Massive MIMO Pilot Distortion Attack and Zero-Startup-Cost Detection: Analysis and Experiments

Xu Zhang and Edward W. Knightly
Department of Electrical and Computer Engineering, Rice University, Houston, Texas, USA
{xu.zhang, knightly}@rice.edu

***Abstract*—Accurate Channel State Information (CSI) is a key requirement for massive MIMO to achieve multi-fold increases in throughput and secrecy rate. Consequently, an adversary targeting the channel sounding process has the potential to significantly degrade performance. In this paper, we first present and model the *Pilot Distortion Attack*, a simple but devastating jamming strategy in which the adversary distorts the AP's CSI measurement of even a single client leading to denial-of-service for all clients associated with the AP. We propose *MACE* as a countermeasure that exploits the AP's large antenna array to detect jamming with zero startup cost and zero additional network overhead. Our key insight is that with many antennas, the AP's variance estimator of client Carrier Frequency Offset (CFO) will significantly increase when there are jamming signals present. We build a testbed with a 72-antenna massive MIMO AP and conduct the first experimental study of the Pilot Distortion Attack. Our results show that a single-antenna adversary jamming no more than $1/60$ of the time and having no more transmit power than any client can cause over 26% reduction of achievable rate of all clients. Moreover, by setting a single threshold, *MACE* can achieve 0.97 true positive at 0.01 false positive for various client/adversary locations and for a wide range of SNR ($5 \sim 35$ dB) and SIR ($-5 \sim 35$ dB) with SNR$-$SIR$\geq$5 dB.**

I. Introduction

Access Points (APs) employing massive MIMO provide new opportunities to scale both throughput and secrecy rate. However, similar to conventional multi-antenna networks such as IEEE 802.11ac, the gain of massive MIMO depends critically on whether the AP can accurately estimate the Channel State Information (CSI) of different clients [1], [2]. Current methods of CSI estimation in massive MIMO networks require clients to transmit pre-defined channel sounding pilots to the AP, which enables the AP to measure the uplink CSI from different clients. Downlink CSI can be obtained in the same way by using channel reciprocity. Previous work has shown that this channel sounding process is vulnerable to jamming attacks: If an adversary transmits jamming signals during both pilot transmission and the subsequent data transmission, network throughput will collapse even when the AP has unlimited antennas [3], [4]. The secrecy rate of clients also rapidly decreases when there is jamming during channel sounding [5]–[7].

In this paper, we analytically and experimentally study the impact and detection of jamming during channel sounding in practical massive MIMO networks. In particular, we first present and model the *Pilot Distortion Attack*, a simple but devastating jamming strategy that can lead to *denial-of-service of all clients associated with the AP*. Different from previous attacks in which the adversary is active during both channel sounding and data transmission, pilot distortion attacks only require the adversary to transmit jamming signals during channel sounding, while keeping silent afterwards. We study both *non-protocol-specific jamming* via Gaussian white noise spread over the entire channel as well as *protocol-specific jamming*, in which jamming signals have the same format as client channel sounding pilots. We show that in practical massive MIMO networks, the distorted CSI of even *a single* client can thwart concurrent uplink MMSE reception at the AP, thereby vastly degrading aggregate throughput.

As a counter mechanism, we propose MAssive MIMO Carrier frequency offset Estimate (*MACE*), a system that exploits variance scaling of Carrier Frequency Offset (CFO) measurements in massive MIMO to detect jamming with *zero startup cost* and *zero additional network overhead*. In other words, *MACE* can detect jamming for even the first packet received by the AP and is compatible with current WiFi and LTE standards. A key insight of *MACE* is that when there are no jamming signals, the CFO estimated by different antennas at the AP are very close to each other, because all estimates share the same true value and are also based on signals in the same carriers. Thus, we develop a model of the variance of CFO estimates and show that without jamming, the normalized variance is independent of the wireless channel, the signal SNR, and the CFO between the AP and the client. In comparison, when there are jamming signals, we show that even if they are sent in exactly the same format as the channel sounding pilots, the normalized variance estimator significantly increases. As this difference increases with the size of the massive MIMO array, *MACE* can detect jamming with zero startup cost, i.e., without *a priori* statistical training. This further enables *MACE* to support highly mobile clients, and prevents the adversary from escaping detection by affecting statistical training. Moreover, because repeated symbols already exist in various wireless standards for CFO estimation, *MACE* does not introduce any additional network overhead. *MACE* also does not require any shared secrets. Consequently, after detection via *MACE*, the AP can use different scheduling and beamforming algorithms to minimize the impact of distorted CSI (e.g., exclude the distorted clients for concurrent uplink transmission).

Furthermore, to prevent the adversary who is aware of the *MACE* mechanism and may foil the detection by imitating the client's CFO when transmitting protocol-specific jamming signals [8], we propose client-side *Per-Frame Random CFO Injection*. In particular, before sending the channel sounding pilots, each client will inject a random CFO in the digital

978-1-5386-0683-4/17/$31.00

domain. The range of this random CFO is computed by the client, such that it does not lead to decoding error at the AP. Moreover, by changing the random CFO per transmission, the adversary cannot estimate its value.

Finally, we build a massive MIMO testbed to evaluate the impact of pilot distortion attacks and the detection performance of *MACE*. We are the first to experimentally study massive MIMO from a security point of view. In particular, we use WARP v3 [9] and the Argos massive MIMO AP [10], [11] that has a 72-antenna array, and collect over 3,000,000 packet measurements in the 5 GHz WiFi band. Our main experimental results can be summarized as follows:

(1) For the pilot distortion attack, a single adversary jamming no more than $1/60$ of the overall airtime and having no more transmit power than any client can lead to 38% to 26% reduction of achievable rate when 4 to 9 clients are grouped for concurrent uplink transmission. In practice, the damage will be even more severe, as limiting throughput reduction to 38% and 26% requires the clients to perfectly adapt their Modulation and Coding Scheme (MCS) to the maximum achievable rate given the attack properties. Otherwise, the attack can degrade throughput to zero due to unrecoverable decoding errors.

(2) Because the variance of the normalized CFO estimates is independent of the wireless channel and the signal SNR, by setting a single detection threshold, *MACE* can achieve 0.97 true positive at 0.01 false positive for various client/adversary locations, and for a wide range of SNR ($5 \sim 35$ dB) and SIR ($-5 \sim 35$ dB) with SNR$-$SIR$\geq$5 dB.

(3) Even with only 16 antennas at the AP and 32 repeated symbols, *MACE* can achieve 0.97 true positive at 0.03 false positive with the same client/adversary locations and SNR/SIR range; consequently, *MACE* can also be used for general-purpose jamming detection, even with a moderate number of antennas and repeated symbols (e.g., cyclic prefix of OFDM symbol).

The rest of the paper is organized as follows. Sec. II describes our threat model. We analyze pilot distortion attacks in Sec. III and present our design of *MACE* in Sec. IV. Experimental evaluations are studied in Sec. V. Sec. VI discusses related work and Sec. VII concludes the paper.

II. Threat Model

As shown in Fig. 1, we consider a threat model with a WLAN setup ,which includes a massive MIMO AP (Alice) that has M antennas, and K single-antenna clients (Bobs). OFDM transmission is employed along with channel sounding with time division to measure CSI between Alice and the K different Bobs. That is, pre-defined channel sounding pilots are transmitted from different Bobs to Alice in orthogonal time slots (sending pilots from Alice to Bobs and feeding back the CSI measurements is infeasible in massive MIMO [11]). However, because there are no standards defining the channel sounding pilots originated from clients, we use the signal format of IEEE 802.11ac, where two identical Long Training Sequences (LTS) are concatenated and broadcasted by Alice for downlink CSI measurement. After Alice receives Bobs' LTS and estimates Bobs' CSI, linear beamforming algorithms like ZF/MMSE are used for concurrent uplink/downlink transmissions. Recent developments have shown that ZF/MMSE can be implemented for massive MIMO [12] and lead to higher throughput than conjugate beamforming [1], [11].

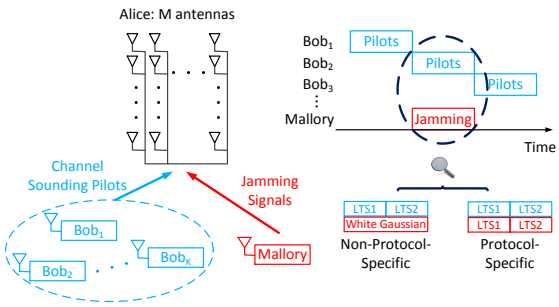

Fig. 1. Threat model: A single-antenna malicious node Mallory distorts CSI measurement of a legitimate client Bob by jamming his channel sounding.

We further consider that there is a single-antenna malicious node Mallory in range of Alice. Mallory is a reactive jammer and can transmit jamming signals during channel sounding (the timing of channel sounding can be estimated by overhearing network control signals). In particular, we consider the following two types of jamming signals: *(1) Non-Protocol-Specific Jamming*: Mallory knows the carriers in which the channel sounding pilots are transmitted, but is unaware of the detailed protocol used by Alice and Bobs. In this case, Mallory transmits white Gaussian noise in the carriers. *(2) Protocol-Specific Jamming*: Mallory knows that each Bob transmits two repeated LTS for CFO/CSI measurement, and is also able to strictly time-synchronize with Bob [13]. Therefore, Mallory can also send repeated jamming signals (in this paper we consider the same repeated LTS as Bob) to distort Bob's CSI measurement at Alice.

III. Pilot Distortion Attacks

A multi-antenna AP can realize concurrent uplink and downlink transmissions to multiple clients. However, with a *Pilot Distortion Attack*, an adversary transmits jamming signals during channel sounding, targeting that the distorted CSI measurement at the AP, will result in large reduction of network throughput. Not only is such an attack difficult to detect due to its small energy and time footprint, it is also powerful because distorting the CSI of a single Bob can lead to denial-of-service for all Bobs associated with the AP.

Distorted CSI can have different influences on uplink and downlink due to properties of beamforming algorithms. Consider ZF beamforming: Denote the channel between an M-antenna Alice and K single-antenna Bobs to be an $M \times K$ matrix H. Thus the beamforming weights are computed by $W = (H^*H)^{-1}H^*$ (H^* is the conjugate transpose of H). In the uplink, inter-client interference is removed by Alice computing $W \cdot H$, while in the downlink, interference is removed by Alice computing $H^T \cdot W^T$ (H^T is the transpose of H). As a result, if Mallory distorts the CSI of Bob_i, which is the i^{th} column of H, only Bob_i receives extra interference in the downlink, while all clients but Bob_i receive extra interference in the uplink. In other words, by distorting the CSI of a single Bob, all concurrent uplink transmission can be thwarted. This also reduces downlink throughput for closed-loop traffic (e.g., TCP) [14].

To further quantify the reduction of uplink throughput when Bob_i's CSI is distorted, we denote the channel from Bob_i and Mallory to Alice to be $H_{Bi} \sim CN(0,1)$ and

$H_{Mal} \sim CN(0,1)$, respectively. During channel sounding, Bob_i's sounding pilot is $X_{Bi,p}$ while Mallory's jamming signal is $X_{Mal,p}$ ($|X_{Bi,p}| = |X_{Mal,p}| = 1$). What Alice receives can thus be written as

$$Y_{i,p} = \sqrt{P_{Bi,p}} H_{Bi} X_{Bi,p} + \sqrt{P_{Mal,p}} H_{Mal} X_{Mal,p} + Z, \quad (1)$$

where $P_{Bi,p}$ and $P_{Mal,p}$ are the signal strength of Bob_i and Mallory at Alice during channel sounding, respectively. Z is random noise with strength N. Here we assume that Bob_i only transmits the channel sounding pilot once without loss of generality. When Bob_i transmits repeated channel sounding pilots, $P_{Mal,p}$ will become the effective jamming strength and have different values for protocol-specific and non-protocol-specific jamming. N will also become the effective noise strength. The MMSE estimate of H_{Bi} given $Y_{i,p}$ is

$$\hat{H}_{Bi} = E\{H_{Bi} X_{Bi,p} Y_{i,p}^*\} E\{Y_{i,p} Y_{i,p}^*\}^{-1} Y_{i,p}, \quad (2)$$

with error $\epsilon_{Bi} = H_{Bi} - \hat{H}_{Bi}$ being Gaussian with variance $\sigma^2_{\epsilon_{Bi}} I$ and

$$\sigma^2_{\epsilon_{Bi}} = \frac{P_{Mal,p} + N}{P_{Bi,p} + P_{Mal,p} + N}. \quad (3)$$

During concurrent uplink transmission, we denote W_j to be the beamforming weights of Bob_j ($j \neq i$). Mallory keeps silent during data transmission. Therefore, after receive beamforming, Alice obtains

$$\begin{aligned} Y_{j,d} = \quad & W_j \sqrt{P_{Bj,d}} H_{Bj} X_{Bj,d} \\ & + W_j \textstyle\sum_{k \neq i,j} \sqrt{P_{Bk,d}} H_{Bk} X_{Bk,d} \\ & + W_j \sqrt{P_{Bi,d}} \left(\hat{H}_{Bi} - \epsilon_{Bi} \right) X_{Bi,d} + W_j Z, \end{aligned} \quad (4)$$

where $P_{Bk,d}$ is the signal strength of Bob_k at Alice during data transmission, and $|X_{Bk,d}| = 1, \ \forall k$. It can be observed in Eq. (4) that, due to Bob_i's distorted CSI, extra interference to Bob_j can be computed as $W_j \sqrt{P_{Bi,d}} \epsilon_{Bi} X_{Bi,d}$. For MMSE estimate, ϵ_{Bi} is independent of $\hat{H}_{Bi}$ and thereby the computed beamforming weights W_j. Therefore, the expected strength of the extra interference with normalized W_j is

$$E\{|W_j \sqrt{P_{Bi,d}} \epsilon_{Bi} X_{Bi,d}|^2\} = \frac{(P_{Mal,p} + N) P_{Bi,d}}{P_{Bi,p} + P_{Mal,p} + N}. \quad (5)$$

Two observations can be obtained from Eq. (5). First, the extra interference does not decrease when Alice has an increasing number of antennas. However, because of the beamforming gain, when Alice has more antennas, Bob_j's signal strength after receive beamforming increases. This makes the impact of the extra interference diminish when Alice's antenna number tends to infinity. Nonetheless, for practical massive MIMO networks, Alice's antenna number is limited. It is shown in Sec. V-B that even if Mallory has no more transmit power than any Bob, pilot distortion attack can still lead to 38% to 26% reduction of per-client achievable rate for concurrent uplink transmission of 4 to 9 Bobs.

Second, if the noise strength N is ignored in Eq. (5), we can further compute that the pilot distortion attack is Δ times more efficient than attacks with the same strength $P_{Mal,p}$ but directly jamming the data transmission, where

$$\Delta = \beta \cdot \frac{P_{Bi,d}}{P_{Bi,p} + P_{Mal,p}}. \quad (6)$$

Here β is the ratio of duration of data transmission over channel sounding. For 20 MHz bandwidth and 2 LTS as channel sounding pilots, each Bob's channel sounding takes $8\mu s$ (including cyclic prefix of the LTS). In comparison, data transmission can be extended within channel coherence time that ranges from $500\mu s$ to more than $1ms$ [11]. This leads to a β no smaller than 60. Consequently, if Mallory has similar power to Bob_i, pilot distortion attack will be over 30 times more efficient than directly jamming the data transmission. In other words, the pilot distortion attack has high efficiency with small energy and time footprint.

IV. Jamming Detection with *MACE*

In this section, we present *MACE*, a system that can detect jamming with *zero startup cost* and *zero additional network overhead*. We introduce the background of CFO estimation, present the architecture of *MACE*, and analyze the variance of CFO estimates without and with jamming signals, respectively. We further study the countermeasure of per-frame random CFO injection.

A. CFO Estimation with a Single Receiving Antenna

CFO commonly exists due to hardware discrepancies between the transmitter and the receiver, and it needs to be estimated and corrected in the early stage of the decoding chain. In current wireless networks, CFO is estimated through repeated training sequences. If we denote $Y = \{Y_1 | Y_2\}$ to be the signals at the receiver (Y_1 and Y_2 are the first and the second half of Y, respectively), we can obtain

$$\begin{aligned} Y_1 &= R + Z_1, \\ Y_2 &= R e^{j\theta} + Z_2, \end{aligned} \quad (7)$$

where R is the received copy of the training sequence, Z_1 and Z_2 are random noise with strength N, and $\theta = 2\pi f t \cdot len(R)$ is the phase rotation due to CFO f and sampling interval t. We define $len(\cdot)$ as the function that returns the length a vector.

The Maximum Likelihood (ML) estimate of θ given Y_1 and Y_2 was derived by Moose in [15], which computes

$$\hat{\theta} = \arg(Y_2 Y_1^*). \quad (8)$$

It was also computed that in high SNR regime,

$$\begin{aligned} E\{\hat{\theta} | \theta, R\} &= \theta, \\ Var\{\hat{\theta} | \theta, R\} &= N/(RR^*). \end{aligned} \quad (9)$$

B. System Architecture of MACE

The architecture of *MACE* is illustrated in Fig. 2. *MACE* employs the CFO estimates of Alice's many antennas to detect jamming signals, because the existence of jamming signals will rapidly increase the variance of CFO estimates, thus enabling detection (since *MACE* targets jamming detection, this is not the optimal CFO estimation for packet decoding). Since CFO estimation is supported by various wireless standards, *MACE* does not introduce any additional network overhead.

As a stand alone module at Alice, there are four steps of computation after *MACE* receives the raw signals from each Bob and before it determines whether jamming signals are present. The four steps are summarized as follows:

(1) SNR Estimation. *MACE* first measures the SNR of each antenna. Particularly, the noise strength is measured when there are no incoming signals.

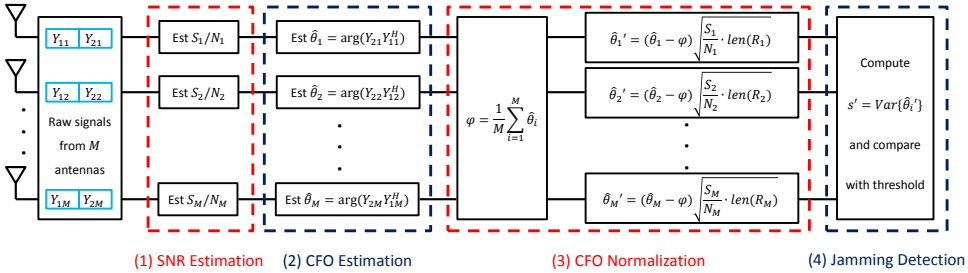

Fig. 2. System architecture of *MACE*: the variance of CFO estimates at Alice greatly increases with jamming signals, which is used by *MACE* for detection.

(2) CFO Estimation. Subsequently, the repeated symbols received by each antenna are used to compute a CFO estimate. We employ the ML estimator discussed in Sec. IV-A.

(3) CFO Normalization. *MACE* then computes the average of these M CFO estimates, and normalizes each CFO estimate by subtracting the average and scaling with the corresponding SNR. Without jamming signals, each normalized CFO estimate can be approximated by a standard Gaussian random variable. The details are discussed in Sec. IV-C.

(4) Jamming Detection. Finally, *MACE* computes the variance of these normalized CFO estimates, which is close to 1 without jamming, but much larger than 1 with jamming. Therefore, a threshold can be set for jamming detection. The details are discussed in Sec. IV-C and Sec. IV-D.

C. Variance of CFO Estimates without Jamming

Because the multiple CFO estimates at Alice share the same true value and are also based on signals in the same carriers, when there is no jamming, the variance of these CFO estimates should be small (in high SNR regime). In the following, we derive an analytical form of this variance.

When Alice has M antennas, we denote the multiple CFO estimates to be $\{\hat{\theta}_1, \hat{\theta}_2, \ldots, \hat{\theta}_M\}$. Therefore, we can compute the average as $\phi = \frac{1}{M}\sum_{i=1}^{M}\hat{\theta}_i$, and the variance as $s = \frac{1}{M}\sum_{i=1}^{M}(\hat{\theta}_i - \phi)^2$. However, it can be observed that, without the knowledge of the distribution of each $\hat{\theta}_i$, the statistics of s can hardly be computed. Therefore, we further make the following 2 assumptions about the CFO estimates at Alice: *(i) Normal Distribution.* Given θ and R_i (which is the the training sequence received by Alice's i^{th} antenna), $\hat{\theta}_i$ is a Gaussian random variable with average θ and variance $N_i/(R_i R_i^*)$. *(ii) Uncorrelated Noises.* We assume that the random noises are uncorrelated among Alice's different antennas. Therefore, $\hat{\theta}_i$ is uncorrelated with $\hat{\theta}_j$ if $i \neq j$.

With the assumption of uncorrelated noises, we can first compute the statistics of the average ϕ, which are

$$\begin{aligned} &E\{\phi|\theta, R_1, \ldots, R_M\} = \tfrac{1}{M}\textstyle\sum_{i=1}^{M} E\{\hat{\theta}_i|\theta, R_i\} = \theta, \\ &Var\{\phi|\theta, R_1, \ldots, R_M\} = \tfrac{1}{M^2}\textstyle\sum_{i=1}^{M} Var\{\hat{\theta}_i|\theta, R_i\}. \end{aligned} \tag{10}$$

Therefore, ϕ is a conditionally unbiased estimate of θ, which also has a small conditional variance (due to the $1/M^2$ factor) when Alice has many antennas. Consequently, we can use ϕ to approximate the true CFO θ. This allows us, together with the assumption of normal distribution, to normalize each CFO estimate $\hat{\theta}_i$ into a standard Gaussian random variable $\hat{\theta}_i'$ by

$$\hat{\theta}_i' = \frac{\hat{\theta}_i - \phi}{\sqrt{Var\{\hat{\theta}_i|\theta, R_i\}}} = (\hat{\theta}_i - \phi)\sqrt{\frac{S_i}{N_i}\cdot len(R_i)}, \tag{11}$$

where S_i and N_i are the signal and noise strength measured by Alice's i^{th} antenna, respectively. Moreover, it is known that the summation of the square of M standard Gaussian random variables is subjected to chi-squared distribution with M degrees of freedom. Therefore, if we denote s' to be the variance of these M normalized CFO estimates, we can compute that

$$\begin{aligned} &E\{s'|\theta, R_1, \ldots, R_M\} = 1, \\ &Var\{s'|\theta, R_1, \ldots, R_M\} = 2/M. \end{aligned} \tag{12}$$

It can be observed in Eq. (12) that, when Alice has more antennas, s' becomes increasingly concentrated around 1. This makes it possible to set a threshold to separate those channel sounding pilots without jamming signals. Furthermore, the conditional statistics of s' is independent of the wireless channel, the signal SNR (as long as in high SNR regime), and the CFO between Alice and Bob. This is the main reason why *MACE* can detect jamming with zero startup cost.

D. Variance of CFO Estimates with Jamming

CFO estimate at each Alice's antenna becomes less accurate in the presence of jamming signals. As a result, the variance of CFO estimates increases, which makes s' larger than 1. In the following, we characterize s' for both non-protocol-specific and protocol-specific jamming.

Non-Protocol-Specific Jamming. Since Mallory transmits white Gaussian noise during channel sounding, if the signal SINR at Alice is not very small, according to Eq. (9), $\hat{\theta}_i$ should have conditional variance $(N_i + J_i)/(S_i \cdot len(R_i))$, where J_i is the jamming signal strength at Alice's i^{th} antenna. As a result, the correct normalization of $\hat{\theta}_i$ should be

$$\hat{\theta}_i' = (\hat{\theta}_i - \phi)\sqrt{\frac{S_i}{N_i + J_i}\cdot len(R_i)}. \tag{13}$$

However, Alice does not know the existence of jamming signals, and thereby treats $S_i + J_i$ as the legitimate signal strength. If we assume that the average ϕ keeps unchanged, Alice will now mistakenly compute

$$\hat{\theta}_i'^{(err)} = (\hat{\theta}_i - \phi)\sqrt{\frac{S_i + J_i}{N_i}\cdot len(R_i)}. \tag{14}$$

Therefore, as long as $J_i > N_i$, we can obtain

$$\frac{\hat{\theta}_i'^{(err)}}{\hat{\theta}_i'} = \sqrt{1 + \frac{S_i J_i + J_i J_i + J_i N_i}{S_i N_i}} > 1. \tag{15}$$

Consequently, the variance of $\hat{\theta}_i'^{(err)}$ also increases.

Protocol-Specific Jamming. When Mallory uses protocol-specific jamming, the jamming signals received by Alice's i^{th} antenna can be written as $\{Q_i|Q_i e^{j\eta}\}$, where η is the phase

rotation due to CFO between Alice and Mallory. Similarly, because Alice is not aware of the jamming signals, she uses Eq. (8) to compute a mistaken CFO estimate $\hat{\theta}_i^{(err)}$, which has conditional average

$$\begin{aligned} & E\{\hat{\theta}_i^{(err)}|\theta, R_i, \eta, Q_i\} \\ = \; & \arg\left((R_i e^{j\theta} + Q_i e^{j\eta})(R_i + Q_i)^*\right) \\ = \; & \arg\left(e^{j\theta}(|R_i|^2 + R_i Q_i^*) + e^{j\eta}(|Q_i|^2 + Q_i R_i^*)\right). \end{aligned} \tag{16}$$

While θ and η are fixed for all of Alice's M antennas, R_i and Q_i will be different. However, because Alice is a massive MIMO AP, Mallory can hardly estimate or control the channel between Alice and herself or the channel between Alice and Bob, and thereby the values of R_i and Q_i. Consequently, the only parameter in Eq. (16) that Mallory can control is η. And as long as $\theta \neq \eta$, the conditional average of $\hat{\theta}_i^{(err)}$ will no longer be the same for Alice's different antennas.

In addition, we can also compute the conditional variance of $\hat{\theta}_i^{(err)}$ when there are protocol-specific jamming signals, which is

$$Var\{\hat{\theta}_i^{(err)}|\theta, R_i, \eta, Q_i\} = \frac{N_i}{(R_i + Q_i)(R_i^* + Q_i^*)}. \tag{17}$$

Combining Eq. (16) and Eq. (17), it can be observed that when Alice normalizes $\hat{\theta}_i^{(err)}$ to $\hat{\theta}_i'^{(err)}$ by using Eq. (11), each $\hat{\theta}_i'^{(err)}$ will have unit variance but non-zero average. This again makes s' larger than 1.

E. Per-Frame Random CFO Injection by Bobs

As discussed in Sec. IV-D, for protocol-specific jamming, as long as $\theta \neq \eta$, the value of s' will be larger than 1. Thus jamming can be detected by *MACE*. In contrast, if $\theta \approx \eta$, s' gets close to 1, which makes the jamming signals hard to be detected. However, it was shown in previous work that it is possible for Mallory to set η close to θ [8] (which may then foil the *MACE* detection). In particular, oscillator frequency remains stable within short durations. By overhearing Bob's packets, Mallory can estimate the CFO between Bob and herself, and thereby compensate for such CFO in the digital domain before sending the jamming signals.

To address this problem, we further propose a counter-measure called *Per-Frame Random CFO Injection*, with which each Bob injects a random CFO in the digital domain before sending his channel sounding pilots. Such random CFO cannot be predicted and thereby imitated by Mallory. Mallory also cannot estimate its value before completely receiving the 2 LTS, because Bob can actually inject the random CFO only to the LTS but not the prepended short training sequences. In the meantime, this random CFO should not lead to decoding error at Alice (exceeds Alice's correcting range, which is defined in standards like IEEE 802.11ac) when there are no jamming signals, which further makes it fully compatible with current WiFi and LTE standards.

The detailed process of per-frame random CFO injection is illustrated in Fig. 3, where f_{Alice}, f_{Bob}, and f_{Mal} are the actual oscillator frequencies of Alice, Bob, and Mallory, respectively. $f(\delta)$ denotes the frequency offset that causes δ phase rotation. First, when Bob overhears packets from Alice, he can estimate $f_{Alice} - f_{Bob}$. Since Bob knows that Alice can correct CFO within $[f(-\pi), f(\pi)]$, he can then compute

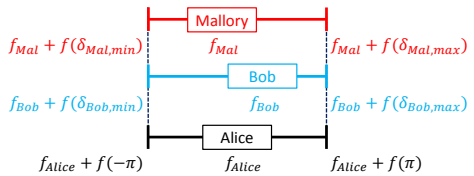

Fig. 3. Per-frame random CFO injection, with which each Bob injects a random CFO in the digital domain before sending his channel sounding pilots.

a range $[f(\delta_{Bob,min}), f(\delta_{Bob,max})]$ in which the additional random CFO will not lead to decoding error at Alice.

Similar to Bob, Mallory can also estimate $f_{Alice} - f_{Mal}$ and thereby compute $[f(\delta_{Mal,min}), f(\delta_{Mal,max})]$. As a result, at Alice both θ of Bob and η of Mallory are between $-\pi$ and π. If Bob uniformly chooses his additional random CFO within $[f(\delta_{Bob,min}), f(\delta_{Bob,max})]$, the best strategy for Mallory is to also uniformly select an additional CFO within $[f(\delta_{Mal,min}), f(\delta_{Mal,max})]$, or to just fix her CFO. In this case, If *MACE* cannot detect protocol-specific jamming signals when $|\theta - \eta| < \omega$, we can computed that

$$P(|\theta - \eta| < \omega) = \frac{\omega}{\pi}. \tag{18}$$

As evaluated in Sec. V-D, ω has a small value in practice.

V. Experimental Evaluation

In this section, we build a testbed and use experiments to evaluate the impact of pilot distortion attacks and to study the detection performance of *MACE* for practical massive MIMO.

A. Experimental Setup

We build a testbed for experimental evaluation by using the WARP v3 [9] and the Argos massive MIMO AP [10], [11], and use the topology shown in Fig. 4(a). It emulates a network with one massive MIMO AP and multiple clients, and a malicious node jams the channel sounding process to reduce the network throughput. In particular, the Argos massive MIMO AP has a 72-antenna array spaced by 6.35 cm (Fig. 4(b)). During each experiment, a single Bob and a single Mallory are selected to transmit signals to Alice, which emulates the channel sounding with time division and with/without jamming signals. Moreover, to emulate different CFO between Bob and Mallory, we add additional CFO to the signals in the digital domain before each transmission. This is because the inherent CFO between Bob and Mallory due to hardware discrepancies is relatively stable over time. We also change the transmit power of Bob and Mallory to explore various combinations of SNR and SIR. All experiments are conducted in the 5 GHz WiFi band with 20 MHz bandwidth. In total, we collect measurements for over 3,000,000 packets.

The detailed format of each transmission from Bob/Mallory to Alice is shown in Fig. 4(b). The first part contains only LTS (defined in IEEE 802.11ac) from the selected Bob, which are used to estimate Bob's CSI/CFO (to Alice) and to compute *MACE*'s output without jamming. In comparison, the second part contains only jamming signals from the selected Mallory: for non-protocol-specific jamming, they are white Gaussian noise within the 20 MHz channel, while for protocol-specific jamming, they are the same LTS that are transmitted by Bob. We use the second part to measure the jamming signal strength and Mallory's CFO (to

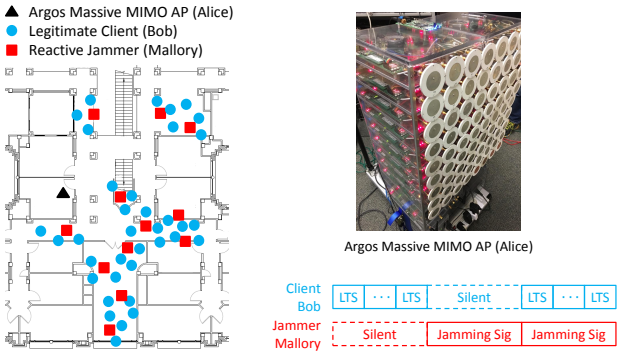

(a) (b)

Fig. 4. (a) Experimental setup with the location of the massive MIMO AP Alice, and part of the locations of the legitimate clients Bobs and the adversary Mallory. (b) The Argos massive MIMO AP (Alice) and the format of signals from Bob/Mallory to Alice.

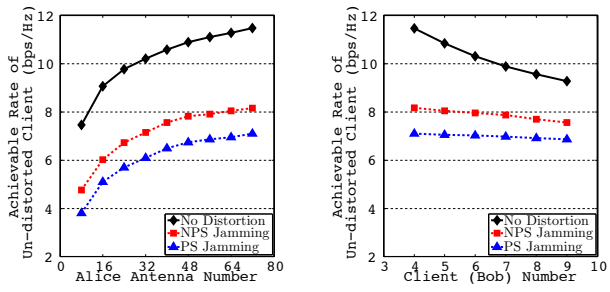

(a) (b)

Fig. 5. Average per-client uplink achievable rate of un-distorted Bobs (a) when Alice has different number of antennas (with 4 Bobs) and (b) when there are different number of concurrent transmitting Bobs (with 72 antennas at Alice). All Bobs' SNR before receive beamforming are around 20 dB ($18 \sim 22$ dB). And a single Bob's CSI is distorted by around 0 dB SIR ($-2 \sim 2$ dB) jamming signals. In the figures, NPS and PS stand for Non-Protocol-Specific and Protocol-Specific jamming, respectively.

Alice). Finally, the third part contains signals from both Bob and Mallory, which are used to measure Bob's distorted CSI and *MACE*'s output with jamming.

In addition, while the Argos massive MIMO AP has 72 antennas, we also study the impact of pilot distortion attacks and the detection performance of *MACE* when Alice has fewer antennas. Particularly, we randomly select M antennas out of the 72 if $M < 72$. For every M, this process is repeated several times to obtain the average results.

B. Achievable Rate Reduction due to Pilot Distortion Attacks

To study the impact of the Pilot Distortion Attacks, we use the Shannon equation $log_2(1 + SINR)$ to compute the achievable rate of Bobs' concurrent uplink transmissions, and compare their values without and with jamming signals. The results with Alice having different number of antennas and using MMSE receive beamforming are shown in Fig. 5(a).

It can be observed that, even if only a single Bob's CSI is distorted, the achievable rate significantly decreases, ranging from 49% to 38% reduction for protocol-specific jamming, and from 36% to 29% for non-protocol-specific jamming, when Alice's antenna number increases from 8 to 72. The main reason that non-protocol-specific jamming leads to a smaller reduction is because its effective jamming strength decreases when repeated LTS are used for CSI measurement. In particular, the resulting average increase of inter-client interference is measured to be 16.2 dB and 13.7 dB for protocol-specific and non-protocol-specific jamming, respectively. In comparison, using experimental data for Eq. (5), we can compute the increase to be 15.3 dB and 12.5 dB, respectively.

Fig. 5(b) further displays the achievable rate when Alice has 72 antennas but the number of Bobs increases from 4 to 9. Because only a single Bob's CSI is distorted, the achievable rates under the pilot distortion attacks do not change much, while the achievable rates without the attack decrease due to increasing inter-client interference. Nonetheless, when there are 9 Bobs transmitting concurrently, we can still observe that the pilot distortion attack with protocol-specific jamming leads to 26% decrease of achievable rate. In practice, the damage will be even more severe, as limiting throughput reduction to 26% requires the clients to perfectly adapt their MCS to the maximum achievable rate given the attack properties. Otherwise, the attack can degrade throughput to zero due to unrecoverable decoding errors. Meanwhile, Mallory can also distort multiple CSI to further reduce the clients' achievable rate. And as more clients tend to be included in concurrent transmissions in massive MIMO networks, the network-wide impact of pilot distortion attack actually increases when Alice has more antennas.

Therefore, for the pilot distortion attacks, a single adversary jamming no more than $1/60$ of the time ($8\mu s$ over $> 500\mu s$ as discussed in Sec. III) and having no more transmit power than any client can lead to 38% to 26% reduction of achievable rate when 4 to 9 clients are grouped for concurrent uplink transmission.

C. Variance of Normalized CFO Estimates without Jamming

To evaluate the detection performance of *MACE*, in the following, we first discuss the CDF of the variance of normalized CFO estimates without jamming signals. Particularly, we compare 2 methods for noise strength estimation at Alice:

(1) Non-Signal-Aided. Alice measures noise strength when there are no incoming signals. This method only allows Alice to measure the noise strength generated by the receiver.

(2) Signal-Aided. Alice knows that for her i^{th} antenna, the incoming signals have a structure of $\{Y_{1i}|Y_{2i}\}$, where $Y_{1i} = R_i + W_{1i}$ and $Y_{2i} = R_i e^{j\theta} + W_{2i}$. Therefore, Alice can first estimate $\hat{\theta}$ and then compute the noise strength as $E\{|Y_{2i}e^{-j\hat{\theta}} - Y_{1i}|^2\}/2$. This method requires an accurate estimation of θ, yet it does not include the noise correlated with signal R_i.

It can be first observed in Fig. 6(a) that, for the Non-Signal-Aided method, the experimental results with high/low SNR deviate from the theoretical calculation. The main reasons are that: when SNR is high, noise strength introduced by the transmitter begins to surpass that generated by the receiver, which results in large normalization error in Eq. (11); when SNR is low, the error in Eq. (9) becomes large. In comparison, when SNR is within $5 \sim 25$ dB, the experimental results are close to the theoretical calculation. The main reason for the long tail is that the SNR of Alice's different antennas vary significantly: in experiments, the average range of SNR difference is 22 dB.

In contrast, as can be observed in Fig. 6(b), the difference

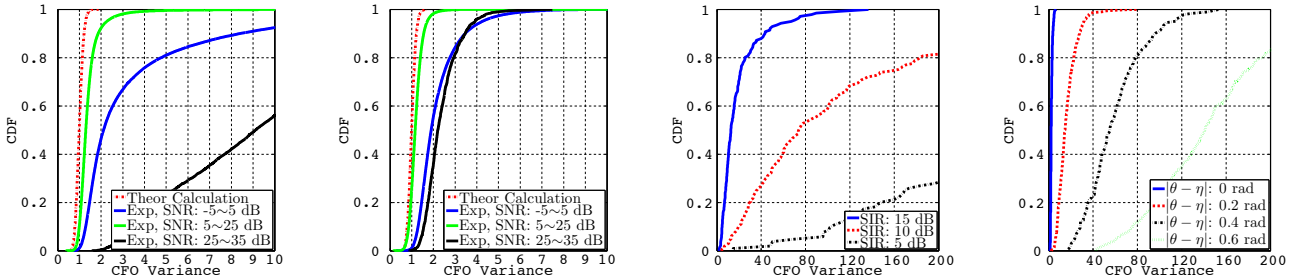

(a) (b) (c) (d)

Fig. 6. When there are no jamming signals, (a) and (b) display the variance of normalized CFO estimates with noise strength measured by the (a) Non-Signal-Aided and (b) Signal-Aided method, respectively. When there are jamming signals, (c) and (d) display the variance of normalized CFO estimates (with noise strength measured by the Non-Signal-Aided method) when there are (c) non-protocol-specific jamming ($\sim$ 20 dB SNR) and (d) protocol-specific jamming ($\sim$ 20 dB SNR, $\sim$ 10 dB SIR, 0.1 rad bin size), respectively. Alice has 72 antennas.

between the experimental results and the theoretical calculation decrease when the Signal-Aided method is used to measure the noise strength. In particular, when the transmitter side noise is included, the experimental results at high SNR become much closer to the theoretical calculation. However, the Signal-Aided method cannot be employed by *MACE*, because it will mistakenly include the white Gaussian jamming signals when computing the noise strength. Therefore, all of the following figures are based on the the Non-Signal-Aided method.

D. Variance of Normalized CFO Estimates with Jamming

When there are non-protocol-specific white Gaussian jamming signals, the variance of the normalized CFO estimates significantly increases. In Fig. 6(c), the x-axis now extends to 200 instead of 10 as in Fig. 6(a) and 6(b). It can be also seen that, when the jamming signals become stronger, the ratio of $\hat{\theta}_i'^{(err)}/\hat{\theta}_i'$ computed in Eq. (15) increases, and therefore the variance of the normalized CFO estimates also increases.

For protocol-specific jamming, we observe a similar trend that the variance of the normalized CFO estimates increases with stronger jamming signals. Due to space limitation, the corresponding results are not shown. Instead, in Fig. 6(d), we display the variance of the normalized CFO estimates when the CFO between Mallory and Bob changes: θ is the phase rotation due to CFO between Alice and Bob, while η is the phase rotation due to CFO between Alice and Mallory. It can be seen that when $|\theta - \eta|$ is small, the variance of CFO estimates is also small, which makes the jamming signals hard to be detected. This is the main reason why per-frame random CFO needs to be injected by Bobs before sending the channel sounding pilots (as discussed in Sec. IV-E). Nevertheless, when $|\theta - \eta|$ increases, the variance also quickly increases.

E. ROC Curves of MACE

To characterize the performance of *MACE*, we plot its ROC curves for both non-protocol-specific and protocol-specific jamming: the false positive is the mistaken detection rate when there are no jamming signals, while the true positive is the correct detection rate when there are jamming signals. For performance evaluation baselines, we also consider the following 3 detectors that employ the repeated symbols received by Alice, and compare their performance to *MACE*:

(1) Raw-CFO. As discussed in Sec. IV-C, *MACE* normalizes the CFO estimates by the corresponding SNR. In contrast, Raw-CFO does not normalize the CFO estimates and directly compute their variance.

(2) MSE-Abs-Value. Without jamming signals, Alice's i^{th} antenna receives $\{Y_{1i}|Y_{2i}\}$, where $Y_{1i} = R_i + W_{1i}$ and $Y_{2i} = R_i e^{j\theta} + W_{2i}$. Therefore, $E\{||Y_{i1}| - |Y_{i2}||^2\}$ should be small and is only related to the noise strength. MSE-Abs-Value normalizes $E\{||Y_{i1}| - |Y_{i2}||^2\}$ by the noise strength of each antenna and computes the average over all antennas.

(3) MSE-Raw-Value. Different from MSE-Abs-Value, MSE-Raw-Value computes $E\{|Y_{i1} - Y_{i2}|^2\}$.

Non-Protocol-Specific Jamming. Fig. 7(a) plots the ROC curves of the 4 detectors with white Gaussian jamming signals, where a single detection threshold is set for a wide range of SNR ($5 \sim 35$ dB) and SIR ($-5 \sim 35$ dB) with SNR$-$SIR$\geq$5 dB. It can be observed that *MACE* achieves 0.97 true positive at 0.01 false positive. In contrast, Raw-CFO only achieves 0.50 true positive at the same false positive. This is mainly because the large variance of SNR at Alice's different antennas leads to a relatively large variance of raw CFO estimates (compared to *MACE*) even without jamming signals.

In comparison, MSE-Raw-Value has even worse detection performance than Raw-CFO, while MSE-Abs-Value has similar detection performance to *MACE*. The main reason is the CFO between Alice and Bob, which makes $E\{|Y_{i1} - Y_{i2}|^2\}$ have a large value even without jamming signals. However, MSE-Abs-Value addresses this problem by taking the absolute value of the received signals (i.e., $E\{||Y_{i1}| - |Y_{i2}||^2\}$).

Protocol-Specific Jamming. As discussed in Fig. 6(d), the detection performance of *MACE* when there is protocol-specific jamming is closely related to the CFO between Bob and Mallory. Therefore, in order to plot the expected ROC curves, we vary $|\theta - \eta|$ between $0 \sim \pi$ in the experiments, where θ is the phase rotation due to CFO between Alice and Bob, while η is the phase rotation due to CFO between Alice and Mallory. After that, we group the data based on $|\theta - \eta|$ by dividing $0 \sim \pi$ into bins with 0.1 rad width. ROC curves of each bin is computed first and then the expected ROC curves over all bins are obtained. The results are shown in Fig. 7(b).

It can be observed that, similar to non-protocol-specific jamming, Raw-CFO and MSE-Raw-Value have relatively poor detection performance. Contrarily, while *MACE* still achieves 0.97 true positive at 0.01 false positive, the true positive of MSE-Abs-Value quickly decreases to 0.78. A main reason is shown in Fig. 6(c) and 6(d), which demonstrate that when $|\theta -$

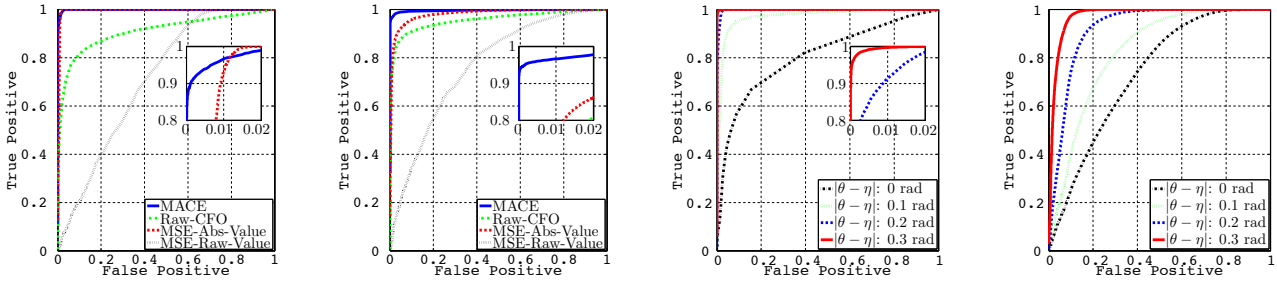

(a) (b) (c) (d)

Fig. 7. ROC curves of (a) non-protocol-specific and (b) protocol-specific (average over different CFO between Bob and Mallory) jamming signals. For protocol-specific jamming signals, we further plot the ROC curves with different CFO between Bob and Mallory for (c) *MACE* and (d) MSE-Abs-Value detector. The range of SNR and SIR are $5 \sim 35$ dB and $-5 \sim 35$ dB, respectively, with SNR$-$SIR$\geq$5 dB. Alice has 72 antennas.

$\eta|$ is small, *MACE* has a much better detection performance than MSE-Abs-Value. This is because for MSE-Abs-Value, the result of $E\{||Y_{i1}| - |Y_{i2}||^2\}$ mainly depends on the noise strength, while for *MACE*, the variance of the CFO estimates is related to the SINR (Eq. (9)). Therefore, even if $|\theta - \eta|$ is small, the change of SINR can still be detected by *MACE*.

Finally, as can be seen in Fig. 7(a) and 7(b), the true positive of *MACE* with protocol-specific jamming converges more slowly to 1 (with increasing false positive) when compared to non-protocol-specific jamming. This is mainly because there is still a chance that $|\theta-\eta|$ is small even if Bob injects per-frame random CFO before sending his channel sounding pilots.

Therefore, for both non-protocol-specific and protocol-specific jamming, by setting a single threshold, MACE can achieve 0.97 true positive at 0.01 false positive for various client/adversary locations, and for a wide range of SNR ($5 \sim 35$ dB) and SIR ($-5 \sim 35$ dB) with SNR$-$SIR$\geq$5 dB.

F. Impact of Number of Antennas and Repeated Symbols

As shown in Fig. 5(a), pilot distortion attacks lead to larger reduction of per-client achievable rate when Alice has fewer antennas. In the following, we explore whether *MACE* can still detect jamming when Alice's antenna number reduces.

Fig. 8(a) shows the true positive (at 0.03 false positive) for both non-protocol-specific and protocol-specific jamming when Alice's antenna number increases from 2 to 72. When the number of antennas increases, the true positives for both types of jamming increase. This is mainly because with fewer antennas, the variance of both ϕ in Eq. (10) and s' in Eq. (12) increases, thereby leading to a larger variance of normalized CFO estimates even without jamming signals. However, for protocol-specific jamming, because there is always a chance that $|\theta - \eta|$ (Eq. (16)) is small, its true positive quickly saturates, and becomes smaller than that of non-protocol-specific jamming afterwards. Nevertheless, for both types of jamming, *MACE* can achieve 0.97 true positive at 0.03 false positive with only 16 antennas. For larger than 5 dB difference between SNR and SIR, an even smaller number of antennas are required at the AP.

Furthermore, we also study the detection performance of *MACE* when fewer than 64 (which is the length of 1 LTS) repeated symbols are employed. In particular, we reduce the number to as few as 1, and the results with Alice having 72 antennas are shown in Fig. 8(b). Compared to Fig. 8(a), it can be seen that, while the true positive decreases with the

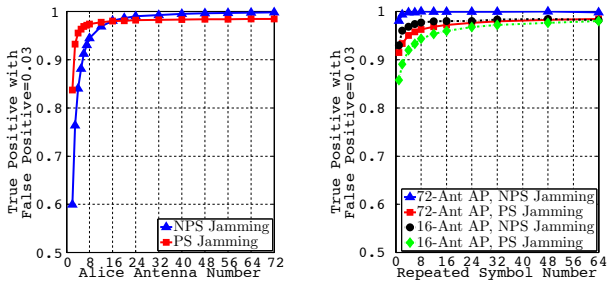

(a) (b)

Fig. 8. True positive at 0.03 false positive (a) when Alice has different number of antennas (with 64 repeated symbols) and (b) when different number of repeated symbols are input into *MACE* (with 16 or 72 antennas). The SNR and the SIR is within $5 \sim 35$ dB and $-5 \sim 35$ dB, respectively, with SNR$-$SIR$\geq$5 dB. In the figures, NPS and PS stand for Non-Protocol-Specific and Protocol-Specific jamming, respectively.

number of repeated symbols, the operational limit of *MACE* is primarily from the number of antennas at Alice. If we set a same threshold with 0.97 true positive at 0.03 false positive, we can observe that *MACE* needs to use at least 16 repeated symbols for a 72-antenna array, or 32 repeated symbols for a 16-antenna array.

Therefore, even with only 16 antennas at the AP and 32 repeated symbols, MACE can achieve 0.97 true positive at 0.03 false positive with the same client/adversary locations and SNR/SIR range; consequently, MACE can also be used for general-purpose jamming detection, even with a moderate number of antennas and repeated symbols (e.g., cyclic prefix of OFDM symbol).

VI. Related Work

Pilot Distortion Attacks. Because the improvements brought by massive MIMO are closely related to the accuracy of clients' CSI at the AP, a smart adversary can significantly degrade network performance by reducing accuracy of CSI measurements. Thus, jamming during channel sounding to aid active eavesdropping in massive MIMO networks was studied in [5]–[7]. Due to the channel sounding pilots from the eavesdropper, the AP now measures a combination of the client's and the eavesdropper's channel, which will significantly reduce the client's secrecy rate. Moreover, if the adversary jams both channel sounding and data transmission, clients' achievable rates were shown to quickly saturate even with unlimited antennas at the AP [3], [4].

In comparison, we present pilot distortion attacks, and show that even if the adversary is active only during channel sounding, which takes no more than $1/60$ of the time, concurrent uplink transmission in practical massive MIMO networks can be thwarted. We further demonstrate by experiments that an adversary having no more transmit power than any client can lead to large reduction of achievable rate of all clients.

Jamming Detection. Various techniques have been proposed to detect jamming in wireless networks. However, when they are applied to pilots in massive MIMO networks, a first problem will be the excessively high startup cost (training time). Because concurrent uplink transmission is employed, much longer time is needed for the AP to collect enough single-user transmissions from a specific client in order to compute a priori statistics of the packet delivery ratio [16], the received signal strength [17], or the angle-of-arrival information [6]. Moreover, the concurrent transmission also makes the AP hard to differentiate packet decoding error due to incorrect CSI from that due to expired CSI, which renders jamming detection based on packet decoding error less effective [18].

Another problem lies in the network efficiency. Different from data packets, channel sounding pilots are management frames that have lengths as short as several training sequences. Consequently, jamming detection should only introduce minimum network overhead. Techniques that are based on embedded secret keys [19], specially designed random PSK symbols [20], and information exchange between AP and clients [21], [22] all add to network overheard.

In comparison, we propose *MACE*, which employs the capabilities of the many antennas at the AP to detect jamming with zero startup cost, zero additional network overhead, and no shared secrets between the AP and the clients. We also implement *MACE* in our testbed and show that it achieves superior detection performance for practical massive MIMO.

Lastly, CFO has been employed to enhance network security, especially for device fingerprinting, e.g., [23], [24]. *MACE* differs from them in that *MACE* does not need to estimate the value of the CFO. Instead, *MACE* uses the variance of the CFO estimates of a single frame at the AP for jamming detection.

VII. Conclusion

In this paper, we present the Pilot Distortion Attacks, which show that an adversary jamming only the channel sounding of even a single client can lead to all-client denial-of-service in practical massive MIMO networks. As a counter mechanism, we propose *MACE*, which detects jamming with zero startup cost and zero additional network overheard, and requires no shared secrets. Our experiments show that *MACE* can achieve 0.97 true positive at 0.01 false positive.

VIII. Acknowledgments

The authors would like to thank Abeer Javed, Jian Ding, and Clayton Shepard for their assistance in performing the experiments. This research was supported by Cisco, Intel, the Keck Foundation, and by NSF grants CNS-1642929, CNS-1514285, and CNS-1444056.

References

[1] H. Ngo, E. Larsson, and T. Marzetta, "Energy and Spectral Efficiency of Very Large Multiuser MIMO Systems," *IEEE Transactions on Communications*, vol. 61, no. 4, pp. 1436–1449, 2013.

[2] J. Zhu, R. Schober, and V. Bhargava, "Secure Transmission in Multicell Massive MIMO Systems," *IEEE Transactions on Wireless Communications*, vol. 13, no. 9, pp. 4766–4781, 2014.

[3] H. Pirzadeh, S. Razavizadeh, and E. Björnson, "Subverting Massive MIMO by Smart Jamming," *IEEE Wireless Communications Letters*, vol. 5, no. 1, pp. 20–23, 2016.

[4] T. T. Do, H. Q. Ngo, T. Q. Duong, T. J. Oechtering, and M. Skoglund, "Massive MIMO Pilot Retransmission Strategies for Robustification Against Jamming," *IEEE Wireless Communications Letters*, vol. 6, no. 1, pp. 58–61, 2017.

[5] X. Zhou, B. Maham, and A. Hjorungnes, "Pilot Contamination for Active Eavesdropping," *IEEE Transactions on Wireless Communications*, vol. 11, no. 3, pp. 903–907, 2012.

[6] D. Kapetanovic, G. Zheng, and F. Rusek, "Physical Layer Security for Massive MIMO: An Overview on Passive Eavesdropping and Active Attacks," *IEEE Communications Magazine*, vol. 53, no. 6, pp. 21–27, 2015.

[7] Y. Basciftci, C. Koksal, and A. Ashikhmin, "Securing Massive MIMO at the Physical Layer," in *Proceedings of IEEE CNS 2015*.

[8] B. Danev, H. Luecken, S. Capkun, and K. Defrawy, "Attacks on Physical-Layer Identification," in *Proceedings of ACM WiSec 2010*.

[9] WARP, http://mangocomm.com/, 2017.

[10] "Argos Many-Antenna Base Station," http://argos.rice.edu/, 2017.

[11] 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*.

[12] K. Li, Y. Chen, R. Sharan, T. Goldstein, J. Cavallaro, and C. Studer, "Decentralized Data Detection for Massive MU-MIMO on a Xeon Phi Cluster," in *Proceedings of IEEE Asilomar 2016*.

[13] 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*.

[14] P. Nayak, M. Garetto, and E. Knightly, "Multi-User Downlink with Single-User Uplink can Starve TCP," in *Proceedings of IEEE INFOCOM 2017*.

[15] P. Moose, "A Technique for Orthogonal Frequency Division Multiplexing Frequency Offset Correction," *IEEE Transactions on Communications*, vol. 42, no. 10, pp. 2908–2914, 1994.

[16] 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*.

[17] S. Im, H. Jeon, J. Choi, and J. Ha, "Secret Key Agreement under An Active Attack in MU-TDD Aystems with Large Antenna Arrays," in *Proceedings of IEEE GLOBECOM 2014*.

[18] M. Strasser, B. Danev, and S. Čapkun, "Detection of Reactive Jamming in Sensor Networks," *ACM Transactions on Sensor Networks*, vol. 7, no. 2, p. 16, 2010.

[19] R. Miller and W. Trappe, "On the Vulnerabilities of CSI in MIMO Wireless Communication Systems," *IEEE Transactions on Mobile Computing*, vol. 11, no. 8, pp. 1386–1398, 2012.

[20] D. Kapetanovic, G. Zheng, K. Wong, and B. Ottersten, "Detection of Pilot Contamination Attack Using Random Training and Massive MIMO," in *Proceedings of IEEE PIMRC 2013*.

[21] D. Kapetanovic, A. Al-Nahari, A. Stojanovic, and F. Rusek, "Detection of Active Eavesdroppers in Massive MIMO," in *Proceedings of IEEE PIMRC 2014*.

[22] S. Im, H. Jeon, J. Choi, and J. Ha, "Robustness of Secret Key Agreement Protocol with Massive MIMO under Pilot Contamination Attack," in *Proceedings of IEEE ICTC 2013*.

[23] K. Cho and K. Shin, "Fingerprinting Electronic Control Units for Vehicle Intrusion Detection," in *Proceedings of USENIX Security 2016*.

[24] D. Moser, P. Leu, V. Lenders, A. Ranganathan, F. Ricciato, and S. Capkun, "Investigation of Multi-Device Location Spoofing Attacks on Air Traffic Control and Possible Countermeasures," in *Proceedings of ACM MobiCom 2016*.