduction in estimation accuracy increases with larger noise strength (or equivalently, smaller signal SNR).

Condition Number of H_{AE}. After $W_E = W_E^{(2)}$ (Eq. (7)) is estimated from known-transmitted-symbol attack, *CSIsnoop* computes Alice's transmit beamforming weights as $W_A P = H_{AE}^{-1} W_E$. Therefore, the accuracy of W_A is also related to the condition number of H_{AE}. To analyze this factor, we consider measurements only from scenarios in which the average signal SNR is 30 dB, and plot the median of C with different condition number of H_{AE} in Fig. 9(b). In our experiments, we use the 2-norm condition number of H_{AE}, which is defined as the ratio of the largest singular value of H_{AE} to the smallest. It can be seen that C decreases when H_{AE} has a larger condition number. In particular, the estimation accuracy becomes worse than the 15 dB SNR point in Fig. 9(a) when the average condition number of H_{AE} exceeds 35.

Therefore, if Eve wants to increase her accuracy, she needs to search for favorable places where the signal SNR is large but the condition number of H_{AE} is small. Given that Alice is often fixed, one of the best strategies for Eve is to find a favorable place near Alice. Moreover, when Eve is near Alice, the change of H_{AE} due to environmental mobility also tends to be smaller.

6.4 Number of Known Symbols and Overheard Packets

Here, we evaluate whether the *CSIsnoop*'s inference accuracy is improved when Eve has more observations. Specifically, Eve can (1) still overhear just one transmission but have more known symbols, or (2) overhear multiple packets within the channel coherence time of H_{AB_j} and combine the observations.

More Known Symbols. When Eve overhears just one transmission but has more known symbols, she can calculate W_E in Eq. (7) with smaller error and thereby increase her estimation accuracy. To study this, we suppose that Eve knows from 10 to 50 symbols and plot the change of C in Fig. 10(a).

It can be observed that when Eve is in a favorable location (where the signal SNR is high while the condition number of H_{AE} is small), having more known symbols does not lead to a large increase of C, especially when the number of known symbols is over 20. This is because *CSIsnoop* already achieves a very accurate estimation with as few as only 10 known symbols (median of C is 0.993). In contrast, when Eve is restricted to a non-favorable location, the estimation accuracy keeps improving when Eve has more known symbols: from 10 to 50 known symbols, the median of C increases from 0.95 to 0.98. Nonetheless, even with 50 known symbols, the estimation is still less accurate compared to the results when Eve is in a favorable place but has only 10 known symbols.

Overhearing Multiple Packets. When Eve overhears multiple transmissions, she can combine them to improve accuracy. Fig. 10(b) and Fig. 10(c) depict C when the number of overheard packets increases from 1 to 5. Specifically, *SubSpaceSearch* indicates that Eve maximizes Eq. (16) when combining the multiple observations, while *SimpAvg* indicates that Eve simply calculates the average.

Similar to Fig. 10(a), it can be observed that when Eve is in a favorable location, the computation based on 1 overheard transmission is already very accurate, and repeated observations do not lead to a large increase in C. In comparison, when Eve is in a non-favorable location, accuracy significantly improves with more overheard transmissions. Yet, again, even with 5 transmissions, the accuracy is still worse than when Eve is in a favorable place but overhears just 1 transmission. Therefore, compared to increasing the number of known symbols or overhearing multiple transmissions, it is more important for Eve to have a favorable H_{AE}.

Furthermore, it can be observed from Fig. 10(b) and Fig. 10(c) that simply averaging the results of multiple overheard packets can actually lead to worse estimation accuracy. This is mainly due to the random phase rotation from fractional timing offset that is added to every single observation. For example, if for 2 overheard packets, the computed channel are H_{AB} and $e^{j\pi} H_{AB}$, *SimpAvg* computes the average value to be $H_{est} = 0$, while *SubSpaceSearch* maximizes $\|H_{est} H_{AB}^H\| + \|H_{est} (e^{j\pi} H_{AB})^H\|$ and obtains $H_{est} = H_{AB}$.

6.5 Number of Clients and Data Streams

In the following, we evaluate the performance of *CSIsnoop* in the generalized scenario of Sec. 3.2 and *CSIsnoop+a* of Sec. 4. Table 1 lists the detailed description and Fig. 11 plots the results.

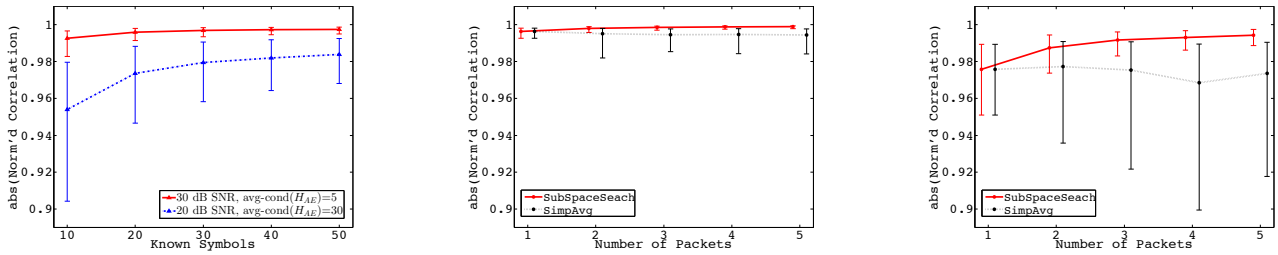

(a) (b) (c)

Figure 10: (a) plots the variation of C over the number of known symbols by Eve. (b) and (c) plot the variation of C over the number of overheard transmissions by Eve (20 known symbols per transmission). In particular, for (b) Eve's average signal SNR is 30 dB and the average condition number of H_{AE} is 5; for (c) Eve's average signal SNR is 20 dB and the average condition number of H_{AE} is 30. We only show the results for 4-antenna Alice.

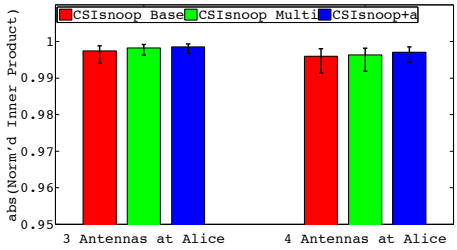

Figure 11: The median of C (with 25th/75th percentile) of the 3 cases of *CSIsnoop* Base, *CSIsnoop* Multi, and *CSIsnoop+a* (defined in Table 1). The average signal SNR at Eve is 30 dB and the average condition number of H_{AE} is 5.

Table 1: Definition of *CSIsnoop* Base, *CSIsnoop* Multi, and *CSIsnoop+a* in Fig. 11.

Case	Definition
CSIsnoop Base	The number of data streams from Alice to the Bobs equals the number of antennas at Alice.
CSIsnoop Multi	Alice has K antennas and there are K Bobs. The 1st packet is to Bob_1 and Bob_2. The 2nd packet is to Bob_1 and Bob_3 ($K = 3$), or Bob_1, Bob_3 and Bob_4 ($K = 4$). Eve computes H_{AB_1}.
CSIsnoop+a	Alice has K antennas but there are $K-1$ Bobs. Eve fakes Bob_K by using her first antenna. So Alice still sends K data streams.

It can be observed in Fig. 11 that the estimation accuracy of the 3 cases of Table 1 are very similar, with the difference among the median of C within 0.001. But *CSIsnoop+a* has a smaller variance than *CSIsnoop*. This is because when Eve joins in the downlink transmission, part of W_E become known to Eve. As a result, Eve can obtain a more accurate W_A and thereby H_{AB}. It should also be noted that there is a significant difference between the "Multi" case here and the discussion in Fig. 10(b-c): in Fig. 10(b-c), each overheard transmission contains k data streams, whereas in the "Multi" case each overheard transmission contains fewer than k data streams. Therefore, even if Eve overhears more than one transmission in the "Multi" case, she may not obtain more information compared to the base case. Thus estimation accuracy may not be improved.

6.6 Computation of H_{AE} with Encrypted Sounding Sequence

Finally, we examine how accurately Eve can estimate H_{AE} when Alice encrypts the channel sounding sequence [19]. Similar to the analysis of H_{AB}, it is equivalent that Eve estimates H_{AE} and $d \cdot H_{AE}$, where d is an unknown complex number. Therefore, we still use the absolute normalized correlation defined in Eq. (17) to evaluate the estimation accuracy of H_{AE}, which here is computed as ($H_{AE,comp}$ and $H_{AE,meas}$ are first expanded into 1-dimensional vectors)

$$C_{AE} = \frac{|H_{AE,comp} \cdot H^H_{AE,meas}|}{\|H_{AE,comp}\| \cdot \|H_{AE,meas}\|}. \quad (18)$$

We first fix the average condition number of H_{AE} to be 5 and plot the variation of C_{AE} over signal SNR in Fig. 12(a). It can be seen that under the same channel condition (in terms of H_{AE}), the computation of H_{AE} is more accurate than that of H_{AB} (by comparing with Fig. 9(a)). Moreover, it can also be observed that, while the accuracy of H_{AE} reduces when the SNR becomes smaller, the impact of noise strength is very small from 20 dB to 30 dB, and the median of C_{AE} is constantly above 0.996.

In comparison, the average condition number of H_{AE} has a larger impact on the estimation accuracy. Specifically, in Fig. 12(b), median of C_{AE} decreases to 0.96 when the condition number increases to 40. The 25th percentile also reduces to below 0.94. Nonetheless, compared to Fig. 9(b), the change of C_{AE} is still small. In other words, the estimation of H_{AE} is more robust to the channel condition between Alice and Eve than the estimation of H_{AB}. This is mainly because that the latter computation includes more steps. With the estimated H_{AE}, Eve can further use *CSIsnoop* to compute H_{AB}, even if Alice encrypts the channel sounding sequence.

7 CSI-BASED ATTACKS

In this section, we study how Eve can use the computed $H_{AB,comp}$ to attack the network, including computing the CSI-based password and selectively reducing the uplink throughput of a target Bob.

7.1 Computing CSI-Based Password

Because CSI decorrelates over half a wavelength in rich scattering environments, which is several centimeters in 2.4/5 GHz WiFi, schemes were proposed to generate a password between a transmitter and a receiver based on the CSI. It was assumed that such a

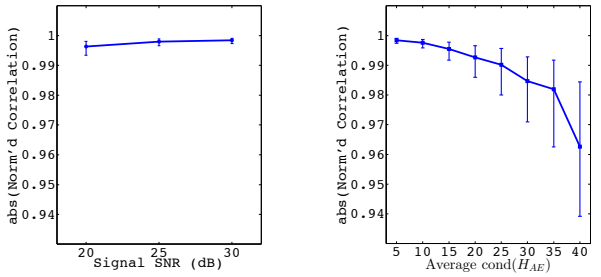

(a) (b)

Figure 12: The median (with 25th/75th percentile) of the absolute normalized correlation between $H_{AE,comp}$ and $H_{AE,meas}$ for 4-antenna Alice over (a) Eve's average signal SNR and (b) the average condition number of H_{AE}.

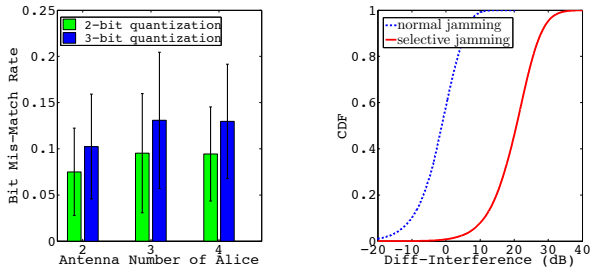

(a) (b)

Figure 13: (a) plots the average bit mis-match rate (with standard deviation) between the password computed by Eve and the password generated by Bob. (b) plots the CDF of interference difference between Bob_j and $Bob_{i\neq j}$ when Eve selectively jams Bob_j in the uplink.

password cannot be estimated by Eve even if Eve is close to Alice or Bob. However, by using *CSIsnoop*, Eve can compute CSI between Alice and Bob and use it to further compute the CSI-based password.

To evaluate how much Eve can estimate about the password, we consider the password generating scheme proposed in [12]: Alice and Bob quantify the relative amplitude of each sub-carrier of the downlink channel to generate password bits. Therefore, for q-bit quantization, the total bits of the password will be $k \cdot n \cdot q$, where k is the number of antennas at Alice and n is the number of sub-carriers that are used. Moreover, we assume that Eve's average SNR is 30 dB and the average condition number of H_{AE} is 5. After estimating the CSI from Alice to Bob, Eve uses the same method to compute the password.

Fig. 13(a) depicts experimental results for the bit mis-match rate between the password computed by Eve and the password generated by Bob. As discussed in Fig. 8, when Alice has more antennas, Eve's estimation accuracy of H_{AB} decreases. Consequently, the bit mis-match rate increases with the number of antennas at Alice, ranging from 7.5% to 9.5% for 2-bit quantization, and from 10.2% to 13.1% for 3-bit quantization.

However, Alice and Bob will also have bit mis-match between them. According to [12], in an indoor environment, the 2-bit scheme has bit mis-match rate between 3.5% and 5%, and the 3-bit scheme has bit mismatch rate between 5.6% and 7.9%. Suppose that the bit mis-match rate between Alice and Bob and between Eve and Bob is x% and y%, respectively. It can be calculated that Eve can estimate at least $(100-x-y)/(100-x)$ of the common bits between Alice and Bob. Therefore, for 2-bit and 3-bit quantization, *CSIsnoop* enables Eve to estimate over 90% and 85% of the password, respectively. For a practical system, there is still a step called "Information Reconciliation" for Alice and Bob to correct their bit mis-match by exchanging some packets over the air [12]. Eve may overhear these packets to further improve her computation of the password.

7.2 Selectively Decreasing Uplink Throughput

Concurrent *uplink* transmission of multiple data streams has been regarded as an important feature in the next generation wireless standard [4, 6, 18]. Receiver-based ZF-BF can be used to remove the inter-stream interference by projecting the desired signals onto the sub-space orthogonal to the interference. Nonetheless, this same property can be exploited by Eve to *selectively jam* the uplink transmission of Bob_j. In particular, once Eve knows the downlink CSI of Bob_j, she can send interference only in the sub-space of Bob_j's uplink signals. This is fundamentally different from current jamming techniques [14, 23], which treat all the concurrent data streams as the same.

In above analysis of this paper, we only consider the downlink CSI of Bob_j, which we denote as $H^d_{AB_j}$ in the following. *CSIsnoop* computes $H^d_{AB_j}$ from Alice's downlink multi-user beamforming transmission. However, the uplink CSI of Bob_j, $H^u_{AB_j}$, will be different from $H^d_{AB_j}$, mainly because of the difference between the receiving and the transmitting chain. Such difference has been shown to be stable over time and can be calibrated [10, 17]. Specifically, we denote $\alpha_{Alice,i}$ and $\alpha_{Eve,i}$ as the calibration coefficient between the receiving and the transmitting chain of Alice's and Eve's i^{th} antenna, respectively, and $Q_{Alice} = diag(\{\alpha_{Alice,1}, \ldots, \alpha_{Alice,k}\})$ and $Q_{Eve} = diag(\{\alpha_{Eve,1}, \ldots, \alpha_{Eve,k}\})$. For Bob_j, $\alpha_{Bob,j}$ has the same definition. Therefore, it can be computed that

$$H^u_{AB_j} = Q_{Alice} \left(H^d_{AB_j}\right)^T \alpha_{Bob,j}. \tag{19}$$

If Eve sends her jamming signals with beamforming weights

$$Q^{-1}_{Eve} \left(\left(H^d_{AE}\right)^T\right)^{-1} \left(H^d_{AB_j}\right)^T, \tag{20}$$

it can be calculated that at Alice, the jamming signals will be in the sub-space spanned by $(1/\alpha_{Bob,j})H^u_{AB_j}$, which is exactly the same signal sub-space of Bob_j's uplink transmission.

To evaluate this attack, we assume that both Alice and Eve have 4 antennas. We also suppose that Eve already computes her calibration coefficients. Fig. 13(b) shows the CDF of the interference difference between Bob_j and $Bob_{i\neq j}$ when selective jamming or normal jamming (for which Eve broadcasts her jamming signals) is used. It can be seen that with selective jamming, Eve is able to direct most of her jamming signal energy towards the uplink transmission of Bob_j. On average, there is a 20 dB difference between the interference at Bob_j and $Bob_{i\neq j}$, which indicates that this selective attack is highly effective.

8 RELATED WORK

CSI-based security schemes. Because CSI decorrelates over half a wavelength, it has been proposed as a security mechanism. One technique is artificial noise, which degrades the eavesdropper's channel by sending artificial noise in the null-space of the desired signals. The secrecy rate achieved by this scheme was analyzed in [5]. Experiments further demonstrated that the eavesdropper consistently has an SINR 15 dB smaller than the desired receiver [1]. Besides artificial noise, the transmitter and the receiver can also use CSI to generate a password between them. The generating framework as well as the secret key extraction speed were studied both theoretically [13] and experimentally [12]. What is more, the AP can also use CSI to authenticate the source of packets: if the packets are from the same client, they should have similar CSI within channel coherence time. Both averaged CSI magnitudes [9] and angle-of-arrival information [22] were proposed as a signature.

However, these schemes assume that the malicious node does not know the CSI between the transmitter and the receiver. Otherwise, they will be no longer safe. Indeed, artificial noise was demonstrated to be removable once CSI is known [16]. Likewise, link signatures based on CSI can be spoofed as long as the attacker has information about the uplink CSI of the client [3]. Furthermore, we demonstrate in this paper that once the attacker knows the CSI, a password generated out of it can also be predicted. Thus, because a malicious node with *CSIsnoop* is able to compute the CSI of any client by overhearing their downlink data transmission, the above designs need to be reconsidered for multi-user WLANs.

Encrypted CSI measurement. Even though CSI quickly decorrelates over distance, the explicit channel sounding process in current beamforming standards [7, 8] provides an opportunity for a malicious node to learn the CSI of clients by overhearing their CSI measurement feedback. One solution is to encrypt the feedback, but it results in additional overhead for the clients and the AP to establish the encryption and decryption key. In comparison, *CSIsec* uses a random sequence unknown to Eve instead of the pre-defined one to sound the channel by the AP [19].

In contrast, *CSIsnoop* uses the downlink data transmission instead of the channel sounding process to compute the CSI of clients. Moreover, while it was discussed in [19] that an encrypted channel sounding sequence can prevent the malicious node from estimating her channel (which is an important pre-filtering step in [16]), we show that the malicious node can still do so by employing her multiple antennas and the L-LTF field prepended to every packet [7, 8].

9 CONCLUSION

In this paper, we describe *CSIsnoop*, a framework by which the malicious node can infer CSI between the AP and clients by overhearing downlink multi-user transmission. We implement *CSIsnoop* in the software defined radio WARP v3 and show that the absolute normalized correlation between the computed CSI by *CSIsnoop* and the measured CSI by clients has an average of over 0.99. We also demonstrate that the malicious node can now use the computed CSI to break CSI-based security schemes, which urges reconsideration of the use of CSI as a shared secret in multi-user MIMO WLANs.

ACKNOWLEDGMENTS

The authors would like to thank Matthias Hollick and Matthias Schulz from Technische Universität Darmstadt, and João Luiz Rebelatto and Richard Demo Souza from Universidade Tecnológica Federal do Paraná for discussions. This research was supported by Cisco, Intel, the Keck Foundation, and by NSF grants CNS-1642929, CNS-1514285, and CNS-1444056.

REFERENCES

[1] N. Anand, S. Lee, and E. Knightly. Strobe: Actively Securing Wireless Communications Using Zero-Forcing Beamforming. In *Proceedings of IEEE INFOCOM 2012.*
[2] E. Aryafar, N. Anand, T. Salonidis, and E. Knightly. Design and Experimental Evaluation of Multi-User Beamforming in Wireless LANs. In *Proceedings of ACM MobiCom 2010.*
[3] S. Fang, Y. Liu, and P. Ning. 2016. Mimicry Attacks against Wireless Link Signature and New Defense Using Time-Synched Link Signature. *IEEE Transactions on Information Forensics and Security* 11, 7 (2016), 1515–1527.
[4] A. Flores, S. Quadri, and E. Knightly. A Scalable Multi-User Uplink for Wi-Fi. In *Proceedings of USENIX NSDI 2016.*
[5] S. Goel and R. Negi. 2008. Guaranteeing Secrecy Using Artificial Noise. *IEEE Transactions on Wireless Communications* 7, 6 (2008), 2180–2189.
[6] IEEE P802.11-Task Group ax. 2017. http://www.ieee802.org/11/Reports/tgax_update.htm. (2017).
[7] IEEE Std 802.11ac. 2013. IEEE Standard for Information Technology, 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.
[8] IEEE Std 802.11af. 2013. IEEE Standard for Information Technology, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, Amendment 5: Television White Spaces (TVWS) Operation.
[9] Z. Jiang, J. Zhao, X. Li, J. Han, and W. Xi. Rejecting the Attack: Source Authentication for Wi-Fi Management Frames using CSI Information. In *Proceedings of IEEE INFOCOM 2013.*
[10] F. Kaltenberger, H. Jiang, M. Guillaud, and R. Knopp. Relative Channel Reciprocity Calibration in MIMO/TDD Systems. In *Proceedings of ICT Future Network & Mobile Summit 2010.*
[11] K. Lin, S. Gollakota, and D. Katabi. Random Access Heterogeneous MIMO Networks. In *Proceedings of ACM SIGCOMM 2011.*
[12] H. Liu, Y. Wang, J. Yang, and Y. Chen. Fast and Practical Secret Key Extraction by Exploiting Channel Response. In *Proceedings of IEEE INFOCOM 2013.*
[13] Y. Liu, S. Draper, and A. Sayeed. 2012. Exploiting Channel Diversity in Secret Key Generation from Multipath Fading Randomness. *IEEE Transactions on Information Forensics and Security* 7, 5 (2012), 1484–1497.
[14] H. Rahbari, M. Krunz, and L. Lazos. Security Vulnerability and Countermeasures of Frequency Offset Correction in 802.11a Systems. In *Proceedings of IEEE INFOCOM 2014.*
[15] T. Schmidl and D. Cox. 1997. Robust Frequency and Timing Synchronization for OFDM. *IEEE Transactions on Communications* 45, 12 (1997), 1613–1621.
[16] M. Schulz, A. Loch, and M. Hollick. Practical Known-Plaintext Attacks against Physical Layer Security in Wireless MIMO Systems. In *Proceedings of NDSS 2014.*
[17] C. Shepard, H. Yu, N. Anand, E. Li, T. Marzetta, R. Yang, and L. Zhong. Argos: Practical Many-Antenna Base Stations. In *Proceedings of ACM MobiCom 2012.*
[18] K. Tan, H. Liu, J. Fang, W. Wang, J. Zhang, M. Chen, and G. Voelker. SAM: Enabling Practical Spatial Multiple Access in Wireless LAN. In *Proceedings of ACM MobiCom 2009.*
[19] Y. Tung, S. Han, D. Chen, and K. Shin. Vulnerability and Protection of Channel State Information in Multiuser MIMO Networks. In *Proceedings of ACM CCS 2014.*
[20] WARP. 2017. http://mangocomm.com/. (2017).
[21] WARPLab 7. 2017. https://warpproject.org/trac/wiki/WARPLab. (2017).
[22] J. Xiong and K. Jamieson. SecureArray: Improving WiFi Security with Fine-Grained Physical-Layer Information. In *Proceedings of ACM MobiCom 2013.*
[23] W. Xu, W. Trappe, Y. Zhang, and T. Wood. The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks. In *Proceedings of ACM MobiHoc 2005.*
[24] T. Yoo and A. 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.
[25] A. Zhou, T. Wei, X. Zhang, M. Liu, and Z. Li. Signpost: Scalable MU-MIMO Signaling with Zero CSI Feedback. In *Proceedings of ACM MobiHoc 2015.*