Feasibility of Passive Eavesdropping in Massive MIMO: An Experimental Approach

Chia-Yi Yeh, Edward W. Knightly
Department of Electrical and Computer Engineering
Rice University, Houston, Texas 77005-1892
{chia-yi.yeh, knightly}@rice.edu

***Abstract*—Massive MIMO has the potential to thwart passive eavesdropping as the signals transmitted by a large antenna array become highly focused. Indeed, the impact of passive eavesdropping has been shown to be negligible when the number of base station (BS) antennas approaches infinity for independent Rayleigh channels. In this paper, we experimentally explore eavesdropping in Massive MIMO incorporating real-world factors including a limited BS antenna array size, potential correlation in over-the-air channels, and adaptation of modulating and coding schemes (MCS) over a discrete and finite set. Using a 96-antenna ArgosV2 BS, we *(i)* explore scaling the array size; *(ii)* identify eavesdropper advantages due to channel correlation and the resulting increase in array size required to mitigate this advantage; *(iii)* identify the "MCS saturation regime" as a vulnerability even with high SNR, *(iv)* characterize transmit power control counter strategies at the BS, and *(v)* explore the impact of a nomadic eavesdropper that moves to find the most favorable position.**

I. Introduction

Wireless links are vulnerable to passive eavesdropping since wireless signals are broadcast into the air, and any device receiving a strong enough signal can overhear the message intended for the target user (Bob). When the transmitter (Alice) has multiple antennas, she can use beamforming to concentrate the signals to Bob attempting to avoid being intercepted by the eavesdropper (Eve). Indeed, in theory, Massive MIMO systems, in which the base station (BS) is equipped with many antennas (order of a hundred) are immune to passive eavesdropping. For example, prior work has shown that as the number of BS antennas approaches infinity, the secrecy rate approaches the channel capacity and therefore the threat of passive eavesdropping is negligible [1][2].

In this paper, we present the first *experimental* evaluation of eavesdropping in a massive MIMO system employing the Rice Argos massive MIMO platform [3] and approximately 120,000 channel measurements. In contrast to prior theoretical studies, we necessarily incorporate several key factors for a practical system. First, the antenna array size in real massive MIMO systems is limited due to cost and space constraints. Second, over-the-air channel measurements can differ from idealized MIMO models with independent channels [4][5]. Lastly, practical systems are constrained by a discrete and limited set of modulation and coding schemes (MCS). For example, when Alice selects a lower MCS, it gives Eve an increased opportunity to decode the transmission at the physical layer, as she requires lower SNR to do so.

In our experiments, we first explore the role of scaling Alice's array size by sub-sampling measurements from the actual 96 element array. For all scenarios, we also perform Monte Carlo simulations using independent Rayleigh channels as a baseline. First, we find that in the moderate-antenna regime, e.g., below 8 antennas, Bob's and Eve's measured SNR scales as predicted by the baseline channel models. However, in the many-antenna regime, Eve obtains a modest advantage over the idealized model due to channel correlation, with the gap between the measured channels and Rayleigh channels increasing with the number of BS antennas. We find that despite the rich multi-path environment of the indoor channels at 2.4 GHz, a significant Line of Sight (LOS) component nonetheless yields correlation that corresponds to a 4 dB advantage for Eve when the BS array size is 96.

Second, we explore MCS adaptation and identify a critical regime that we term *MCS saturation*, In this regime, the channel from Alice to Bob is sufficiently strong that the BS could increase its MCS, yet it cannot because no higher order MCS is available (e.g., 64 QAM rate 3/4 in our experiments). Once the MCS saturates, the BS can no longer take advantage of Bob potentially having a better channel than Eve, risking that they can both decode the packet.

Third, we consider that the BS (Alice) employs power control in order to thwart Eve. Namely, Alice will attempt to provide the maximum power possible in order to maximize the data rate to Bob, yet avoid the MCS saturation regime which would reduce security without providing any gain in data rate. We define secure packet deliver ratio (s-PDR) as the fraction of packets decoded by Bob but not Eve. We show that in the Massive MIMO regime, a large set of transmit powers, spanning approximately 20 dB, can yield s-PDR of approximately 0.95, whereas an 8 antenna system cannot reach a s-PDR of 0.7, even under perfect power allocation.

Finally, we consider a nomadic Eve who attempts to find a better location with increased channel gain to overhear the Alice-Bob transmission. We begin with Eve on the same radius as Bob, and find that, perhaps surprisingly, having a low angular spread to Bob (Eve closer to Bob), does not help Eve. Nonetheless, the positions on the radius have as much as 12 dB spread from worst to best, and if Eve checks all points, she will gain significantly. Likewise, Eve may position herself as close to Alice as possible to improve her channel gain. Such a strategy requires an increasing large array size to counter. For example, necessarily using simulations, we show that Eve can

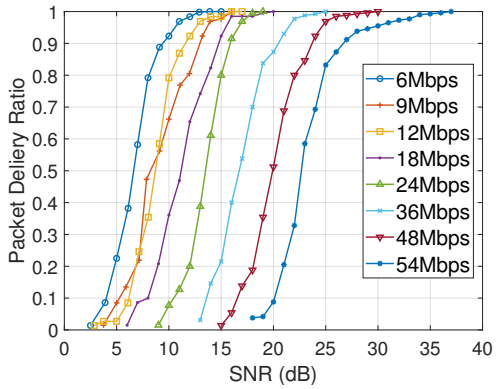

Fig. 1. PDR as a function of SNR. Adapted from [8].

force Alice to require 350 instead of 70 antennas to counter her movement to a position that is 5 dB better than at Bob.

In the following, we first describe the system model and threat model in Section II. In Section III, we describe the methodology we use to obtain the results. We then explore the scaling antenna resources and transmit power adaptation in Section IV and Section V respectively. In Section VI we discuss the threat of a nomadic Eve. And we conclude in Section VII.

II. System Model

A. Massive MIMO Downlink Transmission

We consider a multi-antenna base station (Alice or BS) and a single-antenna user Bob. The BS has M antennas, which can be as large as hundreds to explore the Massive MIMO regime. The BS performs SNR-based rate adaptation in a manner similar to [6], [7], [8]. In particular, the BS maximizes data rate by selecting the highest MCS that can support packet delivery ratio (PDR) of γ close to one, where PDR is a pre-defined function of the measured SNR. Although the SNR-PDR relationship is hardware dependent, the achievable MCS increases with SNR and has a sigmoidal transition period in which PDR rapidly increases. Fig. 1, adapted from [8], illustrates the general trend of the SNR-PDR relationship, and is used in our analysis.

In a downlink transmission, the BS first obtains Bob's CSI. In a TDD massive MIMO system for example, the BS estimates the uplink channel from Bob to the M antennas at the BS using the uplink pilot transmitted by Bob, and the downlink channel is obtained assuming channel reciprocity to avoid the high overhead caused by the downlink training scaling with M [9]. Specifically, the BS compares the received uplink pilot signals with the known pilot sequence and obtain a complex number uplink channel coefficient for each of the M antennas, and the M channel coefficient constitutes the uplink channel vector $\boldsymbol{g}_{ul} \in \mathbb{C}^{M\times 1}$. The downlink channel vector to Bob $\boldsymbol{g}_b \in \mathbb{C}^{1\times M}$, assuming channel reciprocity, is the transpose of the uplink channel vector $\boldsymbol{g}_b = \boldsymbol{g}_{ul}^T$. We can further express the downlink channel as $\boldsymbol{g}_b = \sqrt{\beta_b}\boldsymbol{h}_b$, where $\sqrt{\beta_b}$ and $\boldsymbol{h} \in \mathbb{C}^{1\times M}$ represent the large-scale and the small-scale fading respectively, and the small-scale fading is normalized such that $\mathbf{E}\left[\|\boldsymbol{h}\|^2\right] = M$.

Based on the acquired CSI, the BS estimates Bob's SNR and selects the MCS level according to the known SNR-PDR relationship. The BS then transmits the packet with the selected MCS. The BS transmits using conjugate beamforming to maximize the receive signal strength at Bob. That is, the beamforming weights $\boldsymbol{w} = \frac{\sqrt{\beta_b}\boldsymbol{h}_b^H}{\left\|\sqrt{\beta_b}\boldsymbol{h}_b\right\|}$, where the superscript H denotes Hermitian transpose. The downlink transmission is successful if Bob can decode the packet.

B. Threat Model

We consider a passive eavesdropper (Eve) in range of the BS (Alice), trying to intercept the downlink signals from the BS to Bob. To avoid being discovered, Eve passively monitors the channel without transmitting. We investigate the case of symmetric passive eavesdropping in which Eve has the same capability as Bob. While Eve and Bob can have different capabilities, e.g., Eve may have multiple antennas, the symmetric case enables focus on the effect of the propagation channels. Therefore, we consider Eve to have a single antenna and have the same decoding ability as Bob. In addition, we will first examine the case where Eve has the same pathloss as Bob to exclude the effect of pathloss and focus on beamforming gain. We then extend the discussion to Bob and Eve with different pathloss.

C. Secure Transmission in Practical System

Secured transmission in practical systems is achieved when the target user Bob successfully decodes the signals, while the eavesdropper Eve fails to do so. Therefore, we call the transmission between BS and Bob secured when Bob has high packet delivery ratio (PDR) while Eve has low PDR.

To maintain high PDR at Bob while suppressing PDR at Eve, BS's best strategy is to choose the highest supportable MCS for the predicted Bob's SNR as it forces Eve to have higher SNR to decode the signals for Bob. This strategy, fortunately, aligns with Alice and Bobs incentive to maximize throughput. In the following, we consider that the BS chooses the highest MCS which can achieve PDR of γ, in which γ is close to 1. This strategy ensures a PDR of γ at Bob unless Bob's SNR is too low so that even the lowest MCS cannot be supported.

III. Methodology

In this section, we describe the Rice massive MIMO platform and channel measurement dataset that we use to experimentally evaluate the vulnerability of a practical massive MIMO network to passive eavesdropping. We describe post-processing methods to compute SNR and PDR at Bob and Eve and introduce a Monte Carlo simulation method to study Rayleigh channels as a baseline for comparison. Lastly, we introduce two security metrics that we use throughout the analysis.

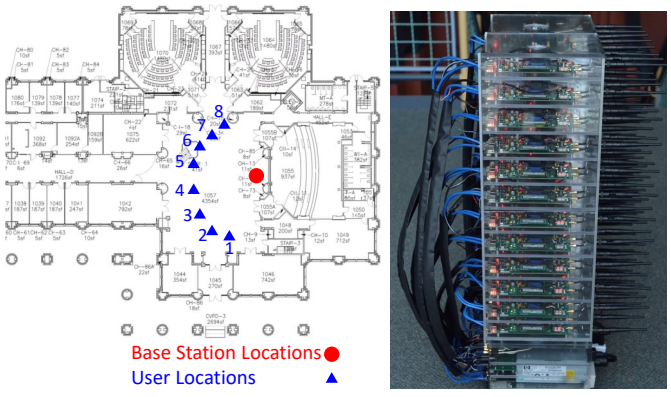

(a) Measurement topology (b) Argos BS

Fig. 2. Topology of the indoor channel measurement and the Argos BS

A. Massive MIMO CSI Measurement Dataset

The channel measurement dataset we use is from [10]. The channel measurement are taken using the ArgosV2 BS [3] with 96 antennas and 8 single-antenna WARP boards [11] in Duncan hall at Rice University in 2.4 GHz band with 20 MHz bandwidth. The 8 WARP boards are placed at the same distance from the ArgosV2 BS, all with a direct path to the ArgosV2 BS, as shown in Fig. 2a. The distance from the BS to each WARP boards is about 13 meters, and the spacing between adjacent WARP boards is about 3 meters. The 8 WARP boards are clients that can serve as Bob or Eve. The 96-antenna ArgosV2 BS is shown in Fig. 2b. Specifically, the 96 antennas are placed 8 in a row with total 12 rows on a plane, with spacing of 6.25 cm, which is half-wavelength of 2.4 GHz.

In each channel measurement, the 8 WARP boards send uplink pilots in the 52 subcarriers in a time-division manner, and ArgosV2 BS estimates the channels from the 8 single-antenna WARP boards to its 96 antennas. The reversed channels can be obtained assuming channel reciprocity. Therefore, a 96x8x52 channel matrix is obtained in each measurement, capturing the channels from the 96 BS antennas to the 8 WARP boards across the 52 subcarriers. The channel measurements are taken every 2 ms for total 30 seconds, resulting total 15,000 snapshots. Since there is no device or environment mobility, CSI remains relatively static over the 30 seconds. Therefore, we treat the 15,000 measurement epochs as one channel state realization, and only use adjacent measurements for normalization.

B. CSI Processing and SNR Calculation

1) Subsampling: Now that we obtain a 96x8 channel measurement from the dataset, we describe how we extract different channel realizations from this channel matrix.

First, we can obtain channel realizations for different Bob-Eve Pairs. Since each of the 8 WARP boards can be viewed as Bob or Eve, there are total $8 \times 7 = 56$ different Bob-Eve combinations.

Also, to explore different numbers of transmit antennas at the BS, we subsample the 96 antennas. To preserve the physical structure of the array, we only subsample adjacent antenna elements to form linear arrays of less than 8 elements or rectangular array with 8 elements in a row.

2) SNR Calculation Using Measured CSI: Once the CSI measurements are subsampled, we obtain the channel vectors for Bob and Eve, for different numbers of BS antennas. With Bob and Eve's channel, the SNR at Bob and Eve can be calculated using

$$\begin{aligned} \text{SNR}_{Bob} &= \frac{p\beta_b}{\sigma^2} \|\boldsymbol{h}_b\|^2 \\ \text{SNR}_{Eve} &= \frac{p\beta_e}{\sigma^2} \frac{\left|\boldsymbol{h}_e \boldsymbol{h}_b^H\right|^2}{\|\boldsymbol{h}_b\|^2} \end{aligned} \tag{1}$$

where p and σ^2 are the BS transmit power noise power; Bob and Eve's channels are represented with large-scale fading β and small-scale fading $\boldsymbol{h}$.

C. MCS Selection and Packet Delivery Based on SNR-PDR Relationship

We assume the SNR-PDR relationship is known to the BS and the SNR-PDR relationship used in our analysis is shown in Fig. 1, which is the result from [8]. Although the SNR-PDR relationship can be different for different systems, Fig. 1 captures the general characteristics so that our analysis can be applied in general systems. Based on the predicted Bob's SNR, the BS selects the highest MCS which achieves PDR of γ, which is set to be 0.9 in our analysis. When Bob's SNR is so low that even the lowest MCS cannot be supported, the BS selects the lowest MCS. Once the MCS is chosen, the success of the transmission depends on the corresponding PDR via Bob and Eve's SNR. Finally, we calculate the overall PDR at Bob and Eve, and also the percentage of packets that are securely delivered to Bob without being intercepted by Eve.

D. Baseline: Independent Rayleigh Fading Channels

While Bob and Eve's SNR behavior for Rayleigh channels are provided in [12], to study the SNR difference between Bob and Eve, as well as Bob and Eve's PDR, we use Monte Carlo simulation with 100,000 instances for the independent Rayleigh fading channel. The Monte Carlo simulations follow the same procedure as above with randomly generated Rayleigh channels rather than over-the-air measurement data.

E. Metrics for Evaluating Secure Transmission in Practical Systems

We employ two metrics to evaluate the empirical security level of a practical system.

1) Secure Packet Delivery Ratio (s-PDR): Bob's and Eve's PDR are both critical when discussing passive eavesdropping in practical systems. Only when Bob has a high PDR and Eve has a low PDR can we say the system is secured. To include both Bob and Eve's PDR in one metric, we define secure PDR (s-PDR) as

$$\text{s-PDR} = \frac{\text{Number of packets decoded by Bob but not Eve}}{\text{Total number of packets}}$$

When s-PDR is low, either Bob cannot decode packets, or both Bob and Eve can successfully decode the packet, which implies the transmission is either unsuccessful or insecure. In comparison, when the s-PDR is close to 1, Bob can decode most packets while Eve can hardly decode any so that the transmission is secure and successful. In the following, we explore how s-PDR is impacted by transmit power and the number of BS antennas.

2) SNR Difference Between Bob and Eve: From the discussion in Section II-C, we can observe that no matter which MCS is chosen, the transmission can be secure once the difference between Bob's and Eve's SNR is large enough. Therefore, the transmission is more likely to be secured when the difference between Bob and Eve's SNR is larger. Furthermore, SNR difference does not depend on the BS transmit power. Based on theses two reasons, we employ SNR difference as a metric to evaluate resilience to passive eavesdropping.

IV. Scaling BS Antenna Resources

Passive eavesdropping is affected by both BS array size and BS transmit power. To explore the factors separately, in this section, we explore the effect of scaling BS antennas by fixing the BS transmit power. In the following, we first examine the SNR and PDR at Bob and Eve, as well as the selected MCS for the transmissions for both Rayleigh channels and the measured channels. Although these results are obtained under a specific BS power constraint and therefore behave differently in different transmit power regimes, a full understanding of one specific scenario helps us relate to other scenarios. After exploring the case of a selected BS transmit power, we present the SNR difference between Bob and Eve that holds under any BS transmit power and further predict Eve's advantage in the over-the-air channels for even larger arrays.

A. Baseline: Independent Rayleigh Fading Channels

We first examine how Bob and Eve's SNR changes as the number of BS antennas increases for a baseline case of independent Rayleigh channels. In particular, this will enable us to quantify for idealized channels the SNR advantage of Bob over Eve as the BS devotes an increasing number of antenna resources to providing a beamforming gain to Bob, which will not benefit Eve.

As described in Section III, we vary the number of BS antennas from 1 to 96. For each antenna size, we generate 100,000 independent Rayleigh channel realizations for Bob and Eve. The BS precodes the signals for Bob with conjugate beamforming, and Bob and Eve's SNR is calculated by Equation (1).

Fig. 3a shows the median of Bob and Eve's SNR in dB vs. the number of transmit antennas for a fixed total transmit power. Also, the 90% confidence interval is shown with the 5 and 95 percentiles. Since SNR depends on BS transmit power, channel gain, and noise power, we choose the BS transmit power so that the median of Bob's SNR falls at 10 dB when the BS uses a single antenna. The choice of 10 dB allows Bob's SNR to fall in the MCS operating range, spanning from approximately 10 dB to 30 dB.

As derived in [12], Bob's SNR has an Erlang distribution with shape M and scale $\frac{p\beta_b}{\sigma^2}$, and Eve's SNR has an exponential distribution with rate $\frac{\sigma^2}{p\beta_e}$. Therefore, Bob and Eve's mean SNR is $M\frac{p\beta_b}{\sigma^2}$ and $\frac{p\beta_e}{\sigma^2}$ respectively. That is, under a fixed total power constraint, we expect Bob's SNR to increase and Eve's SNR to remain the same with an increasing number of BS antennas. Indeed, that trend can be observed in Fig. 3a.

Fig. 3a indicates that Bob's SNR increases by approximately 3 dB when the size of BS antenna array doubles, which is a direct result of beamforming. In comparison, Eve's SNR remains the same no matter how many antennas are at the BS. As a result, as the BS's antenna resources are increased, it can select a higher MCS to take advantage of Bob's SNR gain. However, since Eve's SNR remains the same, she may not be able to decode the higher order MCS. Generally, if the BS can adapt the MCS to the maximum allowed by Bob's SNR, the BS will eventually successfully prevent passive eavesdropping, as it has more and more antennas. Unfortunately, we will show that this is not always possible due to limited MCS levels.

As for the SNR variation, Bob's SNR variation decreases as the BS increases the number of antennas thanks to the law of large numbers. The smaller SNR variation implies that Bob is less likely to encounters deep-fade event, and therefore avoids being forced to use lower MCS for transmissions. In contrast, Eve's SNR variation does not change with the scaling antenna resources as Bob and Eve's channels are independent. The consistent large SNR variation makes Eve encounter deep-fade events, but also gives Eve chances to have good channel conditions for eavesdropping.

Next, we examine how the higher and more converged Bob's SNR affects the MCS selection. Fig. 3b shows the selected MCS when the BS chooses the highest MCS achieving over 90% PDR for transmission, as described in section III-C. We observe that when the BS has only a small number of antennas, the beamforming gain is reduced and the BS must use a reduced MCS. As the number of BS antennas increases, the BS can increase MCS. When the BS has more than 48 antennas, Bob's SNR is so large that the BS always chooses the highest MCS. We refer to this point as "MCS saturation" since while the channel can support a higher MCS, none is supported by the standard. Also, since the variation of Bob's SNR decreases as the number of antennas increases, the variety of the selected MCS also decreases. For example, the BS chooses from MCS-1 to MCS-5 when it has 2 antennas, but only from MCS-6 to MCS-8 when it has 24 antennas.

Now we know the trend of Bob and Eve's SNR, as well as the selected MCS as the BS increases its antennas. Here we examine the resulting PDR at Bob and Eve. Fig. 3c is the overall packet delivery ratio at Bob and Eve based on SNR and MCS shown in Fig. 3a and 3b. When the BS has only a single antenna, Bob and Eve have the same PDR, since they are at the same distance from the BS. We also observe that because of the large SNR variation, Bob sometimes fails to decode packets with even the lowest MCS, which therefore results in PDR lower than 90% (which would otherwise trigger MCS adaptation). As the BS has more antennas, Bob's PDR increases and remains above 90% since Bob can at least decode packets at the lowest MCS given the higher SNR

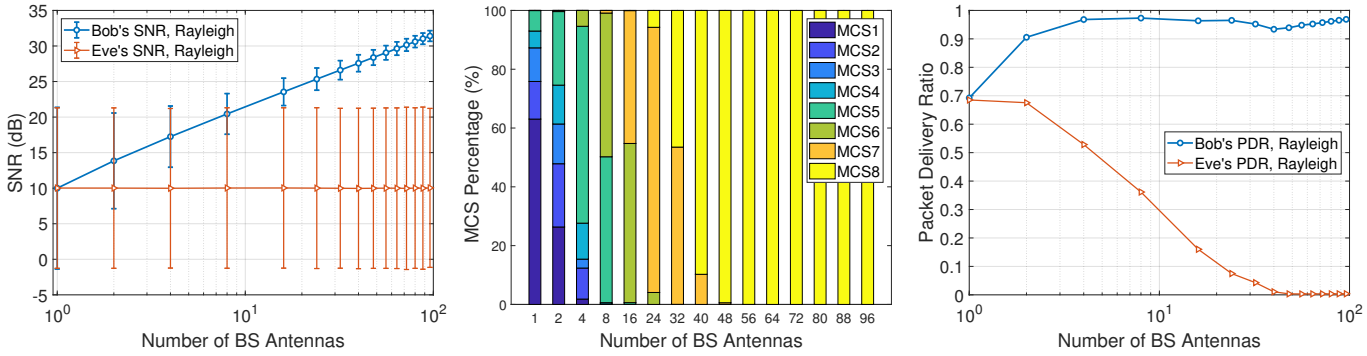

(a) Bob's and Eve's SNR (b) MCS proportion (c) Bob's and Eve's packet delivery ratio

Fig. 3. Independent Rayleigh channel

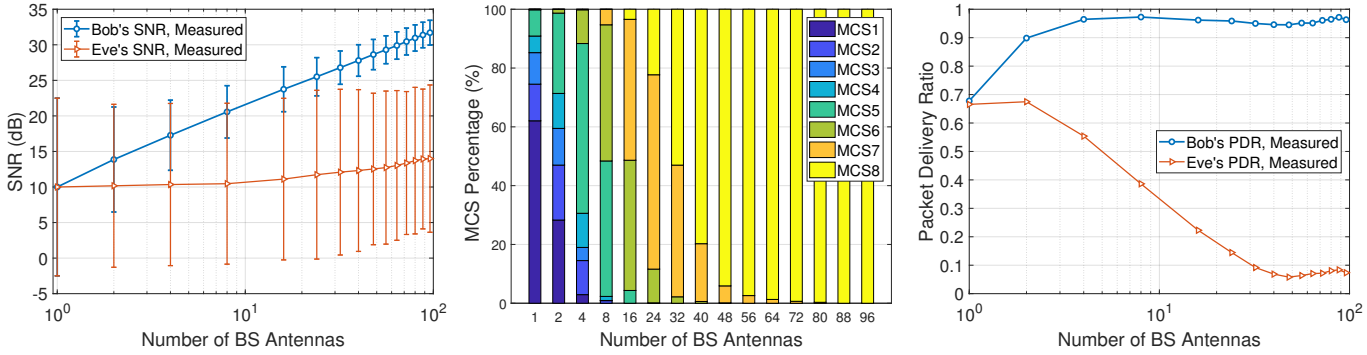

(a) Bob's and Eve's SNR (b) MCS proportion (c) Bob's and Eve's packet delivery ratio

Fig. 4. Indoor LOS

due to beamforming gains. In comparison, Eve's PDR drops with increasing BS antennas since Eve fails to decode packets transmitted with higher MCS. Eve's PDR drops below 10% at 24 antennas, and eventually drops to 0% when the BS has more than 40 antennas.

As the number of antennas and beamforming gains to Bob increase, the BS can utilize higher MCS due to Bobs increasing SNR; Eve eventually fails to decode the high-MCS packets, as her SNR (of Bobs packets) does not increase. Thus, with independent Rayleigh channels, a massive MIMO network becomes highly resistant to a single antenna passive eavesdropper when the BS has enough antennas, 24 in this same-radius eavesdropping scenario.

B. Measured Channels

Here, we apply the same methodology as above to over-the-air channel data instead of independent Rayleigh channels. As described in Section III, the following results for measured channels are based on channel measurements from a 96-antenna BS to 8 same-radius users across 52 subcarriers. Specifically, Bob and Eve can locate at any 2 of the 8 locations, and the the 96 antennas are subsampled to emulate different sizes of antenna array. Therefore, the results include all Bob-Eve location combinations, all sub-array configurations, and 52 subcarriers.

Fig. 4a shows Bob and Eve's SNR with 5, 50, and 95 percentiles. Similar to Rayleigh channels, Bob's SNR increases by approximately 3 dB when the number of BS antennas doubles, with only a slight difference that the variation of Bob's SNR in the measurements is larger than that in the Rayleigh channels. However, in contrast to the Rayleigh case, Eve's SNR based on the measured channels does not remain the same when the BS increases its antennas. Instead, Eve's SNR starts to increase when the BS has more than 8 antennas. When the BS has 96 antennas, the 50 percentile of Eve's SNR has increased from 10 dB to 14 dB.

Both the larger variation of Bob's SNR and the increase of Eve's SNR are due to the LOS component of the channel. Since the LOS component causes some correlation among the channels from the BSs antennas to Bob, it is more likely that Bob's channels experience good or bad channel conditions together, leading to more extreme SNR values and larger variation in SNR. Also, the LOS component causes correlation among the channels from the BSs antennas to Eve. In addition, since Bob and Eve locate at the same distance from the BS with a direct LOS, the correlation pattern of the BS-Bob channel coefficients and that of the BS-Eve channel coefficients are similar. As a result, Eve also receives part of the beamforming gain as the BS beamforms to Bob.

Next, we explore how the larger variation of Bob's SNR and the increase of Eve's SNR affect MCS and the PDR at

Bob and Eve. We expect that both the larger variation of Bob's SNR and the increase of Eve's SNR can have a negative impact on resisting passive eavesdropping. Namely, the SNR at Bob can be particularly low in some cases due to the larger variation. As a result, the BS may be forced to choose a lower MCS, which makes decoding easier at Eve. The increase of Eve's SNR, albeit modest, also enhances Eve's probability of decoding Bobs packets. Therefore, indoor channels with a LOS component can be expected to be less resilient to passive-eavesdropping than independent Rayleigh channels.

Fig. 4b shows the MCS selected for Bob's transmissions as a function of the number of transmit antennas. Similar to Rayleigh fading channels, the measurement data indicates that packets are transmitted with higher MCS as the BS has more antennas as a result of increasing Bob's SNR. However, when comparing the MCS chosen in measured channels (Fig. 4b) to Rayleigh channels (Fig. 3b), packets tend to be transmitted with more widely varying MCS in the measured channels as a reslut of Bob's larger SNR variation. For example, when the BS has 24 antennas, packets are transmitted with MCS 6, 7, and 8, in both measured and Rayleigh channels. However, in Rayleigh channels only few packets are transmitted with MCS-6 and MCS-8 (4% and 6%), whereas in measured channels, a larger portion of packets are transmitted with MCS-6 and MCS-8 (12% and 22%). Thus, larger portion of packets transmitted with lower MCS in the measured channel, indicating that more packets are vulnerable to passive eavesdropping.

Fig. 4c shows Bob and Eve's PDR in the measured channels and we can observe the negative impact of the LOS component. While the trends of Bob and Eve's PDR are similar to the case of Rayleigh channels, Eve's PDR decreases with a slower rate for measured channels. Furthermore, Eve's PDR does not drop to zero even when the BS has as many as 96 antennas; instead, Eve's PDR hovers around 7%. The slower drop of Eve's PDR is due to a higher percentage of lower-MCS packets and a higher SNR at Eve. The higher SNR at Eve also enables Eve to decode part of the highest-MCS packets and makes Eve's PDR remain approximately 7%, even moving towards the many antenna regime.

In summary, experiments with over-the-air transmissions indicate that the real system is less resistant to a single antenna passive eavesdropper compared to Rayleigh channels. For the measured channels, 32 antennas are needed to achieve PDR$<$ 0.1, whereas only 24 are needed in the Rayleigh case.

C. Scaling Beyond 100 Antennas

We study the resistance to passive eavesdropping of systems up to 96 antennas in the previous subsection. Although we only have channel measurements up to 96 antennas, we explore array size larger than 96 using linear regression in this subsection.

We quantize the resistance to passive eavesdropping of the system by SNR difference between Bob and Eve, which holds under any BS transmit power as discussed in Section III-E. Fig. 5 shows the prediction of the median of SNR difference between Bob and Eve up to 200 antennas. The prediction is made using the last 6 data points, which corresponds to BS antennas range from 56 to 96. Note that this prediction can be too optimistic since it does not model the slower growth of SNR difference in the larger antenna regime for the over-the-air channels.

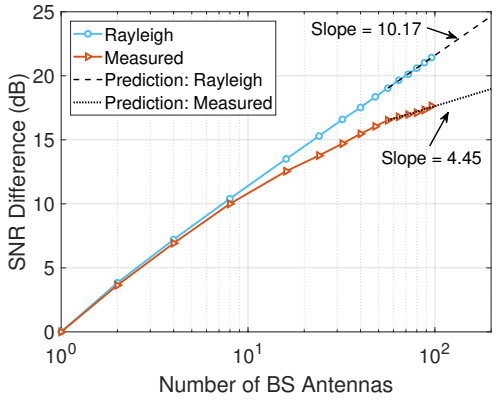

Fig. 5. SNR difference prediction.

In Fig. 5, we observe that when the array size increases by 10 times, the SNR difference increases by 10.17 dB for Rayleigh channels, but only 4.45 dB for the measured channels, which is less than half of the increment in the Rayleigh channels. Under this prediction, 200 antennas in the measured channels will only have similar security level of 60 antennas in the Rayleigh channels. And about 700 antennas are required to match the performance of 96 antennas in the independent Rayleigh channel.

The result indicates that increasing security level in practical systems is a challenging problem. Since the growth of SNR difference slows down in the large-antenna regime for the over-the-air channels, the SNR difference growth predicted by Rayleigh channels will require another order of magnitude of antennas to achieve in the practical systems. Moreover, since the results of Rayleigh channels and the over-the-air channels diverge with the scaling antennas, using Rayleigh channels to model the real channels becomes less and less applicable in the larger antenna regime.

V. Transmit Power Adaptation as a Counter Measure

Thus far, we considered that the BS uses a fixed total transmit power in all scenarios. Here, we discuss how the BS transmit power affects the security level of the transmission. We show that the security level of the transmission does not increase monotonically with Alice's transmit power due to the limited MCS levels. Therefore, Alice and Bob can increase their resilience to Eve via a counter-strategy in which Alice's transmit power is set according to the Alice-Bob channel in the MCS saturation regime. In the following, we define a transmit power performance factor suitable for different numbers of BS antennas and then analyze s-PDR for both Rayleigh and over-the-air channels.

A. Receiver-Normalized Transmit Power

As we have already explored the impact of array size on beamforming gain, here we employ two steps to explore

security as a function of both array size and transmit power. First, we normalize transmit power to the number of antennas such that the BS transmits to Bob with power p when it has a single antenna, and power $\frac{p}{M}$ when it has M antennas.

Second, we define the BS transmit power based on the receive SNR at Bob. In this way, the BS's transmit power is translated to the SNR range that Bob falls into. Specifically, we define the *receiver-normalized BS transmit power* as the transmit power scaled to Bob's median SNR when the BS has a single antenna.

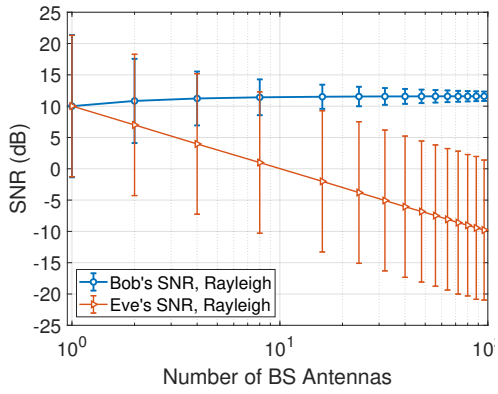

Fig. 6. Independent Rayleigh channels, BS normalize its transmit power to the number of antennas.

Fig. 6 shows an example of receiver-normalized BS transmit power of 10 dB for independent Rayleigh channels. In this example, all settings are the same as in Section IV except that the BS decreases its transmit power proportionally to its number of antennas. As shown in Fig. 6, Bob's median SNR is 10 dB when the BS has a single antenna, and thus Fig. 6 represents receiver-normalized BS transmit power of 10 dB. Since the BS decreases its transmit power proportionally to its number of antennas, the beamforming gain is largely canceled out. Therefore, Bob's normalized SNR remains approximately 10 dB as the BS increases its array size. Furthermore, the variation of Bob's SNR reduces as array size increases due to the law of large numbers. In comparison, Eve's SNR is suppressed by the reduced transmit power as the BS has more antennas. We observe that Eve's SNR decreases 3 dB when the array size doubles, with the same variation. Thus, the figure shows that when BS transmit power scales down with its number of antennas, Bob's SNR falls in similar range and allow us to compare passive eavesdropping across different numbers of BS antennas.

B. Independent Rayleigh Channels

Here, we study the power allocation strategy that makes transmissions more resilient to the passive eavesdropper for Rayleigh channels. In particular, we use s-PDR, the percentage of packets received by Bob but not Eve, as the security performance metric. We vary the number of BS antennas from 2 to 96, and the receiver-normalized BS transmit power from 2 dB to 40 dB, which also impacts Bob's SNR ranges. The results are based on Monte Carlo simulation with 100,000 instances.

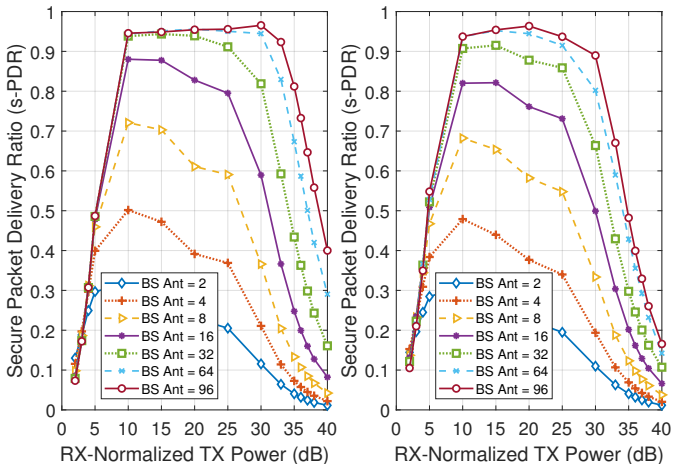

(a) Independent Rayleigh Channels (b) Measured Channels

Fig. 7. Secure packet delivery ratio (s-PDR) varies as BS increases its transmit power.

Fig. 7a depicts s-PDR vs. receiver-normalized transmit power (defined above), with a family of curves for different array sizes. For each array size, s-PDR first increases, and then decreases as the transmit power increases: When the BS transmits with power so low that neither Bob nor Eve can receive the packets, the s-PDR is close to zero. As the BS increases its transmit power, Bob starts to receive packets with the lowest MCS. Since Eve's SNR is generally lower than Bob's SNR when the BS has multiple antennas, Eve can barely decode packets for Bob at this point. Therefore, the s-PDR first grows as the BS increases its transmit power.

As the BS transmit power increases, Eve's SNR also increases. If Bob's SNR is not significantly larger than Eve's, as occurs when the BS has few antennas, it is likely that some packets for Bob can also be decoded by Eve, and the s-PDR is hence reduced. In contrast, if Bob's SNR is significantly larger than Eve's, i.e., when the BS has many antennas, Eve cannot decode packets for Bob even with her improved SNR. Therefore, the s-PDR remains high, and we can observe a plateau when the BS has a large antenna array. However, the s-PDR eventually decreases when the BS transmits with sufficiently high power, even when the BS has many antennas. In this case, Bob and Eve's SNR are both high enough to decode packets modulated with the highest MCS, resulting insecure transmissions. We term this regime *MCS saturation* as having increased MCS levels would defer this effect to higher powers. Nonetheless, Eve will eventually be able to decode all packets for Bob as the BS increases its transmit power when MCS levels are limited.

The results also reveal the impact of scaling the array size by analyzing the family of curves: First, the highest achievable s-PDR grows as the number of antennas increases. For example, the highest achievable s-PDR is only 0.3 when the BS has 2 antennas, vs. 0.5 for 4 antennas, and 0.95 when the BS has more than 32 antennas. That is, more packets can be delivered securely to Bob as the BS has more antennas if the BS chooses transmit power properly. For independent Rayleigh channels, 32 antennas at the BS are enough to securely deliver

95% of the packets.

Second, when the BS has more antennas, the transmissions remain secure for a larger transmit power region. If we consider s-PDR above 0.9 as secure transmissions, the normalized transmit power region which results in secure transmissions is 10-33 dB when the BS has 96 antennas. In comparison, the power region shrinks to 10-25 dB when the BS reduces its array size to 32. A larger antenna array at the BS increases the SNR difference between Bob and Eve, and thus relieves the precision of transmit power allocation.

In summary, a larger antenna array at the BS increases the achievable s-PDR, and also broadens the transmit power range that results in secure transmissions. However, a large antenna array does not guarantee secure transmissions, as the BS needs to be careful not to enter the MCS saturation regime, transmitting with power sufficiently high such that both Eve and Bob can decode the packets modulated with the highest MCS.

C. Measured Channels

Here, we explore the same factors using channel measurement data based on the same 96-antenna BS and 8 same-radius single-antenna as previously. Fig. 7b shows that, similar to Rayleigh channels, s-PDR first increases, and then decreases, as the BS increases its transmit power, and the highest achievable s-PDR grows when the BS has more antennas.

However, the Argos system measurements requires more antennas at the BS to ensure secure transmissions compared to Rayleigh channels. In Rayleigh channels, 16 antennas at the BS can achieve almost 0.9 s-PDR, whereas only 0.82 s-PDR is achieved in the measured channels. To achieve an s-PDR of 0.9, it requires only $\sim$20 antennas for Rayleigh channels, but $\sim$30 antennas for the measured channels. Furthermore, s-PDR drops at lower BS transmit power than for Rayleigh channels. For instance, when the BS has 96 antennas, s-PDR remains above 0.9 when receiver-normalized BS transmit power ranges from 10 to 33 dB for Rayleigh channels, but the s-PDR drops below 0.9 after 29 dB for the measured channels. This also implies the transmission is less secure in the measured channels. For example, when the normalized BS transmit power is 30 dB, having 64 antennas at the BS still ensures secure transmission of 0.95 s-PDR in Rayleigh channels, yet the s-PDR drops to 0.8 in the measured channels.

The less secure transmission in the measured channels is a result of a smaller SNR difference between Bob and Eve in the measured channels. As discussed in Section II-C, a large SNR difference between Bob and Eve makes Eve harder to decode packets for Bob, provided that the system is not in the MCS saturation regime. In the measured channels, Eve also receives a small beamforming gain when the BS beamforms to Bob with a large antenna array. This advantage makes Eve attain the SNR required to decode packets for Bob at a lower BS transmit power.

In summary, secure transmission is possible in the measured channels considering a same-radius passive Eve, but requires more transmit antennas compared to the Rayleigh channels. Also, the BS has to allocate its transmit power more carefully since s-PDR is more sensitive to BS transmit power in the measured channels than Rayleigh channels.

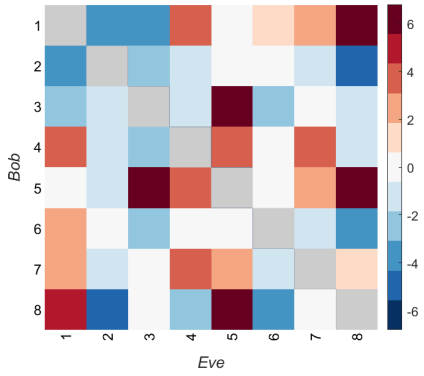

Fig. 8. SNR difference between Bob and Eve of all Bob-Eve pairs when BS has 96 antennas. Values are relative to the average case.

VI. Nomadic Eve

In the previous sections, we examine the overall passive eavesdropping behavior regardless of Bob and Eve's positions. However, there might exist some eavesdropping positions that are especially vulnerable to passive eavesdropping and will be exploited by a nomadic Eve. In addition, we consider a same-radius Eve so far, but a mobile Eve can move closer to the BS to increase her signal strength. Therefore, in this section, we first investigate the variation results from different Bob-Eve positions, in search of potential threatening location patterns in the indoor environment with a LOS component. After that, we discuss the threat as Eve move closer to the BS.

A. Bob and Eve's Relative Position

Bob and Eve's position affects the resistance to the passive eavesdropping, and we wonder whether any position pattern exists for passive eavesdropping in the indoor environment with a LOS component that helps predict weakness in such environment. Therefore, we examine the SNR difference of all Bob-Eve pairs from our 8 same-radius user dataset.

Fig. 8 shows the median of SNR difference between Bob and Eve of each Bob-Eve pair when the BS has 96 antennas. The value shown in Fig. 8 is relative to the median of SNR difference of all pairs in dB. Color red, blue, and white represents a larger, lower, and similar SNR difference compared to the average case. As Bob and Eve cannot be at the same location, the diagonal elements are represented with gray and are excluded from the discussion.

Since (Bob,Eve)=(a,b) and (Bob,Eve)=(b,a) is simply swapping positions, we can observe Fig. 8 is nearly symmetric along the diagonal. Cases in which (Bob,Eve)=(a,b) and (Bob,Eve)=(b,a) do not result in exactly the same SNR difference are due to variation of channel gain.

One natural expectation is that Eve receives higher SNR, or a smaller SNR difference, when Eve is closer to Bob, with smaller angular spread. However, we do not observe this phenomenon in the data. Fig. 8 depicts cases with a smaller

angular separation between Bob and Eve as closer to the diagonal line. If being closer to Bob aided Eve, we would see mostly blue near the diagonal, followed by white and then red when moving further from the diagonal. However, we observe cases which are especially resistant to eavesdropping (red) even when Bob and Eve are adjacent nodes, namely (Bob,Eve) = {(4,5), (5,4)}. Similarly, we also observe cases which are especially vulnerable to eavesdropping (blue) when Bob and Eve have a large angular separation, such as (Bob,Eve) = {(2,8), (8,2)}.

To further explore the role of angular separation, we calculate the correlation between the two variables, SNR difference and angular separation, and find the correlation to be only 0.101. As a result, the transmissions do not become more secure when Bob and Eve have a large angular separation. Indeed, while correlation can exist when the distance between Bob and Eve is on the order of the wavelength (12 cm for of 2.4 GHz), the distance between Bob and Eve is $\sim$3 meters in our measured topology.

B. Eve Closer to the BS

So far, we consider Bob and Eve have similar pathloss. However, a nomadic and shrewd Eve would like to move closer to the BS for higher channel gain and a favorable SNR spread from Bob. In this case, achieving secure transmissions becomes even challenging.

The results of the measured channels are based on same-radius Bob and Eve, and therefore do not represent the near-far scenario. However, we still expect the near-far scenario performs no better than the Rayleigh channels considering the channel correlation in the real world.

When Eve has a higher channel gain than Bob, the SNR difference shown in Fig. 5 is reduced by the channel gain difference between Eve and Bob. As a result, the antenna size that achieves secure transmissions previously in the same-radius Eve case fails to prevent eavesdropping as Eve moves close to the BS. To achieve secure transmissions considering a close Eve, much larger antenna array, possibly a few more hundreds of antennas, is required at the BS. For example, if Eve's channel gain is 5 dB higher than Bob's, $\sim$70 and $\sim$350 antennas are required for Rayleigh channels and measured channels respectively. Similarly, the same issue happens when Bob locates far from the BS.

In summary, although a few dozens of antennas may be enough to prevent a same-radius Eve, the pathloss gap between Bob and Eve brings the required array size to another magnitude. In theory, keep increasing the BS antenna size solve the problem of passive eavesdropping. However, the order of antennas required may not be practical considering a close-Eve or far-Bob scenario and correlated channels in the real world. Thus, other techniques such as sending orthogonal artificial noise to interfere with Eve still needs to be considered in massive MIMO systems.

VII. Conclusion

Being the first experimental study of passive eavesdropping in massive MIMO, we consider practical factors including limited BS antenna array size, potential correlation in over-the-air channels, and adaptation of MCS over a discrete and finite set. Based on channel measurements using a 96-antenna ArgosV2 BS, we have the following findings: *(i)* We find that Eve obtains a modest advantage due to channel correlation, and the gap between the measured channels and Rayleigh channels increases with the number of BS antennas. *(ii)* We identify the "MCS saturation regime" which happens when the high SNR saturates the predefined MCS levels and prevents the BS from utilizing potentially a better channel at Bob compared to Eve, indicating the importance of transmit power allocation. *(iii)* We find that having a low angular spread to Bob does not help Eve when the distance is much larger than the wavelength. However, Eve can take advantage of the high variation among different locations or move closer to the BS to improve her channel gain, forcing the BS to increase hundreds of antennas to counter.

Acknowledgements

This research was supported by Cisco, Intel, the Keck Foundation, and by NSF grants CNS-1642929 and CNS-1514285.

References

[1] 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.

[2] A. Al-Nahari, "Physical layer security using massive multiple-input and multiple-output: passive and active eavesdroppers," *IET Communications*, vol. 10, no. 1, pp. 50–56, 2016.

[3] C. Shepard, H. Yu, and L. Zhong, "Argosv2: A flexible many-antenna research platform," in *Proceedings of the 19th annual international conference on Mobile computing & networking*. ACM, 2013, pp. 163–166.

[4] D.-S. Shiu, G. J. Foschini, M. J. Gans, and J. M. Kahn, "Fading correlation and its effect on the capacity of multielement antenna systems," *IEEE Transactions on communications*, vol. 48, no. 3, pp. 502–513, 2000.

[5] P. Kyritsi, D. C. Cox, R. A. Valenzuela, and P. W. Wolniansky, "Correlation analysis based on mimo channel measurements in an indoor environment," *IEEE Journal on Selected areas in communications*, vol. 21, no. 5, pp. 713–720, 2003.

[6] G. Holland, N. Vaidya, and P. Bahl, "A rate-adaptive mac protocol for multi-hop wireless networks," in *Proceedings of the 7th annual international conference on Mobile computing and networking*. ACM, 2001, pp. 236–251.

[7] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic media access for multirate ad hoc networks," in *Proceedings of the 8th annual international conference on Mobile computing and networking*. ACM, 2002, pp. 24–35.

[8] J. Zhang, K. Tan, J. Zhao, H. Wu, and Y. Zhang, "A practical snr-guided rate adaptation," in *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*. IEEE, 2008, pp. 2083–2091.

[9] E. G. Larsson, O. Edfors, F. Tufvesson, and T. L. Marzetta, "Massive mimo for next generation wireless systems," *IEEE Communications Magazine*, vol. 52, no. 2, pp. 186–195, 2014.

[10] C. Shepard, J. Ding, R. E. Guerra, and L. Zhong, "Understanding real many-antenna mu-mimo channels," in *Signals, Systems and Computers, 2016 50th Asilomar Conference on*. IEEE, 2016, pp. 461–467.

[11] "Warp project." [Online]. Available: http://warpproject.org

[12] C.-Y. Yeh, "Feasibility of passive eavesdropping in massive mimo: An experimental approach," Master's thesis, Rice University, April 2018.