

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/383396909>

Beamforming made Malicious: Manipulating Wi-Fi Traffic via Beamforming Feedback Forgery

Conference Paper · November 2024

DOI: 10.1145/3636534.3690669

CITATIONS

0

READS

65

7 authors, including:



Xin Li

Nanyang Technological University

17 PUBLICATIONS 131 CITATIONS

SEE PROFILE



Jingzhi Hu

Nanyang Technological University

47 PUBLICATIONS 876 CITATIONS

SEE PROFILE



Jun Luo

Nanyang Technological University

218 PUBLICATIONS 8,886 CITATIONS

SEE PROFILE

Beamforming made Malicious: Manipulating Wi-Fi Traffic via Beamforming Feedback Forgery

Mingming Xu^{1*} Yinghui He^{2*} Xin Li² Jingzhi Hu² Zhe Chen³ Fu Xiao^{1†} Jun Luo^{2†}

¹School of Computer Science, Nanjing University of Posts and Telecommunications, China

²College of Computing and Data Science, Nanyang Technological University (NTU), Singapore

³Intelligent Networking and Computing Research Center and School of Computer Science, Fudan University, China

Email: {2020070137, xiaof}@njupt.edu.cn, {yinghui.he, l.xin, jingzhi.hu, junluo}@ntu.edu.sg, zhechen@fudan.edu.cn

ABSTRACT

New Wi-Fi systems have leveraged *beamforming* to manage a significant portion of traffic for achieving high throughput and reliability. Unfortunately, this has amplified certain security risks since beamforming critically relies on the *clear-text* beamforming feedback information (BFI): though similar risks have been exposed using emulation platforms (e.g., USRP), they have never proven realistic till this day. In this paper, we propose BeamCraft, the *first* attack to manipulate traffic in *commodity* Wi-Fi systems; it differs significantly from existing attacks either staying only on emulation platforms with limited real-world applicability or jamming communications by brute force. The core idea of BeamCraft involves corrupting beamforming decisions by injecting crafted BFIs that feed an access point (AP) with erroneous information on channel states. To mount a covert yet purposeful attack, we develop i) a joint location and transmit power selection strategy to evade detection by victims and ii) a novel BFI forgery method to effectively manipulate AP's beamforming decisions. We implement BeamCraft using commodity Wi-Fi devices and perform extensive evaluations with it; the results reveal that BeamCraft effectively manipulates Wi-Fi traffic while maintaining a low exposure rate.

CCS CONCEPTS

- Security and privacy → Mobile and wireless security;
- Networks → Wireless local area networks.

KEYWORDS

Wi-Fi communication, beamforming, physical layer security.

* Both authors contributed equally to this research, which is done when Mingming Xu works as a visiting scholar at NTU, under the funding support offered by Jiangsu Provincial Government Scholarship.

† Fu Xiao and Jun Luo are both the corresponding authors.



This work is licensed under a Creative Commons Attribution International 4.0 License.

ACM MobiCom'24, November 18–22, 2024, Washington D.C., DC, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0489-5/24/11.

<https://doi.org/10.1145/3636534.3690669>

ACM Reference Format:

M. Xu, Y. He, X. Li, J. Hu, Z. Chen, F. Xiao, and J. Luo. 2024. Beamforming made Malicious: Manipulating Wi-Fi Traffic via Beamforming Feedback Forgery. In *The 30th Annual International Conference on Mobile Computing and Networking (ACM MobiCom'24)*, November 18–22, 2024, Washington D.C., DC, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3636534.3690669>

1 INTRODUCTION

As wireless technologies have advanced, modern Wi-Fi systems, starting with the 802.11ac standard, have achieved Gbps links and handle over 50% of network traffic [18, 34]. This vast volume of traffic has become a prime target for various attacks aimed at manipulating or disrupting network operations. Among these, flood [11, 14, 66] and jamming [30, 39] are typical attacks. Designed to overwhelm Wi-Fi networks by incessantly injecting excessive packets or noise, these attacks are *irrational*: they necessitate significant resource expenditure (especially in high-bandwidth scenarios [6, 46]) without bringing much benefit to an attacker. As Wi-Fi bandwidth expands from 20MHz to even 320MHz in the forthcoming Wi-Fi 7 (802.11be) standard [16], the resource¹ required to launch such attacks escalate, thereby discouraging attacks and fortifying network security.

Unfortunately, the incorporation of *beamforming* into Wi-Fi introduces a new vulnerability: its dependence on *clear-text channel feedback* from users. Pivotal in enhancing Wi-Fi communication, *beamforming* enables multi-antenna access points (APs) to perform directional transmission towards users, thereby significantly boosting throughput and reliability [12, 20, 38]. Consequently, it manages nearly all traffic within Wi-Fi systems that support beamforming, as evidenced by our real-world measurements. However, the success of beamforming relies on channel feedback provided by a channel sounding protocol [7]. As the feedback is transmitted in clear-text, it is prone to interception and injection. Upon injecting forged feedback, an attack may potentially

¹Existing jamming attacks demand expensive devices (e.g., USRP) to be deployed for constantly transmitting jamming signals on all possible Wi-Fi channels [33, 39].

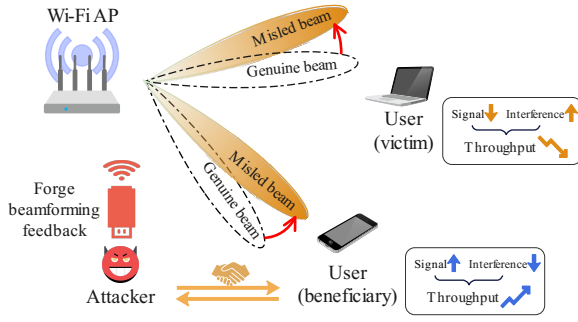


Figure 1: BeamCraft: an attacker forges beamforming feedback to a Wi-Fi AP and the AP is manipulated to align the data transmission towards a misled direction.

manipulate Wi-Fi traffic to its own benefit. Therefore, compared with traditional flood or jamming attacks, *beamforming feedback forgery* attack brings a much higher reward to the attack at a much lower exposure rate and cost (in terms of both hardware and software complexity). Although prior work [52] demonstrated successful injection attacks, they rely on software-defined radio (SDR) [2], hence only staying on emulations and applicable to specific beamforming algorithms, with limited real-world applicability.

In response to these limitations, we introduce BeamCraft, the first attack to manipulate traffic of commodity Wi-Fi systems via forged beamforming feedback. As depicted in Figure 1, the attacker forges the beamforming feedback and injects it into a Wi-Fi system, causing the AP to perform beamforming in an erroneous manner. Of course, developing such an attack in *covert yet purposeful* manner does face two major challenges. On one hand, the omnidirectional broadcast nature of feedback may potentially expose the attack to victims. In particular, as the attacker spoofs the victims' media access control (MAC) address, victims overhearing the feedback can be alerted of an attack. On the other hand, modern Wi-Fi standards only support compressed *beamforming feedback information* (BFI) containing a feedback matrix and average signal-to-noise ratio (SNR) [10], instead of full *channel state information* (CSI). Therefore, existing attack [52] fails to act on Wi-Fi systems as they forge only the whole CSI for a specific beamforming algorithm, yet such algorithms in a Wi-Fi AP are hidden from all users including the attacker.

To solve the first challenge, BeamCraft strategically prevents victims from correctly decoding the forged feedback. To this end, we analyze the relationship between the decoding success rate and SNR, and then develop the joint location and transmit power selection strategy for the attacker. The outcome allows for successfully decoding the forged BFI at the Wi-Fi AP but not the victim, thereby maintaining the attack's covertness. For the second challenge, we focus on

modifying the feedback matrix instead of the average SNR contained in the BFI, which represents channel direction and is pivotal to all types of beamforming algorithms. We further develop a novel BFI forgery method to manipulate the Wi-Fi traffic; it aims to control the correlation between the feedback matrix in the forged BFI and the genuine one, so as to misdirect the beam of the AP and manipulate the resulting SNR of beamforming. Finally, we implement a prototype of BeamCraft using commodity devices and conduct extensive experiments to evaluate the performance. In summary, we make the following major contributions:

- We propose BeamCraft, the first novel attack for traffic manipulation, targeting practical Wi-Fi systems by injecting forged beamforming feedback.
- We design a joint location and transmit power selection strategy, ensuring the covertness of an attacker.
- We develop a novel BFI forgery method to misdirect AP's beamforming and thus manipulate traffic.
- We implement BeamCraft prototype and evaluate it with extensive experiments. The promising results confirm that BeamCraft can manipulate Wi-Fi traffic with a low exposure rate.

The paper is structured as follows. Section 2 introduces the background for Wi-Fi beamforming, the attack model of BeamCraft, and the feasibility study. Section 3 details the design of BeamCraft, including the joint location and transmit power selection strategy and the novel BFI forgery method. Sections 4 and 5 report the experiment setting and performance evaluation results of BeamCraft, respectively. We discuss the limitations of BeamCraft and present defense strategies in Section 6. Related works are briefly captured in Section 7. Finally, we conclude our paper in Section 8.

2 BACKGROUND AND MOTIVATION

In this section, we first provide basic knowledge on Wi-Fi beamforming and present the attack model. Then, we verify the feasibility of forging feedback matrix and compare it with the existing work.

2.1 Wi-Fi beamforming

For the successful implementation of beamforming, it is essential for the Wi-Fi AP to know the channel between itself and the user(s). Therefore, a channel sounding protocol has been defined in the Wi-Fi standards [7], as shown in Figure 2. The protocol can be divided into three parts.

NDP Announcement: The Wi-Fi AP triggers the channel sounding protocol with a control frame, namely *null data packet* (NDP) *announcement*. It is broadcast to the active user(s), along with a unique dialog *token* for security purposes. This announcement prepares the user(s) for the subsequently transmitted NDP.

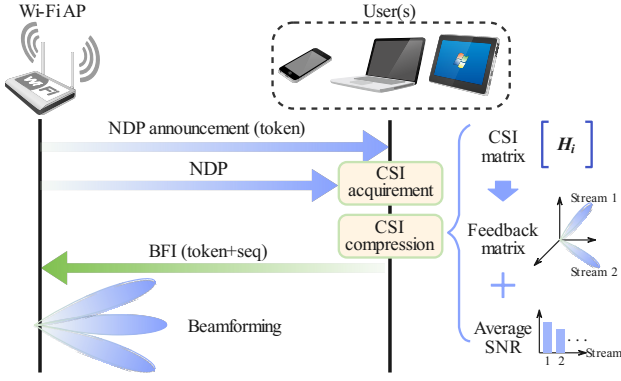


Figure 2: Wi-Fi channel sounding protocol: It starts with an NDP announcement, which is followed by CSI acquisition, demanding a BFI report summarizing the acquired CSI.

CSI Acquisition: Upon receiving the NDP broadcast by the AP, a user can measure the CSI between itself and the AP using the predefined pilot contained in the NDP. Let M^T denote the number of transmit antennas for the Wi-Fi AP and M_i^R denote the number of receiving antennas for the i -th user (which also represents the maximum number of data streams). Then the measured CSI matrix² can be represented by $H_i \in M_i^R \times M^T$.

BFI Report: After measuring H_i , the user needs to feed it back to the AP. Instead of the whole CSI matrix, a CSI compression method is adopted on BFI to reduce the overhead. Specifically, the user first calculates the feedback matrix $V_i \in \mathbb{C}_i^{M^T \times M^R}$ as the right singular vectors of H_i :

$$H_i = U_i \Lambda_i V_i^H, \quad (1)$$

where U_i and V_i are both orthonormal matrices, $\Lambda_i \in \mathbb{C}_i^{M_i^R \times M_i^R}$ is a diagonal matrix with nonnegative real values on the diagonal, and $(\cdot)^H$ denotes the conjugate transpose operation (hence $V_i^H V_i$ results in the identity matrix). Next, V_i is compressed into multiple phase values with Givens Rotation [19], and phase values are further quantified [28]. Beside the quantified phases (which are still denoted by V_i), BFI also contains the average SNR for each stream. Finally, BFI is transmitted to the AP within a management frame, along with the dialog token received from the Wi-Fi AP and a sequence number (seq). Upon receiving the BFI, the AP may act accordingly to perform beamforming towards the user for improving the transmission quality. As the compression method and format are defined in Wi-Fi standards, forging BFI is viable.

²In 802.11 standards, orthogonal frequency-division multiplexing (OFDM) technique is adopted for independently transmitting data via K subcarriers. To keep the model succinct, we ignore the subcarrier and the proposed method in our paper can be readily extended to the multiple-subcarrier case by treating each subcarrier independently.

2.2 System and Attack Models

System Model. We consider a typical communication scenario where several users connect to a Wi-Fi AP and simultaneously engage in respective online activities, such as gaming, streaming videos, and web browsing. This happens often in places covered by Wi-Fi APs, typically including cafeterias, airports, and shopping malls. To improve throughput under such multi-user scenarios, the Wi-Fi AP frequently initiates the channel sounding protocol mentioned in Section 2.1 for collecting BFIs and then acts accordingly. The Wi-Fi AP may apply different beamforming algorithms to increase throughput, including directly taking the feedback matrix as the beamforming matrix [55] or leveraging zero-forcing (ZF) beamforming [64] to reduce inter-user interference. Additionally, it also utilizes a specific transmit power allocation strategy, e.g., equal power allocation, to maintain fairness among users.

Attack Model. Under the aforementioned multi-user Wi-Fi access scenario with selfish human users, it is natural to assume that an attacker aims at a covert traffic manipulation to gain higher throughput than others, by exploiting the vulnerability of beamforming. To this end, the attacker is equipped with two Wi-Fi network interface cards (NICs): one operates in *monitor* mode [17] for sniffing packets and injecting the forged BFI, while the other one (also known as *beneficiary*) enjoys the higher throughput after traffic manipulation. Note that the need for two separated NICs is only an artifact caused by the lack of proper access to Wi-Fi firmware, which can be mitigated with future developments.

Unlike previous approaches [52] that utilize expensive SDR platforms for injecting beamforming feedback not even acceptable to contemporary Wi-Fi NICs, we assume readily available commodity Wi-Fi NICs with very low cost. We further assume that the beamforming algorithm and transmit power allocation strategy are hidden from the attacker, making the attack more meaningful than that proposed in [52]. The attacker is allowed to freely adjust its location to suit its needs, as far as it maintains the connection with the target AP.³ Finally, the attacker is assumed to be resource-constrained, hence with limited ability to perform high-frequency packet injections similar to flood attacks. This limitation underscores the necessity for strategic packet injection over methods that simply exhaust bandwidth.

2.3 Feasibility Study and Existing Work

We now verify the feasibility of injecting forged BFI, by setting up an experiment with two users and one attacker. We let two desktops equipped with two-antenna Wi-Fi NICs act as users (or victims) and connect them to a four-antenna

³The minimum distance between attacker and AP depends on physical layout: for AP at an easily accessible location, zero distance can be viable.




Figure 3: The ratio of traffic managed by the beamforming over different communication tasks.

Wi-Fi AP. The Wi-Fi AP performs beamforming with a bandwidth of 20 MHz towards two users. Meanwhile, the attacker is equipped with one two-antenna Wi-Fi NIC; it aims to sniff the BFI feedback from users and inject forged BFIs, so as to manipulate Wi-Fi traffic. For the sake of comparison, we also reproduce the power attack [52] designed for a non-standard WiFi protocol: though it cannot be directly implemented on commodity Wi-Fi NICs, we emulate it by scaling the magnitude of the whole CSI (i.e., forging average SNR in BFI), which should be practically equivalent to [52].

To report the results, we first analyze the data traffic under four communication tasks, i.e., downloading, playing online games, visiting websites, and watching online videos, as shown in Figure 3. It can be observed that almost all traffic is managed by the beamforming, indicating the potential significance of BFI injection attack. We then compare the effect of forging feedback matrix with the power attack forging only average SNR. In Figure 4 depicting the traffic variations caused by two attacks both start at 50s, one can observe that the normal throughput of 73.4Mbps drops to 49.3Mbps after the attack with forged feedback matrix is launched, whereas no obvious difference in traffic is caused by the power attack. Apparently, forging the feedback matrix can be a much more powerful attack in manipulating Wi-Fi traffic.

The reason for this difference can be attributed to the distinct channel information represented by the average SNR and feedback matrix in BFI. Since the average SNR indicates the path loss between the transmitter and receiver, forging it may affect AP’s power allocation among users (and their respective antennas). However, as common power allocation algorithm implemented by Wi-Fi AP is merely equal




Figure 4: The effect of traffic manipulation under both forged feedback matrix and scaled SNR.

allocation, such an attack virtually has no impact on Wi-Fi traffic. On the contrary, a feedback matrix characterizes the direction of a transmission path, it has a definite impact on AP’s beamforming, especially when the AP directly takes the feedback matrix as its beamforming matrix for data transmission. Consequently, a properly crafted feedback matrix can cause the AP to beam towards a wrong direction, hence effectively reducing the received SNR at a victim. Built upon this discovery, we set out to design BeamCraft in order to concretely realize the BFI injection and better exploit it.

3 THE DESIGN OF BEAMCRAFT

In this section, we first give the design overview of BeamCraft and then introduce the details.

3.1 Overview

Aiming to enable traffic manipulation attack on commodity Wi-Fi communication systems, BeamCraft follows the whole workflow shown in Figure 5. Specifically, an attacker first sniffs the packets from the Wi-Fi AP and users to measure the path loss and then determines the location and transmit power for injection (Section 3.2). Subsequently, among the sniffed packets, the BFI packets from users are first decompressed and transformed into feedback matrices, based on which the forged feedback matrices are crafted. How exactly a feedback matrix is crafted depends on the choice out of two types of the traffic manipulation: i) *traffic disruption* (Section 3.3) that the throughput of users (victims) is throttled, and ii) *traffic plunder* (Section 3.4) that the throughput of a certain user (beneficiary) is boosted by sacrificing that of another user (victim). Upon getting the feedback matrix ready and then detecting the AP’s NDP announcement, the attacker promptly updates the sounding dialog token and seq in the BFI packet and injects this forged BFI. The forged BFI packet is flagged with a “retry” to indicate a retransmission, so that it would be accepted by the AP and take effect even if the arrival time of the forged BFI is behind the genuine one. Realizing the workflow of BeamCraft requires us to tackle two main challenges:

- Whereas the injected BFI needs to be correctly decoded by the Wi-Fi AP so that the attack can take effect, voiding alerting the victim should be an equally important objective. This latter objective is made very challenging because the BFI packet is transmitted omnidirectionally by a commodity NIC. To this end, we analyze the decoding success rate and propose a strategy (Section 3.2) by jointly considering the location and transmit power of the attacker, in order to prevent the victim from correctly decoding the forged BFI.
- As we introduce in Section 2.3, injecting random feedback matrix can affect the Wi-Fi traffic, even without




Figure 5: Overview of BeamCraft.

the knowledge of beamforming algorithms adopted by an AP. However, the remaining challenge is how to carry out the purposeful attacks of traffic disruption and plunder. To tackle this challenge, our methods to forge feedback matrices in Sections 3.3 and 3.4 exploit the shared principle across almost all beamforming algorithms, i.e., enhancing signal power while suppressing interference.

3.2 Location and Power Selection

We address the first challenge in this section, i.e., to maintain the covertness of the attack towards the victim while ensuring successful reception of the forged BFI at the Wi-Fi AP. To this end, the forged packet should be correctly decoded by the Wi-Fi AP but not by the victim. The successful decoding mainly depends on whether the receiver's SNR exceeds the decoding threshold. Specifically, the threshold of decoding increases with the Modulation Coding Scheme (MCS) index that represents the selected modulation scheme and coding scheme. A higher MCS index, indicating more sophisticated modulation types and higher coding rates, makes the system more susceptible to noise, thereby raising the SNR threshold. Therefore, the attacker should leverage a high MCS index to shorten the duration of transmitting the forged BFI, thereby reducing the impact caused by the BFI injection on normal communications. Setting the MCS index to k , the necessary SNR threshold for correctly decoding is denoted by θ_k^{SNR} , and the SNR at the Wi-Fi AP should be higher than $\theta_k^{\text{SNR}} - \delta^{\text{SNR}}/2$, where $\delta^{\text{SNR}} > 0$ is the SNR margin.⁴ To further ensure the covertness of the attacker, the received SNR of the forged BFI at the victim should be lower than $\theta_k^{\text{SNR}} - \delta^{\text{SNR}}/2$.

Now, we focus on the received SNR that is determined by three elements, i.e., the transmit power, the propagation attenuation, and the noise power. To manipulate the SNRs

at both the victim and Wi-Fi AP, we aim to: i) manage propagation attenuation to create the SNR margin, and ii) adjust transmit power to meet the SNR threshold. Propagation attenuation ℓ (in dB) depends on distance d (in meters) as

$$\ell = \alpha \times 10 \log_{10}(d) + \ell_1, \quad (2)$$

where ℓ_1 denotes the propagation attenuation when d being 1 m and α is the path loss exponent [40]. Let d^{ap} denote the distance between the attacker and the Wi-Fi AP and d^v denote the distance between the attacker and the victim. Then, to ensure the SNR margin, d^{ap} and d^v should satisfy

$$d^v \geq 10^{\delta^{\text{SNR}}/(10\alpha)} d^{\text{ap}}. \quad (3)$$

Typically, $\alpha \in [2, 4]$ and here we consider the worst case (i.e., $\alpha = 2$) [40]. Based on Eqn. (3), we can plot the feasible locations for the attacker under different δ^{SNR} requirements when the distance between the victim and the Wi-Fi AP is set as 10 m, as shown in Figure 6. When δ^{SNR} is 3 dB, the boundary of the feasible region is at most 24 m away from the Wi-Fi AP. Note that physical obstructions (e.g., wall and




Figure 6: The feasible regions and their corresponding boundaries given different values of δ^{SNR} . The distance requirement for decoding the genuine feedback transmitted by the victim is virtually not an limiting factor, since its low MCS index enables the attacker to decode victim's feedback from far away.

⁴The table of MCS index and related SNR threshold can be found in Wireless LAN Professionals' site (<https://wlanprofessionals.com/revolution-wifi-mcs-to-snr-levels/>).

glass windows) negligibly affect the feasible area as they similarly attenuate signals from both the victim and Wi-Fi AP. In addition, we need to adjust the transmit power for ensuring the SNR at the victim below $\theta_k^{\text{SNR}} - \delta^{\text{SNR}}/2$ and that at the Wi-Fi AP above $\theta_k^{\text{SNR}} + \delta^{\text{SNR}}/2$. Given the SNR threshold, propagation attenuation, and noise power, the transmit power can be uniquely determined.

In practice, an attacker may have no prior knowledge of the parameters α and ℓ_1 in Eqn. (2), and thus s /he can vary her/his location and then estimate those parameters first. After that, the attacker can first select the location by considering both the physical layout of the scenario and the feasible location indicated by Eqn. (3) with a required δ^{SNR} . By sniffing packets from both the victim and the Wi-Fi AP, the attacker can then measure the received power levels with popular sniff tools: for example, Wireshark [9] and Aircrack-ng [1] allow for an accurate assessment of the SNR margin based on the principle of channel reciprocity [51]. If the measured SNR margin is adequate (i.e., higher or equal to δ^{SNR}), the attacker can then focus on selecting an appropriate transmit power level. Otherwise, the attacker turns back to select the location again within the feasible region defined by a slightly higher SNR margin (e.g., $\delta^{\text{SNR}}+1$ dB): as the required SNR margin δ^{SNR} remains intact, selecting the location within a more “conservative” region should have a better chance to meet the requirement. To determine the transmit power, the attacker needs to again rely on Aircrack-ng for acquiring the noise power.

3.3 Traffic Disruption

In this subsection, we aim to forge the feedback matrix to realize the traffic disruption. Specifically, there are two modes of beamforming in Wi-Fi systems, that is single-user multiple-in multiple-out (SU-MIMO) and multi-user multiple-in multiple-out (MU-MIMO). The former demands that the Wi-Fi AP transmits signals towards one user at a time and improves the SNR via beamforming, whereas the latter enables the AP to serve multiple users at a time and uses beamforming to separate the signals towards different users. Two modes are easily distinguished by observing how many users are reporting BFI: only one user reports the BFI in each channel sounding of SU-MIMO mode, as opposed to multiple users reporting in that of MU-MIMO mode. Here, we first study the SU-MIMO mode and then consider the MU-MIMO.

In the SU-MIMO mode with one active user, after receiving the feedback matrix \mathbf{V} in the BFI carried by the channel sounding protocol, the Wi-Fi AP first precodes transmission data $\mathbf{s} = [s_m] \in \mathbb{C}^{M^R \times 1}$ using beamforming matrix $\mathbf{C} \in \mathbb{C}^{M^T \times M^R}$ and then transmit $\mathbf{C}\mathbf{s}$ towards the active user.⁵ After

⁵We ignore the subscript i in Eqn. (1) for the SU-MIMO mode, since there is only one active user at any given point in time.

undergoing the wireless channel $\mathbf{H} \in \mathbb{C}^{M^R \times M^T}$, the received signal at the user can be expressed as

$$\mathbf{y} = \mathbf{H}\mathbf{C}\mathbf{s} + \mathbf{n} = \mathbf{U}\mathbf{\Lambda}\mathbf{V}^H\mathbf{C}\mathbf{s} + \mathbf{n}, \quad (4)$$

where $\mathbf{n} \in \mathbb{C}^{M^R \times 1}$ is the Gaussian noise and $\mathbf{U}\mathbf{\Lambda}\mathbf{V}^H$ is the SVD of \mathbf{H} introduced in Section 2.1. The above equation indicates, under beamforming, the equivalent channel is $\mathbf{U}\mathbf{\Lambda}\mathbf{V}^H\mathbf{C}$ from the perspective of the user. As we mentioned before, the beamforming matrix \mathbf{C} is determined by the beamforming algorithm hidden for an attacker. Fortunately, as there is no inter-user interference in the SU-MIMO mode, all beamforming algorithms aim to achieve one purpose: maximizing the received signal power. Therefore, they all directly take $\mathbf{C} = \mathbf{V}$ because \mathbf{V} maximizes the product of $\mathbf{V}^H\mathbf{C}$ (hence the power). Here, we ignore the average SNR information in the BFI as our study in Section 2.3 indicates forging average SNR hardly impacts the traffic.

Towards the goal of traffic disruption, an attacker injects a forged feedback matrix, denoted by \mathbf{V}^f , to reduce the equivalent channel gain of $\mathbf{U}\mathbf{\Lambda}\mathbf{V}^H\mathbf{V}^f$. It seems that the most direct method is to let \mathbf{V}^f be a zero matrix $\mathbf{0}$. However, this method does not work with commodity Wi-Fi systems, since the feedback matrix has been compressed into phases and the matrix would not be zero even if all phases are set to zeros. Therefore, we turn to force the equivalent channel $\mathbf{U}\mathbf{\Lambda}\mathbf{V}^H\mathbf{V}^f$ to approach $\mathbf{0}$, i.e., $\min \|\mathbf{U}\mathbf{\Lambda}\mathbf{V}^H\mathbf{V}^f\|$. As \mathbf{U} is unitary, $\min \|\mathbf{U}\mathbf{\Lambda}\mathbf{V}^H\mathbf{V}^f\|$ is equal to $\min \|\mathbf{\Lambda}\mathbf{V}^H\mathbf{V}^f\|$. Recall that $\mathbf{\Lambda}$ is a diagonal matrix with nonnegative real values on the diagonal. Then, the optimal forged matrix to the above problem, denoted by $\mathbf{V}^{f,\star}$, should lie in the nullspace of genuine feedback matrix \mathbf{V} , i.e., $\mathbf{V}^H\mathbf{V}^{f,\star} = \mathbf{0}$.⁶ To calculate $\mathbf{V}^{f,\star}$, we can use Gram–Schmidt orthonormalization method [50]. With $\mathbf{V}^{f,\star}$ contained in the injected BFI, the AP’s beamforming direction becomes orthogonal to the channel direction, causing the received signal power at the user (victim) being minimal and leading to the throughput being almost zero.

It is now clear that the correlation between the beamforming direction and channel direction affects the throughput: a correlation 1 achieved by \mathbf{V} leads to the highest throughput, while $\mathbf{V}^{f,\star}$ yields a correlation 0 and hence a throughput of almost zero. Intuitively, we should be able to manipulate the traffic at will by adjusting the correlation between \mathbf{V}^f and \mathbf{V} , denoted by $\rho^f \in [0, 1]$. Such a flexible traffic disruption attack is feasible if a forged feedback matrix satisfies two requirements: i) the correlation being ρ^f , i.e., $\mathbf{V}^H\mathbf{V}^f = \rho^f\mathbf{I}$, and ii) being an orthonormal matrix required by the BFI compression method, i.e., $(\mathbf{V}^f)^H\mathbf{V}^f = \mathbf{I}$. To this end, we construct the forged feedback matrix as $\mathbf{V}^f = \sqrt{1 - (\rho^f)^2}\mathbf{V}^{f,\star} + \rho^f\mathbf{V}$. Consequently, the resulting channel becomes $\rho^f\mathbf{U}\mathbf{\Lambda}\mathbf{V}^H\mathbf{V}$ and the

⁶The nullspace of \mathbf{V} is not empty, since the number of user’s antennas (usually being 1 or 2) is less than that of the Wi-Fi AP (usually being 4).

received signal power is $(\rho^f)^2$ times that under the normal beamforming without attack. Based on the above analysis, the attacker can adjust ρ^f contained in V^f to manipulate the throughput of the victim within the range from the normal level without attack (enabled by setting $\rho^f = 1$) to almost zero (setting $\rho^f = 0$).

As the aforementioned attack works for the SU-MIMO mode, we need to further extend the attack to the MU-MIMO mode. The MU-MIMO allows for simultaneous transmitting different data streams to distinct users, and the underlying “separation” is enabled via channel directions specified by the feedback matrices. Taking advantage of this feature, the attacker can inject the feedback matrix of the victim copied from another active user, and then the received signal at the victim would be a mixture of two different signals with equal power for two users (victim and another user), since they appear to be located at the same direction from the perspective of the Wi-Fi AP. Consequently, the signal-to-interference-plus-noise ratio (SINR) at the victim is reduced to around 0 dB, leading to a high bit error rate and interrupting the communication.

3.4 Traffic Plunder

To realize the attack of traffic plunder, we focus on the MU-MIMO mode that allows for simultaneous transmitting to multiple users over shared bandwidth, since it makes sense to “plunder” victims’ throughput to the benefit of the attacker under such circumstances. Let us first consider a simple case where there are two single-antenna users in the MU-MIMO mode, the first one being the beneficiary (and the attacker too) and the second one being the victim, with the attacker aiming to increase its own throughput by sacrificing that of the victim. Further extensions to more general cases will be presented based on this simple one.

With each user having only one antenna, the AP transmitted data for i -th user can be simplified to a scalar s_i , and the corresponding beamforming matrix is $\mathbf{c}_i \in \mathbb{C}^{M \times 1}$. Therefore, the total transmit signal at the AP is $\mathbf{c}_1 s_1 + \mathbf{c}_2 s_2$, since the data for the two users are transmitted simultaneously. After undergoing the wireless channel $\mathbf{h}_i \in \mathbb{C}^{1 \times M}$ from the AP to the i -th user, the received signal at i -th user can be expressed as

$$y_i = \mathbf{h}_i(\mathbf{c}_1 s_1 + \mathbf{c}_2 s_2) + n_i. \quad (5)$$

The resulting SINR of i -th user can be calculated as

$$\eta_i = \frac{|\mathbf{h}_i \mathbf{c}_i|^2 p_i^T}{|\mathbf{h}_i \mathbf{c}_{3-i}|^2 p_{3-i}^T + \sigma^2}, \quad (6)$$

where p_i^T is the AP’s transmit power to the i -th user. To improve the throughput, there are only two feasible ways: increasing signal power $|\mathbf{h}_i \mathbf{c}_i|^2$ and suppressing interference $|\mathbf{h}_i \mathbf{c}_{3-i}|^2$. All MU-MIMO beamforming algorithms seek the

balance between two ways. As BeamCraft aims to manipulate traffic under all beamforming algorithms, we need to forge a feedback matrix of the victim that both increases signal power and suppresses interference for the beneficiary. To this end, we aim to design the forged feedback matrix that can work with two typical beamforming algorithms, that is, direct beamforming and ZF beamforming. The former directly takes the feedback matrix as the beamforming matrix for maximizing the signal power ignoring the interference while the latter forces the inter-user interference to be zero.

For the direct beamforming algorithm, $\mathbf{C} = [\mathbf{v}_1, \mathbf{v}_2]$ where the feedback matrix \mathbf{v}_i can be simplified to $\mathbf{h}_i / \|\mathbf{h}_i\|$ for single-antenna users, and the SINR at the beneficiary becomes:

$$\eta_1^d = \frac{\|\mathbf{h}_1\|^2 p_1^T}{\|\mathbf{h}_1\|^2 |\rho^M|^2 p_2^T + \sigma^2}, \quad (7)$$

where $\rho^M = \mathbf{v}_1^H \mathbf{v}_2$ describes the correlation between the channel of two users. For the ZF beamforming, \mathbf{C} can be expressed as $(1 - |\rho^M|^2)^{-\frac{1}{2}} [\mathbf{v}_1 - (\rho^M)^* \mathbf{v}_2, \mathbf{v}_2 - \rho^M \mathbf{v}_1]$ where $(\cdot)^*$ denotes the conjugate operation [65], and the SINR at the beneficiary becomes

$$\eta_1^{ZF} = \frac{\|\mathbf{h}_1\|^2 (1 - |\rho^M|^2) p_1^T}{\sigma^2}. \quad (8)$$

By observing the above two SINR expressions, one can clearly find that the correlation ρ^M influences the SINR under two typical algorithms. Since all beamforming algorithms aim to balance between increasing signal power (i.e., direct beamforming) and suppressing interference (i.e., ZF beamforming), reducing $|\rho^M|$ can improve the SINR under all types of beamforming algorithms. Therefore, an attacker can forge the feedback matrix of the victim (denoted by \mathbf{v}_2^f) being orthogonal to that of the beneficiary, i.e., $(\mathbf{v}_1)^H \mathbf{v}_2^f = 0$; the forged feedback matrix can again be obtained via Gram–Schmidt orthonormalization method. With the forged \mathbf{v}_2^f , all beamforming algorithms apply the same beamforming matrix, i.e., $\mathbf{c}_1 = \mathbf{v}_1$ and $\mathbf{c}_2 = \mathbf{v}_2^f$, since there is no inter-user interference, and the SINR of the beneficiary can be boosted to $\eta_1^f = \|\mathbf{h}_1\|^2 p_1^T / \sigma^2$ while that of the victim is reduced to $\eta_2^f = \|\mathbf{h}_2\|^2 (1 - |\rho^M|^2) p_2^T / (\|\mathbf{h}_2\|^2 |\rho^M|^2 p_1^T + \sigma^2)$.

Based on the simple two-user single-antenna case, we can proceed to consider general cases with more than two users and each user equipped with multiple antennas. The intuitive attack is to forge all BFI except the beneficiary, however, this would increase the possibility of exposure as the users are randomly distributed in the scenario. Alternatively, we can sniff BFI from all users and calculate the correlation of the feedback matrix between the beneficiary and any other user. Indicated by comparing η_1^f to η_1^d and η_1^{ZF} , the profit of forging increases with $|\rho^M|$. Therefore, we prioritize the user with the highest correlation as the victim, and determine the attacker’s location and transmit power accordingly. The

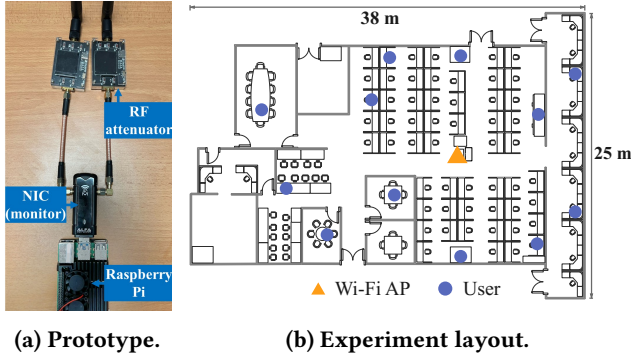


Figure 7: BeamCraft implementation: (a) hardware configurations and (b) experiment layout in an office area.

remaining users may be selected as victims if the forged packet cannot be correctly decoded by them either. Besides, when the users are equipped with multiple antennas, the correlation is a matrix rather than a scalar, so we need to calculate the norm of the correlation matrix and select the highest one. Because this selection procedure effectively brings any complicated cases back to the simple one analyzed earlier, the proposed attack procedure can still be applied, including the Gram–Schmidt orthonormalization method, to derive the forged feedback matrix.

4 IMPLEMENTATION AND SETUP

In this section, we elaborate on BeamCraft’s implementation, as well as introduce the experiment setup.

Implementation. We use a Raspberry Pi 4 [41] with 4GB RAM, one Wi-Fi NIC (Alfa AWUS036ACM [3]), and two radio frequency (RF) digital step attenuators (ranging from 0dB to 31.75dB with the step being 0.25dB) to act as an attacker, as shown in Figure 7a. The Wi-Fi NIC works in *monitor* mode and is used for sniffing and injecting packets. Due to our inability to control the transmit power via NIC driver, the two RF digital step attenuators are used for adjusting the transmit power. Note that the attenuation setting latency (microsecond level [4]) of RF digital step attenuator is much smaller than the time interval of channel sounding (millisecond level) so that the attenuation can be controlled in a timely manner to satisfy the need of the attacker. The Raspberry Pi 4 is used for executing the joint location and transmit power strategy and the proposed BFI forgery method. Moreover, the beneficiary could be the Raspberry Pi or any user who cooperates with the attacker.

We implement the software framework of BeamCraft in low level C++ language running on the Raspberry Pi 4. Following the workflow introduced in Figure 5, the attacker first utilizes libpcap 1.10.3 [47] to capture packets from the Wi-Fi AP and other users with the attenuation of the attenuator being 0dB. After measuring path attenuation, the proposed

joint location and transmit power strategy is used for determining the location of the attacker and transmit power of the forged BFI packet. Here, we set the SNR margin to 3dB since the SNR threshold for two consecutive MCS indices differs by 2-4dB. With these preparations, the attacker starts sniffing the BFI from the users and the captured BFI is decompressed into the feedback matrix following the Wi-Fi standards [7]. With the proposed BFI forgery method, the forged feedback matrix is constructed using Eigen 3.3.7 [21] and then is compressed into a BFI packet. Finally, the forged BFI is injected using libpcap with the specific transmit power once the attacker detects the NDP announcement.

Experiment Setup. We conduct experiments in an office area, as shown in Figure 7b. The Wi-Fi AP located in the center of the area is Xiaomi Redmi Router AC2100 [59] with four antennas and operates at 5GHz with 20MHz bandwidth under the Wi-Fi 5 standard. The users randomly distributed in the area consist of six types: two smartphones (iPhone 15 [5] with two antennas and Xiaomi 13 Pro [58] with two antennas) and four Wi-Fi NICs for the laptop (Intel 8265 [26] with two antennas, MediaTek MT7921 [32] with two antennas, Realtek RTL8821CU [43] with one antenna, and Realtek RTL8812BU [42] with two antennas). Iperf3 [48] is used for measuring the throughput between the Wi-Fi AP and users. To show the average performance, the number of user locations is 12 and the test time of each location is around 15 minutes. The total test time is more than 30 hours.

Metrics. We adopt two metrics, namely reduction rate and increase rate, to quantify attack performance. The former is used for describing the performance of the traffic disruption and is defined as the reduction percentage of the victim’s throughput, i.e., $(\tau^N - \tau^A)/\tau^N$, where τ^N and τ^A denote the throughput without/with attacks, respectively. The latter is used for describing the performance of the traffic plunder and is defined as the increasing percentage of the beneficiary’s throughput, i.e., $(\tau^A - \tau^N)/\tau^N$.

5 EVALUATION

We start with two micro-benchmark studies to demonstrate the effectiveness of BeamCraft and then present the overall performance under traffic plunder and traffic disruption.

5.1 Micro-benchmark Studies

5.1.1 Low Exposure Rate. To validate our joint location and power selection strategy for the attacker (detailed in Section 3.2), we begin by assessing the impact of transmit power at a predetermined location. Specifically, we arrange the attacker, the Wi-Fi AP, and the victim in a linear configuration, maintaining a 5m distance between the Wi-Fi AP and each of the other two parties. Figure 8a shows the decoding probabilities of forged BFI at the Wi-Fi AP and the victim




Figure 8: Probabilities of correctly decoding at the Wi-Fi AP and victim, with different (a) transmit power and (b) distance between the attacker and the AP.

across various transmit power and MCS indexes. Following the proposed strategy, the selected transmit power is 9dBm when the SNR margin δ^{SNR} is 3dB and the MCS index is 6. This setup results in a 0% decoding probability at the victim and a 100% probability at the Wi-Fi AP, as depicted in Figure 8a, thereby confirming the strategy's effectiveness. The difference between the decoding probability for victim and AP can be attributed to distinct propagation attenuation. Additionally, a higher MCS index, correlating with an increased SNR requirement, further reduces decoding probability. One may expect to detect the forged BFI via its high MCS index, but the AP cannot make such an identification as it may receive BFIs modulated with all MCS indices, while the victim cannot decode the forged BFI given our location and power selection strategy.

Maintaining the distance between the victim and the Wi-Fi AP, we also explore the feasible region of the attacker by controlling the distance between the attacker and Wi-Fi AP, as shown in Figure 8b. Using the proposed strategy, the transmit power is dynamically adjusted to ensure that the probability at the victim remains at zero. Remarkably, the probability at the AP stays at 100% even when the distance between the attacker and the AP is 14m, 2.8 times of that between the victim and the AP. This confirms the extensive feasible region and zero exposure risk achieved by our proposed location and power selection strategy.

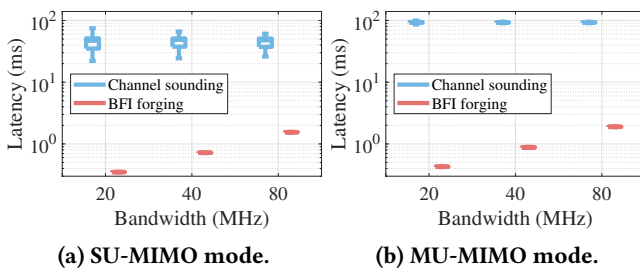


Figure 9: Comparison between the computational latency and the channel sounding interval in (a) SU-MIMO mode and (b) MU-MIMO mode.

5.1.2 Real-time Capability. To avoid the adverse effect of channel variation on the effectiveness of the forged feedback matrix, the injected BFI is forged based on the sniffed BFI in the last channel sounding. This approach necessitates that the total computational latency, including the stages of decompression, feedback matrix forgery, and compression, should be shorter than the channel sounding interval. To this end, Figure 9 presents the latency under SU-MIMO and MU-MIMO modes with different bandwidths for a two-antenna user. One can clearly observe that even under the most demanding conditions (80MHz bandwidth in MU-MIMO mode), the average computational latency is merely 1.9ms. Though the overall latency may grow to around 2.5ms on average by further considering the overhead of sniffing and forgery, this value is still significantly below the minimum channel sounding interval of 43.7ms in all tested scenarios. This demonstrates BeamCraft's real-time operational feasibility and efficiency. Looking forward, the adoption of parallel processing techniques promises to further decrease latency, potentially enabling BeamCraft to effectively work in the upcoming Wi-Fi 7 hardware with up to 320MHz bandwidth.

5.2 Performance of Traffic Disruption

To show the overall performance of traffic disruption, we conducted experiments involving six types of users as victims. Figure 10 presents the normal throughput without attack and the throughput with attacks using the optimal forged feedback matrix and the random matrix. The throughput drops significantly after injecting the optimal feedback matrix. Despite inherent differences in throughput performance among users, attributed to the varied software and hardware configurations of NICs by different manufacturers, the average throughput experiences a substantial decrease to 21.3Mbps from 112.9Mbps for users with two antennas, and to 12.3Mbps from 60.5Mbps for user with one antenna. The average reduction rate achieved by the optimal forged feedback matrix is about 80.6%, seriously hindering normal Wi-Fi communications. This demonstrates that our BFI forgery method is universally effective, not limited to a specific user. Additionally, injecting the random feedback matrix also diminishes throughput by an average rate of 15.8% and the decline is considerably less severe than that with the optimal




Figure 10: Overall performance of traffic disruption.

forged feedback matrix, confirming the superiority of our BFI forgery method. To delve deeper into the performance of the traffic disruption attack, we study the impact of practical factors using RTL8812BU as the representative.

Impact of Bandwidth. We evaluate the impact of the Wi-Fi system's bandwidth, and Figure 11a shows the throughput with and without attacks across various bandwidths. Although the normal throughput increases with the bandwidth, the reduction rate caused by the optimal forged feedback matrix remains at the same level of roughly 80.4%, significantly outperforming that with the random feedback matrix, approximately 21.1%. Note that the cost of BeamCraft still remains low irrespective of bandwidth increases, only needing to inject forged BFI 10~30 times per second. This is attributed to the nearly constant channel sounding interval under different bandwidths as shown in Figure 9. These findings indicate that BeamCraft can effectively attack high-bandwidth Wi-Fi systems at a sufficiently low cost.

Impact of Distance. We evaluate the impact of the distance between the victim and the Wi-Fi AP. As shown in Figure 11b, the throughput with/without the attack decreases with the distance due to the increase of the propagation attenuation. Nevertheless, the reduction rate increases with the distance, being 83.3%, 86.2%, and 88.1% for distances of 10 m, 20 m, and 30 m, respectively. This result arises because the forged BFI directly reduces the received signal power at the victim, and the relationship between throughput and signal power follows a logarithmic function. Therefore, the throughput reduction becomes more pronounced at a longer distance, where signal power is naturally lower.




Figure 11: Impact of practical factors on traffic disruption: (a) AP's bandwidth, (b) distance between the victim and AP, (c) mobility of the victim, and (d) ρ^f for manipulating the traffic at will.

Impact of Mobility. In practice, the user may move. To evaluate the impact of the mobility, the victim moves on a prescribed route, maintaining an average distance of 20m from the Wi-Fi AP. Figure 11c illustrates throughput variations over time with and without the attack. User's mobility naturally decreases the accuracy of the beamforming direction, reducing the average normal throughput to 75.3Mbps (from 98.6Mbps in static state). Under the attack, the throughput further declines to 22.4Mbps, affirming the traffic disruption attack's effectiveness even amidst user movement.

Impact of ρ^f . In Section 3.3, we propose to manipulate the traffic at will by controlling ρ^f . The efficacy of adjusting ρ^f is demonstrated in Figure 11d, where one can clearly observe that the reduction rate increases with ρ^f . The relationship between them follows a monotone convex function. It is because the received signal power at the victim quadratically decreases with ρ^f and the throughput logarithmically increases with the power. Consequently, the throughput logarithmically decreases with $(\rho^f)^2$, leading to the monotone convex function between the reduction rate and ρ^f . The result verifies that the attacker can precisely control victim's traffic, demonstrating the flexibility of BeamCraft.

5.3 Performance of Traffic Plunder

In this section, we conduct experiments to evaluate the performance of traffic plunder. Two users equipped with two-antenna RTL8812BU simultaneously connect to the Wi-Fi AP and request downlink traffic. One user acts as the attacker (also the beneficiary) and executes the traffic plunder attack, boosting its throughput at the expense of the other user (the victim). Figure 12a shows the performance of the traffic plunder attack. With the attack, the beneficiary's throughput increases to 87.2Mbps from 72.9Mbps, an increase rate of 19.6%, while the victim's throughput drops to 43.1Mbps from 71.9Mbps, a reduction rate of 40.1%. This result verifies the effectiveness of the traffic plunder. Additionally, using a random feedback matrix significantly reduces the throughput for both the beneficiary and the victim due to increased interference and decreased received signal power, resulting in lower SINRs at both users. Furthermore, we examine the




Figure 12: Throughput with/without traffic plunder for users with (a) two antennas and (b) only one antenna.




Figure 13: Impact of practical factors on traffic plunder.

attack’s impact on two single-antenna users (RTL8821CU), as shown in Figure 12b. The results closely resemble those from the two-antenna scenarios, though with a milder impact: the beneficiary’s throughput increases by 15.1%, and the victim’s decreases by 20.9%. Given that the rationale behind our method is to eliminate the interference at the beneficiary, the lower inter-user interference among single-antenna users can explain this difference. The results demonstrate our proposed attack’s effectiveness across different user configurations. In the following, we study the impact of practical factors using RTL8812BU as the representative.

Impact of Bandwidth. We examine how the performance of the traffic plunder varies with the bandwidth of the Wi-Fi AP. Figure 13a shows the reduction rate of the victim and the increase rate of the beneficiary. Both of them remain almost unchanged under different bandwidths, around 20.5% and 37.9%, respectively, since the data transmission over the whole bandwidth is managed by the beamforming and the forged BFI misleads the beamforming direction of each subcarrier uniformly. Meanwhile, as shown in Figure 9b, the attack towards MU-MIMO mode requires only 10~20 forged BFI injections per second, indicating that BeamCraft can realize the attack with low cost regardless of bandwidth.

Impact of Correlation. We also evaluate the impact of the correlation $|\rho^M|$ between the feedback matrices of two users. To achieve different correlations, we vary the angle formed by the AP-victim and AP-beneficiary lines and fix the distance between the Wi-Fi AP and the victim/beneficiary. Figure 13b indicates that both the beneficiary’s increase rate and the victim’s reduction rate increase with the correlation. The result accords with the analysis in Section 3.4 that higher correlation leads to higher interference and the attack can yield greater benefits since it eliminates the interference at the beneficiary.

5.4 Real-World Experiment

To demonstrate BeamCraft’s real-world applicability, we conduct an experiment involving 10 users (covering all types mentioned in Section 4) randomly distributed within an office area and concurrently connected to a single Wi-Fi AP. To further verify the wide applicability of BeamCraft, we test it with another two representative Wi-Fi APs.




Figure 14: Performance of BeamCraft in a real-world scenario: (a) traffic disruption and (b) traffic plunder.

Traffic Disruption. Figure 14a illustrates the throughput for four representative users, with the rest exhibiting similar patterns. Due to compatibility issues with MU-MIMO mode among some devices, the Wi-Fi AP switches between SU-MIMO and MU-MIMO modes to accommodate all users. Since multiple users share the bandwidth, the average throughput over users is 28.8Mbps. Under the attack, the throughput experiences a significant reduction to 7.5Mbps, severely impacting the viability of normal communication services. Specifically, we assess the performance of the attack on video streaming and online gaming when the Wi-Fi AP operates at 5GHz with 40MHz bandwidth, as shown in Table 1. The resolution and bitrate of the testing video are 3840×2160 pixels and 51 Mbps, respectively. For video streaming, its quality of experience (QoE) is represented by the fluency of the video, i.e., rebuffering percentage [37] that is the percentage of the total streaming time spent rebuffering. For online gaming, we use the latency between the client and the game server. It can be observed that the QoE of the two tasks drops significantly under the attack, further confirming the effectiveness of BeamCraft.

Table 1: Performance on two communication tasks.

Task	Video	Game
Metric	Rebuffering percentage	Latency
W/o attack	0.0%	13.7 ms
W/ attack	78.0%	55.1 ms

Traffic Plunder. Since the Wi-Fi AP transmits data up to four users simultaneously in the MU-MIMO mode, user 1 (RTL8821CU) is selected as the beneficiary and the other three users are regarded as potential victims. For the case with more than two users, we propose to select the user with the highest correlation as the victim in Section 3.4. To verify it, we present the throughput of the beneficiary in three cases: i) without attack, ii) with only one victim, and iii) with three victims together, as shown in Figure 14b. The correlation between the feedback matrices of the three potential victims (users 2, 3, and 4) and that of the beneficiary is 0.43, 0.52, and 0.58, respectively. From Figure 14b, compared to selecting user 2 or 3, selecting user 4 as the victim yields higher profit

Table 2: Traffic disruption on different APs.

AP	TP-Link XDR3020	Xiaomi AX3000T
W/o attack	135.3Mbps	132.6Mbps
W/ attack	23.2Mbps	21.8Mbps
Reduction rate	-82.6%	-83.6%

Table 3: Traffic plunder on different APs.

AP	TP-Link XDR3020	Xiaomi AX3000T
W/o attack	76.4Mbps	75.7Mbps
W/ attack	89.6Mbps	87.9Mbps
Increase rate	+16.9%	+18.3%

and the throughput benefit from the attack increases with the correlation, verifying the effectiveness of the victim selection policy. Moreover, simultaneously selecting users 2, 3, and 4 as victims surpasses the outcome of only one victim.

Different APs. To further verify the performance of BeamCraft, we conduct experiments with another two different Wi-Fi APs, i.e., TP-link Router TL-XDR3020 [49] and Xiaomi Router AX3000T [60]. Tables 2 and 3 illustrate the performance of traffic disruption and traffic plunder, respectively. Due to the page limit, we only show the performance of one user (RTL8812BU), and other users show similar performance. For the two different Wi-Fi APs, the average reduction rate of the victim’s throughput achieved by the traffic disruption is 82.6% and 83.6%. With the traffic plunder, the average increase rate of the beneficiary’s throughput is 16.9% and 18.3%, respectively. Both of them are similar to the results in Sections 5.2 and 5.3. These experiments evidently prove that BeamCraft can apply to different Wi-Fi APs.

6 LIMITATIONS AND DEFENSE

In this section, we first discuss the limitations of BeamCraft with extended experiments, then present the generalizability of BeamCraft towards Wi-Fi 6 systems. Finally, defense strategies against BeamCraft are proposed.

6.1 Discussions with Extended Experiments

In our experiments, we observe the instability of the MU-MIMO mode, as shown in Figure 15a. The Wi-Fi AP may exit the MU-MIMO mode and shift to SU-MIMO due to the channel variation triggered by movements, indicated by the correlation between feedback matrices in two consecutive BFI packets. This phenomenon also occurs after launching the attack for a while and we guess the reason is the “retry” indication in the forged beamforming feedback packet. As the AP continuously receives these packets, it interprets the situation as ongoing significant channel variations, prompting a switch to SU-MIMO mode to avoid serious inter-user interference. To measure the effect caused solely by the attack, we

modify the Realtek drivers for RTL8821CU and RTL8812BU, enabling detailed packet information and throughput measurement exclusive to MU-MIMO mode. Consequently, the experiment results of traffic plunder focus on RTL8821CU and RTL8812BU. To prevent unintended mode switches, one possible way is to add the forecast for the NDP announcement so that the forged BFI can arrive at the Wi-Fi AP before the genuine one without the “retry” indication, which we leave as a future exploration.

Our experiments focus on the Wi-Fi 5 standard that is the most widely adopted one. In the up-to-date Wi-Fi 6 standard, orthogonal frequency-division multiple access (OFDMA) technique is applied. Each active user only occupies a part of the whole bandwidth and beamforming is used for the data transmission over the allocated bandwidth. Consequently, BeamCraft should still work with Wi-Fi 6 systems. Figure 15b illustrates the throughput reduction caused by the traffic plunder attack in a Wi-Fi 6 system consisting of Xiaomi Router AX3200 [61] and MT7921, verifying the generalizability of BeamCraft.

6.2 Defense and Security Analysis

We hereby present three defense strategies.

Encryption. The first strategy involves encrypting the BFI to prevent forged BFI packets from bypassing security check at the Wi-Fi AP. Despite its potential for enhancing security, this method incurs overheads due to the necessity for frequent key exchanges and the encryption/decryption process. Existing work on the overhead of data frames encryption [36] has proved that it severely reduces throughput and introduces high latency. Such negative impacts may also be present when encrypting the BFI, making it less suitable for environments with high user dynamics, such as shopping malls and cafeterias.

Channel Verification. The second strategy leverages channel reciprocity for security verification [56]. The Wi-Fi AP can measure the CSI between the user and itself from the pilot contained in the BFI packet and the measured CSI is




Figure 15: Extended experiments: (a) instability of the MU-MIMO mode and (b) traffic disruption with the Wi-Fi 6 system.

ought to match the received compressed CSI. This countermeasure should work well for most Wi-Fi NICs, but attackers equipped with SDR could modify the pilot signals within BFI packets, thereby passing the security verification.

Pilot Randomization. The third strategy makes use of random pilots within NDP in each channel sounding and the user feeds received signals back to the AP, making it challenging for attackers to sniff the CSI [52]. This approach potentially disrupts the traffic plunder attack by obscuring real CSI, although it might not fully prevent traffic disruption attacks since random feedback can still alter the beamforming direction inaccurately.

Remarks. Given the limitations and benefits of these strategies, we opt to combine channel verification with random pilots. This hybrid strategy aims to obscure the channel information from attackers and strengthen the verification of BFI packets, enhancing the Wi-Fi network's resilience to both plunder and disruption attacks without compromising user experience or network performance.

7 RELATED WORKS

Our work is related to traffic attacks towards Wi-Fi systems. As a network type, Wi-Fi systems face threats from traditional higher-layer network attacks including denial-of-service (DoS) attacks [11, 53, 62]. Bogdanoski *et al.* [11] focus on the SYN (synchronize) flood attack and analyze its impact on the service quality of voice, video, and data in Wi-Fi systems. Tushir *et al.* [62] launch cascade DoS attacks in a multi-hop Wi-Fi network and congest the entire network by sending a few malicious packets. However, these network-layer attacks can be easily detected and defended by traffic filtering strategies. Meanwhile, Wi-Fi systems, relying on wireless radio signals for transmission, are also vulnerable to various jamming attacks, including generic jamming attacks [8, 44, 54, 63] and Wi-Fi-specific jamming attacks [15, 27, 45, 68]. The former is to overwhelm wireless signals by injecting high-power random interference signals and can be used in other types of wireless systems. The latter is dedicated to the signal processing pipeline and media access protocols of the Wi-Fi system, such as the channel estimation [15], frequency orthogonality for OFDM [68], and rate adaptation algorithms [35], thus being more efficient and energy-saving. However, these attack methods generally require strict clock synchronization between the attacker and the victim, so they all stay on emulations with limited practical significance.

Besides the method described, attacks against beamforming can also effectively disrupt traffic while being significantly more energy-efficient, requiring the forgery of only a small number of packets. Tung *et al.* [52] propose power attack that misleads power allocation at the Wi-Fi AP in the

MU-MIMO mode by reporting the falsified scale of genuine CSI. Zhang *et al.* [67] further analyze the theoretical performance of power attack for massive MIMO systems. Different from them, Hou *et al.* [23] propose to undermine the user selection of MAC layer in MU-MIMO mode, causing severe disruption; it differs from BeamCraft that targets physical layer under fixed user selection. Besides, forging CSI feedback can also be used to enable eavesdropping by misleading the beamforming direction, allowing unauthorized listening to ongoing transmissions [56, 57]; such objectives diverge from our focus on traffic manipulation. Finally, the effectiveness of these attacks is confined to SDR-driven emulations under non-standard Wi-Fi protocols and specific beamforming algorithms, resulting in limited real-world applicability.

The security aspect of Wi-Fi *integrated sensing and communication* (ISAC) [13, 22] is a marginally related topic, as it shares the same Wi-Fi technology with BeamCraft. With Wi-Fi sensing recently gaining the capability for handling multiple persons [25, 29], users are increasingly concerned about its vulnerability to exploitation. As a demonstration, WiKI-Eve [24] leverages keystroke-induced BFI variations to eavesdrop numerical passwords typed on smartphones. Meanwhile, MIMOcrypt [31] is the first attempt to thwart such attacks via physical encryption. Inspired by the BFI exploitation of WiKI-Eve, BeamCraft further suggests that BFI can be exploited to threaten Wi-Fi communications.

8 CONCLUSION

In this paper, we have explored a critical vulnerability in beamforming due to clear-text CSI feedback. We introduce BeamCraft, the first attack to manipulate traffic in commodity Wi-Fi networks by forging beamforming feedback. Our approaches include a joint location and transmit power selection strategy for the attack's covertness, alongside a novel BFI forgery method capable of conducting both traffic disruption and plunder attacks in the face of hidden beamforming algorithms. The extensive evaluations with commodity devices have evidently demonstrated BeamCraft's capability to significantly diminish the traffic of victim while simultaneously boosting the traffic of beneficiary. These findings highlight a pressing need for advanced defense strategies for Wi-Fi systems to safeguard against such risks.

ACKNOWLEDGEMENT

This research is supported by National Research Foundation (NRF) Future Communications Research & Development Programme (FCP) grant FCP-NTU-RG-2022-015, MOE Tier 1 grant RG16/22, as well as National Science Fund for Distinguished Young Scholars of China under grant No.62125203 and The Key Program of the National Natural Science Foundation of China under grant No.61932013.

REFERENCES

- [1] Aircrack-ng Team. 2024. Aircrack-ng. <https://www.aircrack-ng.org/>. Online; accessed: 16 February 2024.
- [2] Rami Akeela and Behnam Dezfouli. 2018. Software-defined Radios: Architecture, State-of-the-art, and Challenges. *Computer Communications* 128 (2018), 106–125.
- [3] ALFA Network. 2024. AWUS036ACM - Long Range Dual-Band AC1200 Wireless USB 3.0 Wi-Fi Adapter. <https://www.alfa.com.tw/products/awus036acm?variant=39477234597960>. Online; accessed: 16 February 2024.
- [4] Analog Devices. 2024. HMC1119: 0.25 dB LSB, 7-Bit, Silicon Digital Attenuator, 0.1 GHz to 6.0 GHz Data Sheet (Rev.C). <https://www.analog.com/media/en/technical-documentation/data-sheets/hmc1119.pdf>. Online; accessed: 16 February 2024.
- [5] Apple Inc. 2024. iPhone 15. <https://www.apple.com/sg/shop/buy-iphone/iphone-15>. Online; accessed: 16 February 2024.
- [6] Mohamed A Aref, Sudharman K Jayaweera, and Esteban Yezpez. 2020. Survey on Cognitive Anti-Jamming Communications. *IET Communications* 14, 18 (2020), 3110–3127.
- [7] IEEE Standards Association et al. 2016. IEEE Std 802.11-2016, IEEE Standard for Local and Metropolitan Area Networks—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications.
- [8] Suresh Bandaru. 2014. Investigating the Effect of Jamming Attacks on Wireless LANS. *International Journal of Computer Applications* 99, 14 (2014), 5–9.
- [9] Jay Beale, Angela Orebaugh, and Gilbert Ramirez. 2006. *Wireshark & Ethereal Network Protocol Analyzer Toolkit*. Elsevier.
- [10] Oscar Bejarano, Edward W Knightly, and Minyoung Park. 2013. IEEE 802.11 ac: From Channelization to Multi-User MIMO. *IEEE Communications Magazine* 51, 10 (2013), 84–90.
- [11] Mitko Bogdanoski, Tomislav Suminoski, and Aleksandar Risteski. 2013. Analysis of the SYN Flood DoS Attack. *International Journal of Computer Network and Information Security* 5, 8 (2013), 1–11.
- [12] G Charis and N Showme. 2017. Beamforming in Wireless Communication Standards: A Survey. *Indian Journal of Science and Technology* 10, 5 (2017), 1–5.
- [13] Zhe Chen, Tianyue Zheng, Chao Hu, Hangcheng Cao, Yanbing Yang, Hongbo Jiang, and Jun Luo. 2022. ISACoT: Integrating Sensing with Data Traffic for Ubiquitous IoT Devices. *IEEE Communications Magazine* 61, 5 (2022), 98–104.
- [14] Zicheng Chi, Yan Li, Xin Liu, Wei Wang, Yao Yao, Ting Zhu, and Yanchao Zhang. 2020. Countering Cross-Technology Jamming Attack. In *Proc. of the 13th ACM WiSec*. 99–110.
- [15] T Charles Clancy. 2011. Efficient OFDM Denial: Pilot Jamming and Pilot Nulling. In *Proc. of the IEEE ICC*. 1–5.
- [16] Cailian Deng, Xuming Fang, Xiao Han, Xianbin Wang, Li Yan, Rong He, Yan Long, and Yuchen Guo. 2020. IEEE 802.11 be Wi-Fi 7: New Challenges and Opportunities. *IEEE Communications Surveys & Tutorials* 22, 4 (2020), 2136–2166.
- [17] Luis Martin Garcia. 2008. Programming with Libpcap-Sniffing the Network from Our Own Application. *Hakin9-Computer Security Magazine* 2 (2008), 2008.
- [18] Matthew S Gast. 2013. *802.11 ac: A Survival Guide: Wi-Fi at Gigabit and Beyond*. O'Reilly Media, Inc.
- [19] Wallace Givens. 1958. Computation of Plain Unitary Rotations Transforming a General Matrix to Triangular Form. *J. Soc. Indust. Appl. Math.* 6, 1 (1958), 26–50.
- [20] Michelle X Gong, Brian Hart, and Shiwen Mao. 2015. Advanced Wireless LAN Technologies: IEEE 802.11 ac and Beyond. *GetMobile: Mobile Computing and Communications* 18, 4 (2015), 48–52.
- [21] Gaël Guennebaud, Benoit Jacob, et al. 2024. Eigen. https://eigen.tuxfamily.org/index.php?title=Main_Page. Online; accessed: 16 February 2024.
- [22] Yinghui He, Jianwei Liu, Mo Li, Guanding Yu, Jinsong Han, and Kui Ren. 2023. SenCom: Integrated Sensing and Communication with Practical WiFi. In *Proc. of the 29th ACM MobiCom*. 1–16.
- [23] Tao Hou, Shengping Bi, Tao Wang, Zhuo Lu, Yao Liu, Satyajayant Misra, and Yalin Sagduyu. 2022. MUSTER: Subverting User Selection in MU-MIMO Networks. In *Proc. of the 41th IEEE INFOCOM*. 140–149.
- [24] Jingyang Hu, Hongbo Wang, Tianyue Zheng, Jingzhi Hu, Zhe Chen, Hongbo Jiang, and Jun Luo. 2023. Password-Stealing without Hacking: Wi-Fi Enabled Practical Keystroke Eavesdropping. In *Proc. of the 30th ACM CCS*. 239–252.
- [25] Jingzhi Hu, Tianyue Zheng, Zhe Chen, Hongbo Wang, and Jun Luo. 2023. MUSE-Fi: Contactless MUlti-person SENSing Exploiting Near-field Wi-Fi Channel Variation. In *Proc. of the 29th ACM MobiCom*. 1–15.
- [26] Intel. 2024. Intel Dual Band Wireless-AC 8265 Specifications. <https://www.intel.com/content/www/us/en/products/sku/94150/intel-dual-band-wirelessac-8265/specifications.html>. Online; accessed: 16 February 2024.
- [27] Matthew J La Pan, T Charles Clancy, and Robert W McGwier. 2012. Jamming Attacks Against OFDM Timing Synchronization and Signal Acquisition. In *Proc. of the IEEE MILCOM*. 1–7.
- [28] Xin Li, Jingzhi Hu, and Jun Luo. 2024. Efficient Beamforming Feedback Information-Based Wi-Fi Sensing by Feature Selection. *IEEE Wireless Communications Letters* (2024).
- [29] Xin Li, Hongbo Wang, Zhe Chen, Zhiping Jiang, and Jun Luo. 2024. UWB-Fi: Pushing Wi-Fi towards Ultra-wideband for Fine-Granularity Sensing. In *Proc. of the 22nd ACM MobiSys*. 42–55.
- [30] Zhuo Lu, Wenye Wang, and Cliff Wang. 2013. Modeling, Evaluation and Detection of Jamming Attacks in Time-Critical Wireless Applications. *IEEE Transactions on Mobile Computing* 13, 8 (2013), 1746–1759.
- [31] Jun Luo, Hangcheng Cao, Hongbo Jiang, Yanbing Yang, and Zhe Chen. 2024. MIMOCrypt: Multi-User Privacy-Preserving Wi-Fi Sensing via MIMO Encryption. In *Proc. of the 45th IEEE S&P*. 1–19.
- [32] MediaTek. 2024. MediaTek MT7921 Wi-Fi 6 Featured in New ASUS ROG and TUF Gaming Notebooks. <https://www.mediatek.com/blog/mediatek-wi-fi-6-chipset-powers-new-asus-gaming-notebooks>. Online; accessed: 16 February 2024.
- [33] Aristides Mpitiopoulos, Damianos Gavalas, Charalampos Konstantopoulos, and Grammati Pantziou. 2009. A Survey on Jamming Attacks and Countermeasures in WSNs. *IEEE Communications Surveys & Tutorials* 11, 4 (2009), 42–56.
- [34] PwC Network. 2023. Global Telecom and Entertainment & Media Outlook 2023–2027. <https://www.pwc.com/gx/en/industries/tmt/assets/pwc-gto-2023.pdf>. Online; accessed: 16 February 2024.
- [35] Guevara Noubir, Rajmohan Rajaraman, Bo Sheng, and Bishal Thapa. 2011. On the Robustness of IEEE 802.11 Rate Adaptation Algorithms against Smart Jamming. In *Proc. of the 4th ACM WiSec*. 97–108.
- [36] Alina Olteanu and Yang Xiao. 2010. Security Overhead and Performance for Aggregation with Fragment Retransmission (AFR) in Very High-Speed Wireless 802.11 LANS. *IEEE Transactions on Wireless Communications* 9, 1 (2010), 218–226.
- [37] Ozgur Oyman and Sarabjot Singh. 2012. Quality of Experience for HTTP Adaptive Streaming Services. *IEEE Communications Magazine* 50, 4 (2012), 20–27.
- [38] Eldad Perahia and Michelle X Gong. 2011. Gigabit Wireless LANS: An Overview of IEEE 802.11 ac and 802.11 ad. *ACM SIGMOBILE Mob. Comput. Commun. Rev.* 15, 3 (2011), 23–33.
- [39] Hossein Pirayesh and Huacheng Zeng. 2022. Jamming Attacks and Anti-Jamming Strategies in Wireless Networks: A Comprehensive

- Survey. *IEEE Communications Surveys & Tutorials* 24, 2 (2022), 767–809.
- [40] Theodore S Rappaport. 2024. *Wireless Communications: Principles and Practice*. Cambridge University Press.
- [41] Raspberry Pi Foundation. 2024. Raspberry Pi 4 Model B. <https://www.raspberrypi.com/products/raspberry-pi-4-model-b/>. Online; accessed: 16 February 2024.
- [42] Realtek. 2024. RTL8812BU. <https://www.realtek.com/en/products/communications-network-ics/item/rtl8812bu>. Online; accessed: 16 February 2024.
- [43] Realtek. 2024. RTL8821CU. <https://www.realtek.com/en/products/communications-network-ics/item/rtl8821cu>. Online; accessed: 16 February 2024.
- [44] Matthias Schulz, Francesco Gringoli, Daniel Steinmetzer, Michael Koch, and Matthias Hollick. 2017. Massive Reactive Smartphone-based Jamming Using Arbitrary Waveforms and Adaptive Power Control. In *Proc. of the 10th ACM WiSec*. 111–121.
- [45] Shabnam Sodagari and T Charles Clancy. 2012. Efficient Jamming Attacks on MIMO Channels. In *Proc. of the IEEE ICC*. 852–856.
- [46] Krushang Sonar and Hardik Upadhyay. 2014. A Survey: DDoS Attack on Internet of Things. *International Journal of Engineering Research and Development* 10, 11 (2014), 58–63.
- [47] Tcpdump Group. 2024. Tcpdump& Libpcap. <https://www.tcpdump.org/>. Online; accessed: 16 February 2024.
- [48] Ajay Tirumala. 1999. Iperf: The TCP/UDP Bandwidth Measurement Tool. <http://dast.nlanr.net/Projects/Iperf/> (1999).
- [49] TP-Link. 2024. TP-Link Router XDR3020. https://www.tp-link.com.cn/product_2419.html. Online; accessed: 16 February 2024.
- [50] Lloyd N Trefethen and David Bau. 2022. *Numerical Linear Algebra*. Vol. 181. Siam.
- [51] D Tse. 2005. Fundamentals of Wireless Communication. *Cambridge University Press* 2 (2005), 614–624.
- [52] Yu-Chih Tung, Sihui Han, Dongyao Chen, and Kang G Shin. 2014. Vulnerability and Protection of Channel State Information in Multiuser MIMO Networks. In *Proc. of the 21st ACM CCS*. 775–786.
- [53] Bhagyashri Tushir, Yogesh Dalal, Behnam Dezfouli, and Yuhong Liu. 2020. A Quantitative Study of DDoS and E-DDoS attacks on WiFi Smart Home Devices. *IEEE Internet of Things Journal* 8, 8 (2020), 6282–6292.
- [54] Mathy Vanhoef and Frank Piessens. 2014. Advanced Wi-Fi Attacks Using Commodity Hardware. In *Proc. of the 30th ACM ACSAC*. 256–265.
- [55] Peng Wang and Li Ping. 2011. On Maximum Eigenmode Beamforming and Multi-User Gain. *IEEE Transactions on Information Theory* 57, 7 (2011), 4170–4186.
- [56] Sulei Wang, Zhe Chen, Yuedong Xu, Qiben Yan, Chongbin Xu, and Xin Wang. 2019. On User Selective Eavesdropping Attacks in MU-MIMO: CSI Forgery and Countermeasure. In *Proc. of the 38th IEEE INFOCOM*. 1963–1971.
- [57] Xiaoshan Wang, Yao Liu, Xiang Lu, Shichao Lv, Zhiqiang Shi, and Limin Sun. 2017. On Eavesdropping Attacks and Countermeasures for MU-MIMO Systems. In *Proc. of the IEEE MILCOM*. 40–45.
- [58] Xiaomi. 2024. Xiaomi 13 Pro. <https://www.mi.com/global/product/xiaomi-13-pro/>. Online; accessed: 16 February 2024.
- [59] Xiaomi. 2024. Xiaomi Redmi Router AC2100. <https://www.mi.com/rm2100>. Online; accessed: 16 February 2024.
- [60] Xiaomi. 2024. Xiaomi Router AX3000T. <https://www.mi.com/global/product/xiaomi-router-ax3000t/>. Online; accessed: 16 February 2024.
- [61] Xiaomi. 2024. Xiaomi Router AX3200. <https://www.mi.com/global/product/xiaomi-router-ax3200/>. Online; accessed: 16 February 2024.
- [62] Liangxiao Xin, David Starobinski, and Guevara Noubir. 2020. Cascading Attacks on Wi-Fi Networks: Theory and Experiments. *IEEE Transactions on Control of Network Systems* 7, 4 (2020), 1757–1768.
- [63] Qiben Yan, Huacheng Zeng, Tingting Jiang, Ming Li, Wenjing Lou, and Y Thomas Hou. 2014. MIMO-based Jamming Resilient Communication in Wireless Networks. In *Proc. of the 33th IEEE INFOCOM*. 2697–2706.
- [64] Taesang Yoo and Andrea Goldsmith. 2005. Optimality of Zero-Forcing Beamforming with Multiuser Diversity. In *Proc. of the IEEE ICC*. 542–546.
- [65] Taesang Yoo and Andrea Goldsmith. 2006. On the Optimality of Multiantenna Broadcast Scheduling Using Zero-Forcing Beamforming. *IEEE Journal on Selected Areas in Communications* 24, 3 (2006), 528–541.
- [66] Saman Taghavi Zargar, James Joshi, and David Tipper. 2013. A Survey of Defense Mechanisms against Distributed Denial of Service (DDoS) Flooding Attacks. *IEEE Communications Surveys & Tutorials* 15, 4 (2013), 2046–2069.
- [67] Zhazhan Zhang, Yin Sun, Ashutosh Sabharwal, and Zhiyong Chen. 2018. Impact of Channel State Misreporting on Multi-User Massive MIMO Scheduling Performance. In *Proc. of the 37th IEEE INFOCOM*. 917–925.
- [68] Shangqing Zhao, Zhuo Lu, Zhengping Luo, and Yao Liu. 2019. Orthogonality-Sabotaging Attacks against OFDMA-based Wireless Networks. In *Proc. of the 38th IEEE INFOCOM*. 1603–1611.