A First Look at Wi-Fi 6 in Action: Throughput, Latency, Energy Efficiency, and Security

RUOFENG LIU[*], Bosch Research, USA
NAKJUNG CHOI, Nokia Bell Labs, USA

This paper presents a first-of-its-kind performance measurement of Wi-Fi 6 (IEEE 802.11ax) using real experiments. Our experiments focus on multi-client scenarios. The results reveal the impact of the new channel access mechanisms (i.e., OFDMA and TWT) on the spectrum efficiency, energy consumption, latency, and network security. (i) A comparison with the legacy CSMA/CA scheme shows that the commodity Wi-Fi 6 achieves 3× overall throughput and dramatically reduce the latency (5×) when coexisting with legacy Wi-Fi network. (ii) However, the current OFDMA implementation significantly increases the power consumption (6×), implying a design tradeoff between throughput and latency gain versus the cost of energy consumption. (iii) Finally, TWT negotiating procedure is vulnerable to various malicious attacks. We believe that our findings provide critical insights for the scheduling algorithm design, power optimization, and security protection of the next-generation WLANs.

CCS Concepts: • **Networks → Network measurement**; **Network performance analysis**.

Additional Key Words and Phrases: IEEE 802.11ax, Wi-Fi 6, OFDMA, TWT

ACM Reference Format:
Ruofeng Liu and Nakjung Choi. 2023. A First Look at Wi-Fi 6 in Action: Throughput, Latency, Energy Efficiency, and Security. *Proc. ACM Meas. Anal. Comput. Syst.* 7, 1, Article 25 (March 2023), 25 pages. https://doi.org/10.1145/3579451

1 INTRODUCTION

Due to the ever-growing number of connected devices and deployments in dense environments, the traditional contention-based Wi-Fi network suffers severely from communication overhead, leading to inefficient utilization of the spectrum resources. To cope with these issues, the 6^{th} generation Wi-Fi protocol named IEEE 802.11ax [23] (also known as High-Efficiency Wi-Fi) is proposed. Wi-Fi 6 features two new channel access schemes: orthogonal division multiple access (OFDMA) and target wake time (TWT) [29]. These new schemes allow an AP to schedule clients to communicate in parallel with a frequency-division manner, i.e., OFDMA, or operate in separate time slots with a time-division manner, i.e., TWT, which can mitigate contention and communication overhead.

Despite the promising results from theory and simulations [10], the performance of the real Wi-Fi 6 network remains largely unknown. Recently, both commercial-off-the-self (COTS) Wi-Fi 6 AP (e.g., ASUS RT-AX58U) and client devices (such as Galaxy S10 smartphone and Intel AX210

[*]This work has been done during summer internship at Nokia Bell Labs when Ruofeng Liu was at the University of Minnesota.

Authors' addresses: Ruofeng Liu, ruofeng.liu@us.bosch.com, Bosch Research, USA; Nakjung Choi, Nokia Bell Labs, USA, nakjung.choi@nokia-bell-labs.com.

2476-1249/2023/3-ART25 $15.00
https://doi.org/10.1145/3579451

PCI-E NIC) became available, which motivates us to build a Wi-Fi 6 testbed with the latest Wi-Fi AP and 10 Wi-Fi clients including smartphone and PCI-E devices.

Using the established testbed, we conduct, to the best of our knowledge, the *first* comprehensive experimental evaluation of Wi-Fi 6 network. Our study focuses on multi-client scenarios and benchmarks several critical performance metrics including spectral efficiency (i.e., throughput), latency, power consumption, and security. The results reveal both the benefit of new channel access schemes as well as potential issues in current commodity devices. We envision that these results will provide important insights for future Wi-Fi system design (e.g. scheduling algorithm, energy optimization, and security protection).

Technically, benchmarking the multi-client network performance of Wi-Fi 6 is very challenging. First of all, the multi-client overall performance of OFDMA depends on not only the number of clients but also the characteristics of the application (e.g., traffic intensity and latency tolerance). For example, short transmissions of real-time VOIP and intensive traffic produced by ultra HD video streaming suffer from a significantly different amount of communication overhead, and thus the efficiency gains obtained from OFDMA are very distinct. Conventional experimental methods designed for a single-user Wi-Fi protocol (e.g., conducting the saturation throughput test using `iperf` for throughput and the latency test using `ping`) fail to produce real characteristics of these diversified traffic types so the results cannot represent the effectiveness of Wi-Fi 6 in the real world (could even lead to an elusive result that there are no significant throughput and latency difference between OFDMA and legacy CSMA/CA). Furthermore, understanding the power consumption of Wi-Fi 6 is also non-trivial because it is impacted by a large number of factors including the selected channel access scheme, allocated bandwidth of OFDMA clients, and again the traffic characteristics of the application.

To address these challenges, we propose several novel measurement methodologies. First, to examine throughput gain on real-world Wi-Fi traffic, we carefully manipulate the low-level frame aggregation procedure of Wi-Fi to generate traffic with various characteristics (e.g., short IoT transmission and intensive streaming). Secondly, We further highlight the latency performance of OFDMA in dense environments by intentionally producing external contention with a coexisting Wi-Fi 5 network. Additionally, we demystify the power consumption models by building power profiles for both OFDMA and TWT with an extensive measurement of the power consumed in various TWT states and OFDMA bandwidths. Using these new methodologies, we perform a comprehensive experiment across diverse scenarios and key observations are as follows.

- *Throughput (Section 4).* The traffic characteristics have a significant impact on OFDMA throughput gain. The commodity Wi-Fi 6 can achieve up to 3.2× and 2.6× throughput gain with UL and DL OFDMA for dense clients with short transmissions (a payload data size of fewer than 1.6 Kbytes). The throughput gain becomes less notable for transmissions produced by intensive traffic or when there are few clients.
- *Latency (Section 5).* OFDMA significantly reduces the latency due to channel access (5×) when there are severe contentions with a coexisting Wi-Fi 5 network. In addition, UL OFDMA mitigates the internal contention among Wi-Fi 6 clients.
- *Power consumption (Section 6).* The current commodity Wi-Fi devices are not effectively optimized for small bandwidth allocated by OFDMA (i.e., not energy-proportional to bandwidth), which could lead to high energy consumption in the dense-client scenarios. We observe a 6× and 2.4× increase in energy consumption for intensive traffic and short transmission, respectively. In contrast, TWT maintains a stable power consumption. The results suggest a tradeoff design in which the Wi-Fi scheduler can combine OFDMA and TWT to balance throughput and latency improvement and the extra energy consumption cost.

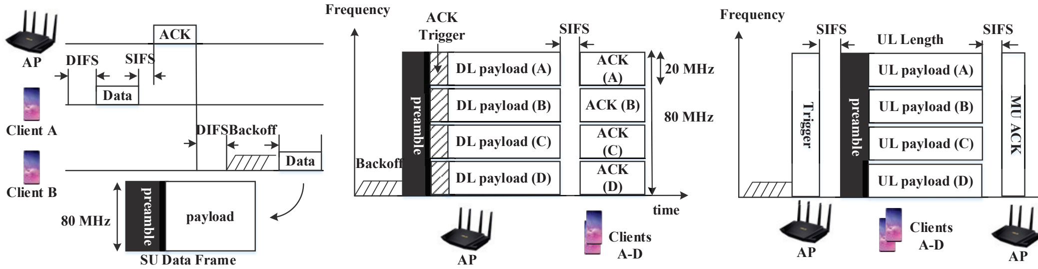

(a) CSMA/CA (b) DL OFDMA (c) UL OFDMA

Fig. 1. Legacy CSMA/CA vs. OFDMA in Wi-Fi 6.

- *Security (Section 7).* Unencrypted TWT negotiation procedures in commodity Wi-Fi devices are vulnerable to various types of malicious attacks that cause battery depletion, disconnection, and fairness issues.

Finally, we conduct a case study of Wi-Fi network slicing as an example of applying our insights to Wi-Fi network design (Section 8). By combining OFDMA and TWT, we can achieve resource isolation and quality of service while ensuring throughput, latency, and energy consumption.

The paper is organized as follows. We provide an overview of new channel access schemes in Wi-Fi 6 (Section 2) and introduce our experiment setup (Section 3). The methodology of the throughput test and results are given in Section 4. Section 5 discusses the latency experiment design and results. The power consumption of Wi-Fi 6 devices is presented in Section 6 in detail. Section 7 investigates the security vulnerability of TWT, followed by a case study of Wi-Fi slicing (Section 8). The limitations and future works are discussed in Section 9. We wrap up the paper with the related work (Section 10) and conclusion (Section 11).

2 BACKGROUND

This section introduces issues stemming from legacy CSMA/CA and new mechanisms (i.e., OFDMA and TWT) proposed in Wi-Fi 6 that are used to achieve high efficiency.

CSMA/CA. First, we briefly recapitulate CSMA/CA scheme used in the legacy Wi-Fi protocol. As Fig. 1(a) illustrates, devices (both AP and clients) in the legacy WLAN contend for the channel in a fully distributed manner. A client (i.e., Client B) defers its transmission until the current transmitting client (i.e., Client A) completes. Then, in order to avoid collision with other clients, it generates a random backoff delay. The backoff timer is decreased when the medium is sensed as idle for a 34-microsecond DIFS (DCF inter-frame space) and frozen when any ongoing transmission is detected on the medium. When the timer reaches 0, the client transmits its packet. The structure of the packet is depicted at the bottom of Fig. 1(a). In the frequency domain, each packet occupies the entire frequency band of the channel. In the time, it consists of a preamble with a fixed duration (typically 30.4 us) and a payload of variable duration.

CSMA/CA incurs significant communication overhead because every packet needs to pay an additional toll for preamble transmission and channel contention (e.g., DIFS and backoff). This traditional scheme leads to low spectral efficiency, especially in dense environments, motivating new channel access mechanisms.

OFDMA. To reduce the communication overhead, OFDMA divides the entire frequency band into resource units (RU) and allocates multiple clients to communicate in parallel. In specific, Fig. 1(b) shows an example of downlink OFDMA (DL OFDMA), in which the AP contends for channel access for all clients (A-D) and delivers the message to each client in the disjointed smaller band (i.e., RU). By doing this, the communication overhead is amortized by a single transmission to multiple

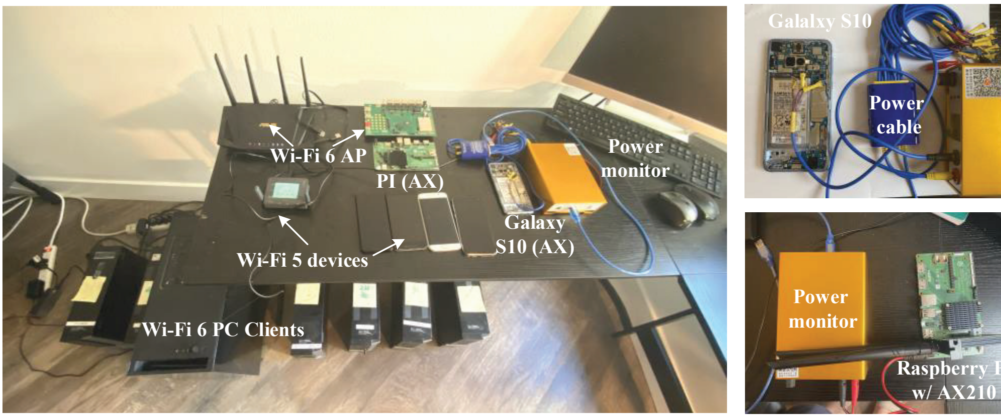

(a) Wi-Fi 6 testbed overview (b) Energy consumption measurement testbed

Fig. 3. Wi-Fi 6 measurement testbed overview.

clients at the same time. Moreover, the downlink frame triggers the client to acknowledge in the parallel using uplink OFDMA (UL OFDMA) which is illustrated in Fig. 1(c) in more detail. In UL OFDMA, the AP sends a trigger frame to synchronize multiple clients to transmit simultaneously in disjointed RUs, thus eliminating the overhead due to the contention among clients. The trigger frame specifies the allocations of RU and the duration of uplink transmission (i.e., UL length). The size of RU (Table 1) varies from 26 to 996, allowing the AP to allocate various numbers of OFDMA clients (from 1 to 37) in one transmission. A smaller RU size provides lower maximum data rates due to the smaller bandwidth.

RU size	#Clients	Max Data rates
26-tone	37	23.6 Mbps
52-tone	16	47 Mbps
106-tone	8	100 Mbps
242-tone	4	286.8 Mbps
484-tone	2	573.6 Mbps
996-tone	1	1200.8 Mbps

Table 1. OFDMA resource unit (80 MHz, 2 streams).

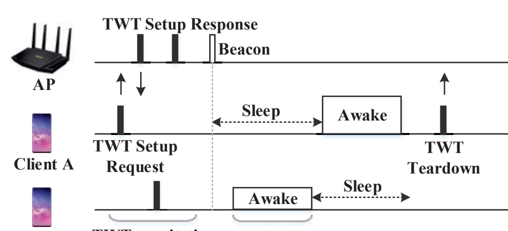

Fig. 2. TWT setup, operation and teardown

TWT. In contrast with the frequency-division approach of OFDMA, target wake time (TWT) allows an AP to allocate clients to operate at non-overlapping times. It minimizes medium contention between clients. It also reduces the required amount of time in which a client needs to be awake, thus saving energy. TWT procedure is demonstrated in Fig. 2 where the client negotiates an individual TWT agreement with the AP that specifies the service periods. The device then wakes up and exchanges messages only during the predefined period. Both AP and clients can tear down the agreement when the service does not satisfy its need. Alternatively, a client can join an existing broadcast TWT agreement that is advertised in the Wi-Fi beacon.

3 EVALUATION SETUP

3.1 Testbed overview

Fig. 3 presents our Wi-Fi 6 evaluation testbed, consisting of Wi-Fi 6 AP, clients, and auxiliary devices. More specifically, we experiment with two representative Wi-Fi 6 AP, i.e., ASUS RT-AX58U

Table 2. Wi-Fi 6 AP and client devices in the measurement testbed.

Device	Role	Chipset	SS	BW	#clients	Manuf.	OS
RT-AX58U	AP	Broadcom BCM6750	2×2	160 MHz	8	ASUS	asuswrt
DR6018	AP	Qualcomm IPQ6018	2×2	160 MHz	8	Wallys	QSDK
Galaxy S10	Client	Broadcom BCM4375	2×2	80 MHz	N/A	Samsung	Android
Intel Gig+	Client	Intel AX200/210	2×2	160 MHz	N/A	Intel	Ubuntu

[9] (based on Broadcom BCM6750 Wi-Fi chipset) and DR6018 [12] (based on Qualcomm IPQ6018 chipset). Both APs support Wi-Fi 6 features including UL/DL OFDMA, TWT, and 1024 QAM. In addition, the current commercial AP can serve up to 8 clients in an OFDMA transmission[1].

On the client side, we have 10 Wi-Fi 6 devices including 6 Dell OptiPlex 790 desktops [15] equipped with Intel AX210 Wi-Fi 6 PCI-E NIC [22], 2 Lenovo Thinkpad T470p laptops with Intel AX200 PCI-E NIC, 1 Raspberry Pi cm4 with AX210, and 1 Samsung Galaxy S10 with integrated Wi-Fi chipset (BCM4375). Ubuntu 20.04 is used as the operating system for clients with Intel PCI-E NICs. To support the latest Wi-Fi 6 features, we install backport-iwlwifi driver and deploy the up-to-date iwlwifi firmware (i.e., iwlwifi-ty-a0-gf-a0-71.ucode). In the backhaul part of the network, an ASUS H170 Pro desktop is connected to the AP via 1Gbps Ethernet. Finally, a coexisting Wi-Fi 5 network is established using GL-AR750S 802.11ac AP [2] and 4 Nexus 5X smartphones. The specification of Wi-Fi devices in our testbed is summarized in Table 2.

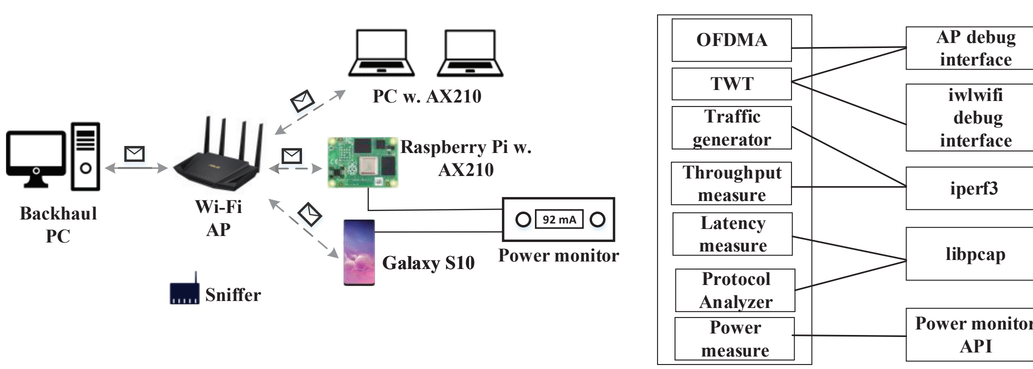

(a) Wi-Fi 6 Testbed hardware diagram (b) software architecture

Fig. 4. The hardware and software architecture of Wi-Fi 6 measurement platform.

3.2 Evaluation methodology

The section provides the general evaluation methodology. More detailed methods of each metric (throughput, energy, and latency) are given in their respective sections. In general, measuring Wi-Fi performance in multi-client scenarios is a non-trivial task because manually coordinating a large number of Wi-Fi devices is very labor-intensive and error-prone, especially when we have to repeat the experiment for a variety of combinations of configurations. To address the challenge, we develop an automated multi-client experiment platform (Fig. 4(a)), which automatically configures the parameters of devices on the test, generates the traffic, and collects the data for analysis[2]. The components of the platform depicted in Fig. 4(b) include:

[1]Wi-Fi standard allows 37 clients for 80 MHz and 74 clients for 160 MHz. We will combine our empirical results of 8 clients and theoretical model to estimate the performance on a larger scale.
[2]We release the measurement platform and collected data sets at https://github.com/liux4189/wifi-ax-measurement.

- *OFDMA manager.* The component can disable/enable OFDMA and configure resource unit (RU) allocation through the debug interfaces provided by iwlwifi, QSDK, and asuswrt (e.g., `wl` and `nvmram`). It supports four representative RU configurations: (1) SU: the entire bandwidth (996 tones) is allocated to one client, (2) 2MU: 996 tones are equally shared by 2 clients using 484 tones each, (3) 4MU: four 242-tones RU allocated to 4 clients, and (4) 8MU: eight 106-tone RU allocated to 8 clients. OFDMA manager also configures other Wi-Fi parameters that impact OFDMA performance (e.g., frame aggregation and MCS).
- *TWT manager.* The manager sets up and tears down TWT sessions from clients using `twt_setup` debug interface provided by iwlwifi driver. Alternatively, TWT can also be configured by the AP using `wifitool` provided by QSDK.
- *Traffic generator.* After OFDMA or TWT manager completes the configuration, the traffic generator produces UL or DL UDP traffic using `iperf3`. Without additional explanations, each experiment takes 60 seconds. We repeat a test 20 times and report the statistical results.
- *Throughput, latency, and consumption analysis.* Our testbed automatically collects the throughput and latency during the experiment with `iperf3` and `libpcap`. For energy consumption analysis, Raspberry Pi and Galaxy S10 are connected to an EKA1080p power monitor [16] (shown in Fig. 3(b)) with a 100 kHz sample rate. To measure the S10 smartphone, we remove its back cover and modify the power supply interface to adapt it to the power monitor.

Our experiment is conducted in a home environment during the night time to minimize interference by external factors. During the experiment, the clients are placed in a circle with a 4-meter distance from the AP. Channel 161 in 5 GHz is used.

4 THROUGHPUT

4.1 Methodology

We evaluate the empirical spectral efficiency improvement fulfilled by the commodity Wi-Fi 6 network by measuring the multi-client aggregated throughput achieved by OFDMA using our novel measurement methodology.

Why the conventional saturated throughput test is ineffective? Conventionally, a throughput measurement is done by directly performing a saturated throughput test using tools (e.g., `iperf3`). The traditional method cannot characterize the performance of OFDMA for real-world traffic because Wi-Fi communication overhead highly depends on the characteristics of the application (e.g., traffic intensity and latency tolerance). For example, real-time VOIP produces short transmissions of 200 bytes which are dominated by the preamble and channel access delay (> 80%), whereas mobile VR and UHD video generate more intensive traffic and thus the time consumed by the overhead is less. Yet, in a conventional saturated throughput test, low-level Wi-Fi protocol could aggregate a large number of (up to 256) MAC Protocol Data Unit (MPDU) into a huge aggregated MPDU (AMPDU), which results in impractically long transmissions that suffer very little from communication overhead. This makes the benefits of OFDMA invisible, which explains the illusion in the recent industrial report [21] that OFDMA does not offer a perceivable throughput gain.

Challenge of producing traffic characteristics. We cannot simply disable AMPDU because DL OFDMA inherently depends on AMPDU to piggyback a trigger frame for UL acknowledgment in downlink transmissions (as we introduced in Section 2). Therefore, we found that an AP automatically turns off OFDMA mode if AMPDU is disabled.

Solutions for downlink. To produce the diversified traffic characteristics in the downlink transmission, we carefully exploit the block acknowledgment parameters. As Fig. 5(a) depicts, before data exchanges, an AP sends an `AddBA` request to a client for the maximum number of AMPDU the client supports (denoted as AMPDU size). We intentionally modify the AMPDU size in the client's

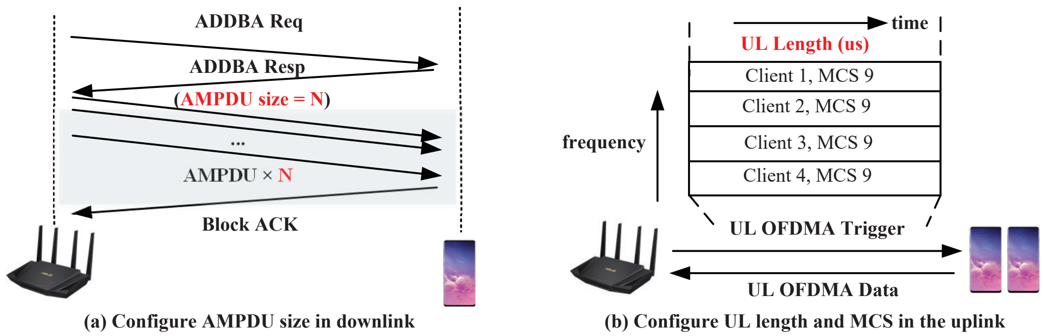

Fig. 5. Techniques to generate real aggregated packets (a) Downlink, (b) Uplink.

response to indirectly configure the number of MPDU aggregated by AP. By choosing different AMPDU sizes, we produce transmission with various aggregated payload lengths representing different traffic characteristics. For example, setting AMPDU size to 1 produces short transmissions (e.g., IoT and VOIP) while configuring the size to 10 produces intensive traffic by mobile VR. We found that this parameter is configurable in high-level Wi-Fi software (e.g., `buf_size` in mac80211).

Solutions for uplink. For uplink, we figure out the number of MPDU aggregated at clients is decided by the maximum duration of uplink transmission (i.e., `UL length` parameter in the UL OFDMA trigger frame from AP) and the modulation order of the data. Therefore, we develop a table that maps the desired combination of AMPDU size and MCS to the value of `UL length`. Finally, we configure the value via debug tool of the AP, so that the clients naturally limit the AMPDU size according to the maximum duration.

4.2 Results

4.2.1 SU throughput. We first examine the basic single-user Wi-Fi 6 throughput with various bandwidth, MCS, and AMPDU settings. Fig. 6(a) depicts the throughput using the default AMPDU configuration. For 20, 40, and 80 MHz bandwidth, the throughput increases linearly with MCS whereas the throughput of the 160 MHz setting stops growing after MCS 6, achieving a maximum data rate of 967 Mbps. We cross-validate the results on both RT-AX58U and DR6018 and figure out that the maximum throughput is constrained by the backhaul (i.e., 1Gbps Ethernet port of AP). To avoid the impact of Ethernet, we mainly experiment on 80 MHz in this study.

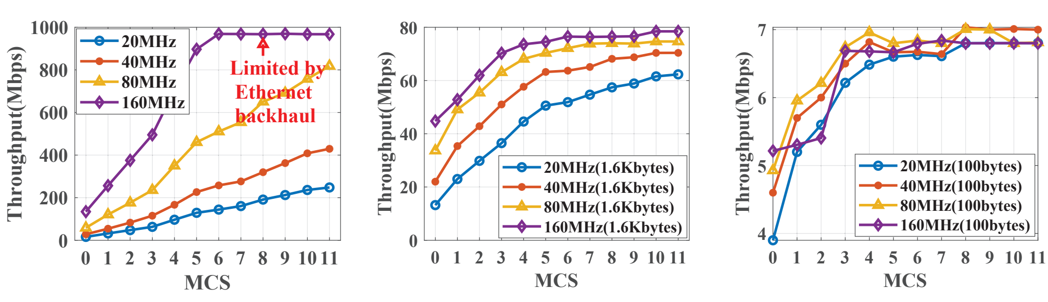

Fig. 6. SU throughput (a) Default AMPDU configuration, (b) 1.6 Kbytes (c) 100 Kbytes.

Fig. 6 (b) and (c) show the results when the AMPDU size is adjusted to 1. With a packet length of 1.6 Kbytes (maximum length of an MPDU), we obtain at most 74 Mbps data rate which decreases to 9.05% of the throughput we obtain with the default AMPDU configuration. The throughput

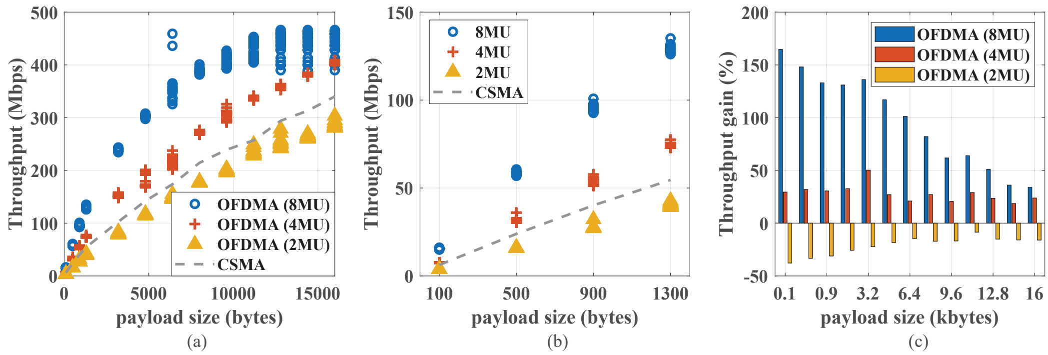

Fig. 7. Downlink aggregated throughput: (a) OFDMA vs. CSMA, (b) short transmission, (c) OFDMA gain.

is further reduced to 6.8 Mbps when the packet size is decreased to 100 bytes. The results show that our manipulation of AMPDU size is effective. In addition, increasing the modulation order and bandwidth offers marginal throughput improvement. This is due to severe communication overhead in such transmissions and the next subsection will examine if OFDMA can improve it.

4.2.2 DL OFDMA throughput. In this section, we report the throughput gain of downlink OFDMA. To produce real traffic characteristics of various intensities, AMPDU size is adjusted from 1 to 10, generating various aggregated packet lengths from 100 bytes to 16 Kbytes. We measure the aggregated throughput using 3 multi-user OFDMA settings (2, 4, and 8 clients) and CSMA/CA setting. For the CSMA/CA, various numbers of clients produce similar throughput results. For the sake of simplicity, we report the results of 8 clients as a representative. More specifically, Fig. 7 (a) compares the absolute aggregated throughput achieved by OFDMA and legacy CSMA/CA (grey dashed line). We zoom in for the detailed results of short transmissions (AMPDU size = 1) in Fig. 7 (b). Fig. 7 (c) further depicts the relative throughput gain of OFDMA, which is defined as Equation (1), where TP_i^{OFDMA} and TP_i^{CSMA} are OFDMA and CSMA/CA throughput of i^{th} client.

$$Gain_{TP} = \frac{\sum_i TP_i^{OFDMA} - \sum_i TP_i^{CSMA}}{\sum_i TP_i^{CSMA}} \tag{1}$$

We observe that in the optimal scenario (i.e., short transmissions and 8 clients), DL OFDMA achieves $2.6\times$ **overall throughput compared to CSMA/CA.** For instance, with a 1.3 Kbyte payload, OFDMA improves the throughput from 54 Mbps to 128 Mbps, which is equivalent to doubling the capability or the number of devices to serve. The throughput gain gradually decreases as the payload size increases. For example, the maximum throughput improvement of 16 Kbytes is 35%. The results confirm that short Wi-Fi transmissions suffer most from the communication overhead (e.g., preamble and channel access delay), while OFDMA effectively amortizes this overhead across the clients. **The gain drops dramatically for 4MU and 2MU configuration.** For all payload sizes, the benefit of 4MU OFDMA clients is always less than 50% because the overhead is shared by less number of devices. Counter-intuitively, OFDMA even negatively reduces overall throughput when serving 2 clients in parallel. We found that this is due to the increased interframe space in OFDMA implementation (to be discussed in Section 4.2.4).

4.2.3 UL OFDMA performance. The results of UL OFDMA measurement are depicted in Fig. 8. Overall, UL OFDMA provides a slightly larger throughput gain than DL OFDMA. Specifically, UL OFDMA achieves a $3.2\times$ aggregated throughput for short transmissions in 8MU. The additional improvement comes from the extra advantage of OFDMA in the uplink, i.e., scheduled parallel uplink

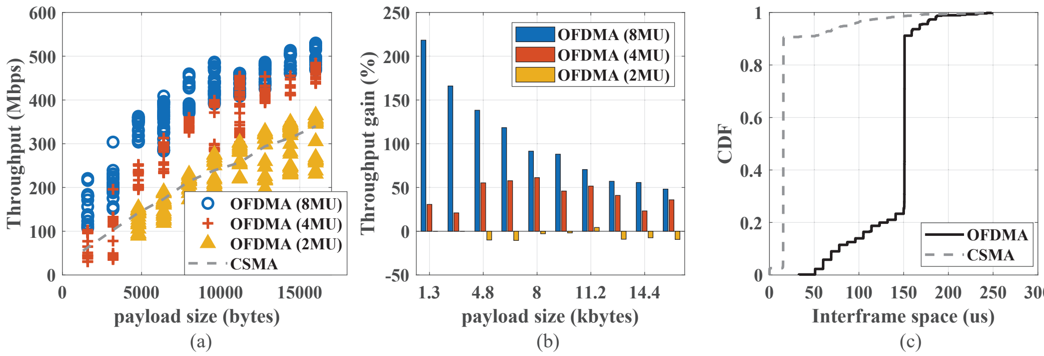

Fig. 8. Uplink aggregated throughput: (a) OFDMA vs. CSMA, (b) OFDMA gain, (c) Interframe space

transmissions eliminate channel contention among clients. In contrast, the downlink transmission is scheduled by the AP and thus doesn't suffer from contention among clients. Similar to the downlink, the gain significantly drops for 4MU and 2MU configurations.

In addition to OFDMA, we also evaluate the throughput using individual TWT. As expected, the TWT throughput is proportional to the duty cycle (i.e., the ratio of wake duration). Because the results are intuitive, we gloss over the details in the paper.

4.2.4 Interframe space. In the previous sections, we observe that OFDMA reduces the throughput when there are few clients (e.g., 2MU). Through a detailed investigation, we found the current Wi-Fi 6 implementation significantly increases the interframe space when OFDMA is enabled, which weakens the strength of OFDMA. Fig. 8(c) compares the cumulative distribution function (CDF) of the idle time between SU frames and OFDMA frames. We found for all OFDMA configurations, successive OFDMA frames are separated by a constant space (approximately 150 us), which is much larger than the 34 us DIFS of SU. We cross-validate the phenomenon on both RT-AX58U (Broadcom BCM6750) and DR6018 (Qualcomm IPQ6018) and obtain similar observations. Without the access to the low-level states in commodity Wi-Fi AP, we do not have a conclusion for the reason. Yet, we hypothesize that AP might reserve extra time for MU operations (e.g., MU scheduling and processing MU acknowledgment).

4.3 Implication to applications

As we show above, OFDMA is especially beneficial when there are a large number of clients with short transmissions. To understand the potential benefits of OFDMA to various applications, we collect the real Wi-Fi trace of several representative applications on Samsung Galaxy S10 to analyze their traffic patterns and aggregated packet lengths. Examples of various traffic patterns are visualized in the first row of Fig. 9 and AMPDU sizes are shown in the second row where each data point denotes an AMPDU. The statistics are presented in Table 3, where "Rate", "Length", "AMPDU" columns represent the data rates, the average aggregated packet length, and the average number of aggregated MPDU in a Wi-Fi transmission. Detailed analyses are as follows.

VOIP (Voice call). Voice calls (e.g., Skype, Viber, Zoom) feature low traffic intensity and strict latency requirements. The first column of Fig.9 shows the Wi-Fi traffic we collect on S10 running Zoom with the video turned off. We observe that packets of voice calls are transmitted immediately and thus rarely aggregated because frame aggregation could significantly increase the latency. The average aggregated Wi-Fi packet lengths are 391 and 273 bytes in the uplink and downlink respectively and there is usually only a single MPDU in a transmission (AMPDU size = 1). As a result, OFDMA can increase the overall throughput of concurrent VOIP traffic (3×), improving the spectrum efficiency in dense environments (e.g., offices with simultaneous voice calls).

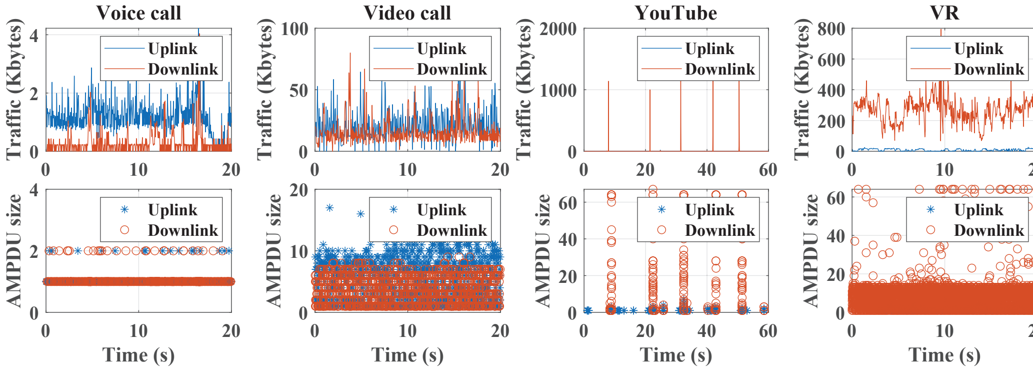

Fig. 9. Traffic Pattern and aggregated Wi-Fi transmission of representative applications.

IoT. IoT communications (e.g., smart home monitoring, and machine-to-machine communication in the industry) feature massive devices that produce a small amount of sensing data or control commands. We replay real IoT traffic in the public IoT dataset [1] (captured with a Philips HUE smart LED lamp, and a Somfy smart door lock) on our testbed using Tcpreplay[3]. In both devices, the average packet length is below 352 bytes and the maximum packet length is 1.6 Kbytes. OFDMA can significantly improve the spectral efficiency of dense IoT networks, achieving around 3× throughput gain and serving massive IoT devices with $\frac{1}{3}$ spectrum resource of CSMA/CA.

Video call. Video conferences (e.g., Zoom, Teams, Meet) produce real-time traffic with higher intensity than VOIP and IoT. As the second column of Fig. 9 depicts, we measure video conference traffic produced by Zoom in high-resolution mode (UHD). We observe due to the traffic intensity of video, Wi-Fi aggregates 3-5 video frames in one transmission. The average aggregated packet lengths are 4.5 and 6.6 Kbytes for downlink and uplink. Using OFDMA can achieve a maximum throughput gain of 99% when there are more than eight concurrent video calls.

YouTube. In contrast to video conferences, the traffic of video streaming services (e.g., YouTube) is non-real-time and unidirectional. We measure a 4K HDR YouTube video [4] playing on S10. Interestingly, although YouTube has a comparable downlink data rate as video call, a Wi-Fi AP produces much larger aggregated packets (with 20 MPDUs and a packet length of 27.6 Kbytes on average). This is caused by the latency tolerance of video streaming shown in the third column of Fig. 9 - Data are buffered at the server and delivered to the client in a batch every 10 seconds. The client consumes the data and requests the next batch. Due to the large aggregated frames (thus low overhead), OFDMA provides marginal throughput gain for YouTube.

VR. VR features intensive traffic with hybrid real-time characteristics (both prefetched downloads and real-time updates). We use Tcpreplay to replay real Virtual Reality traffic captured with Oculus Quest VR headset by the authors in [5]. The fourth column of Fig. 9 shows that the VR produces intensive downlink traffic (44 Mbps), which leads to AMPDUs with 8 aggregated packets on average. We also observe several very large AMPDUs with up to 64 MPDUs (by prefetched data). OFDMA can improve the efficiency of real-time data of VR with a throughput gain of 50%.

4.4 Summany and Takeaways

Our empirical study shows that commodity Wi-Fi 6 can achieve a maximum of 3.2× and 2.6× throughput gain with UL and DL OFDMA while the real-world benefit highly depends on the characteristics of applications such as traffic intensity and latency tolerance. By analyzing the aggregated Wi-Fi packet of various applications, we find that the application with low traffic intensity and high latency sensitivity (e.g., VOIP and IoT) produces very short aggregated packets (a few hundred bytes), which benefit most from OFDMA. Real-time applications with more intensive traffic (e.g., video calls) produce aggregated packets of a few Kbytes and thus can also take advantage of OFDMA when there are dense clients. Finally, the throughput gain is less significant for the non-real-time application with high traffic intensity (e.g., YouTube).

App	$Rate_{DL}$	$Rate_{UL}$	$Length_{DL}$	$Length_{UL}$	$AMPDU_{DL}$	$AMPDU_{UL}$
VOIP	44.59 Kbps	185.38 Kbps	273 bytes	391 bytes	1.10 MPDUs	1.02 MPDUs
Doorlock	1.93 Kbps	904 bps	352 bytes	190 bytes	1 MPDU	1 MPDU
HUE	34 bps	367 bps	95 bytes	270 bytes	1 MPDU	1 MPDU
Zoom	2.56 Mbps	3.30 Mbps	4501 bytes	6661 bytes	3.53 MPDUs	5.12 MPDUs
YoTube	1.42 Mbps	18.49 Kbps	27630 bytes	509 bytes	20.19 MPDUs	1.30 MPDUs
VR	44.42 Mbps	2.21 Mbps	11966 bytes	1651 bytes	8.00 MPDUs	1.06 MPDUs

Table 3. Characteristics of representative applications.

5 LATENCY

Wi-Fi 6 proposes UL and DL OFDMA to mitigate channel contention, which aims at reducing network latency. The section examines the latency performance of Wi-Fi 6.

5.1 Methodology

Why the conventional latency test is ineffective? Conventionally, the latency of a network is commonly measured by generating saturated traffic and using existing tools (e.g., `ping`) to measure round-trip delay. However, this method has several limitations for Wi-Fi 6. First, the results combine uplink and downlink latency. Wi-Fi 6 uses different OFDMA mechanisms on the uplink and downlink so their impacts should be evaluated separately. Furthermore, OFDMA is most effective to handle severe channel contention in dense environments. Without contention with external competitors (i.e., stations of another network), the benefit could be hard to perceive (especially for DL).

Measurement methodology. To examine the benefits of OFDMA in dense environments, we intentionally create external competitors by deploying a Wi-Fi 5 network with 4 legacy clients. The legacy clients send uplink traffic with various data rates and produce different levels of contentions. In each test, we gradually increase the data traffic of Wi-Fi 6 towards saturation. We repeat the experiment with OFDMA and CSMA/CA and compare their performance under the same data rate and level of contention. Additionally, to separately obtain one-way latency for UL and DL, we install a Wi-Fi sniffer on the device transmitting the packet thus avoiding time synchronization issues. A timestamp t_0 is made when a packet enters the network stack and another timestamp t_1 is collected when the sniffer detects that the packet is successfully delivered. $\Delta t = t_1 - t_0$ represents the delay due to channel contentions.

5.2 Results

5.2.1 Downlink Latency. First, we evaluate the downlink latency without the external competing network. We turn off Wi-Fi 5 network and generate saturated downlink traffic from the backhaul PC to a various number of clients (2 to 8). Fig. 10(a) compares the average latency achieved by OFDMA and CSMA/CA. **Without external competitors, DL OFDMA and CSMA/CA achieve comparable latency in the downlink (below 4 milliseconds).** This is because the downlink traffic of clients is totally scheduled by the AP and thus do not content with each other. With few clients ($\leq$ 4), OFDMA latency is even slightly higher than CSMA/CA due to the increased interframe space discussed in Section 4.2. This result confirms our measurement methodology that external competitors are needed for fully evaluating the OFDMA.

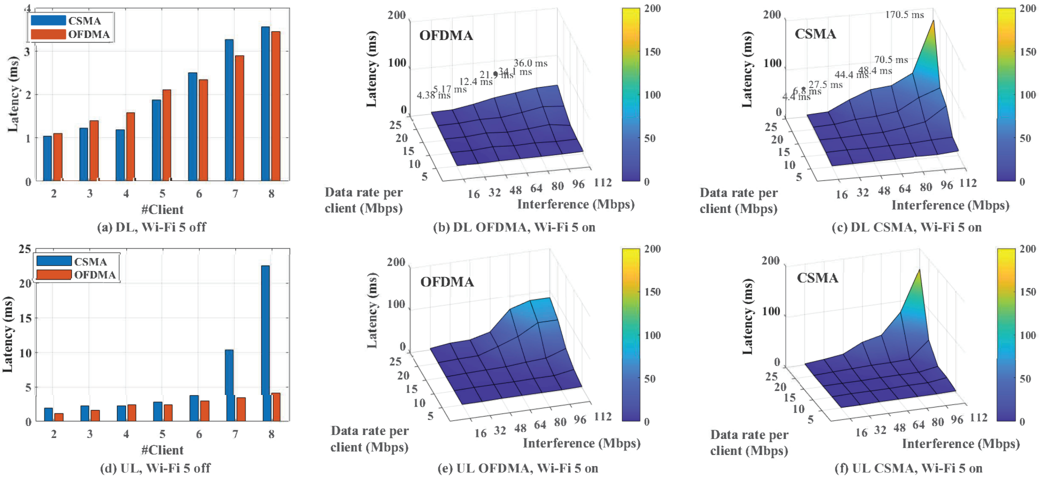

Fig. 10. UL and DL latency benchmark. (a) DL latency w/o external contention: OFDMA vs. CSMA, (b) DL OFDMA latency w. external contention, (c) DL CSMA latency w. contention, (d) UL latency w/o contention: OFDMA vs. CSMA, (e) UL OFDMA latency w. contention, (f) UL CSMA latency w. contention,

Fig. 10(b) and (c) compare OFDMA and CSMA/CA latency with various data rates and contention levels (interference axis). **When there is competing traffic from Wi-Fi 5 network, OFDMA demonstrates significant latency reduction.** As Fig. 10(c) shows, the latency of CSMA/CA rises up dramatically with the increasing contention with Wi-Fi 5 network. For example, CSMA/CA leads to an average latency of 48.5 ms on 25 Mbps Wi-Fi 6 traffic and 64 Mbps external contention, while it surges sharply to 170.5 ms on 112 Mbps contention. In comparison, OFDMA latency (depicted in Fig. 10(b)) grows slowly to a maximum of 35.58 ms, effectively achieving 5× latency mitigation. The results demonstrate the advantage of using DL OFDMA to mitigate severe competition among coexisting Wi-Fi networks. DL OFDMA allows AP to deliver data to multiple clients with one channel access, whereas CSMA/CA competes for the channel for each client's downlink packet.

5.2.2 Uplink Latency. In contrast to downlink, **we observe UL OFDMA effectively reduces the latency even when there are no external competitors.** As Fig. 10(d) shows, the latency of CSMA/CA increases from 1.8 ms to 22.4 ms when the number of clients increases from 2 to 8. This is because in the uplink there are internal contentions among clients, which leads to backoff delay and even packet collisions when the number of clients is large. In comparison, the latency of OFDMA grows much more slowly because the AP schedules the uplink traffic from clients.

Fig. 10 (e) further demonstrates OFDMA latency when the Wi-Fi 5 produces external contention. We observe that the latency remains low when the data rate of contention is lower than 64 Mbps. When further increasing the Wi-Fi 5 data rate, the latency rises to 60.2 ms. With the same amount of external contention, UL OFDMA incurs higher latency than the downlink. One possible reason is that in UL OFDMA, AP needs to frequently acquire the buffer status of clients in order to decide the clients to solicit in the trigger frame. The external contention might interrupt and postpone the buffer status report exchanges, incurring an extra scheduling delay. However, compared with the results of CSMA/CA in Fig. 10(e), UL OFDMA still achieves latency reduction (up to 2.2×).

5.3 Takeaways

The empirical results show that OFDMA effectively reduces the uplink and downlink channel access delay when there is an external competing network. Moreover, UL OFDMA offers extra benefits of mitigating internal contention among clients.

6 ENERGY EFFICIENCY

6.1 Methodology

The energy efficiency of Wi-Fi devices plays an important role in the user experience and the sustainability of the network. So far, we only examine the benefits that new features of Wi-Fi 6 brings about (such as spectral efficiency and latency reduction). Yet, their impact on energy consumption is not clearly understood. This important but missing piece motivates us to conduct a comprehensive evaluation of the energy efficiency of Wi-Fi 6 devices.

6.1.1 Why energy efficiency of 802.11ax is complicated? Understanding the energy consumption of Wi-Fi 6 is significantly more challenging than the prior versions of Wi-Fi protocol because it is impacted by several factors. We discuss the key factors as follows:

Various channel access approaches. The devices in Wi-Fi 6 networks may communicate with the AP using three distinct MAC schemes including OFDMA, TWT, and CSMA/CA. Each MAC scheme has a unique energy consumption model. Systematically comparing the pros and cons of these schemes is non-trivial.

OFDMA RU size. Among the three schemes, the energy consumption of OFDMA is the most sophisticated. OFDMA allows an AP to allocate a proportion of the overall bandwidth (specified by the resource unit size) to a user, which could make two conflicting impacts on energy consumption. **1)** On the one hand, the devices transmit and receive on a smaller bandwidth. This gives a Wi-Fi radio the opportunity to turn off unnecessary electronic components and effectively cut down the instant power output. **2)** On the other hand, with a smaller bandwidth, the data rate is also reduced, leading to a longer time to complete TX and RX. This adversely increases the overall energy consumption. Therefore, to obtain the overall impact, an empirical measurement must be conducted over various OFDMA RU sizes (i.e., various small bandwidths).

Traffic characteristics. We show in Section 4 that OFDMA also increases the overall throughput. This could, to some extent, mitigate the energy costs when the device is allocated to a small bandwidth by OFDMA. However, as we discuss in Section 4, the throughput improvement varies significantly with different traffic intensities (e.g., IoT and video conferencing).

Other parameters. Finally, there are other parameters (e.g., Modulation and Coding Scheme (MCS)) in a Wi-Fi radio that impact power consumption. They should also be considered to obtain a comprehensive result.

6.1.2 Measurement methodology. To analyze the energy consumption in various multi-device scenarios, we use `iperf3` to produce the parallel traffic from AP to several clients (for DL test) and from clients to AP (for UL test). During the experiment, one client device (Raspberry Pi or S10) is connected to the power monitor. The throughput and instant power are collected from the device and the energy consumption (energy consumed for one bit, in nJ/bit) is derived via the ratio of the average instant power (in J/s) over the average throughput (in bits/s) as indicated in equation 2.

$$Energy\ Efficiency = \frac{\sum_t Power(t))}{\sum_t Throughput(t)} \tag{2}$$

To compare various channel access schemes, we repeat the experiment for OFDMA, TWT, and CSMA/CA. During the OFDMA test, we vary the number of current OFDMA clients such that the clients are allocated with various RU sizes (i.e., small bandwidth). The average TX/RX power of the radio operating on each RU size is recorded to build a power profile. We analyze the profile to understand the impact of RU size on the instant power (details in Section 6.2.2). In the TWT test, we use the developed TWT manager to establish individual TWT sessions between clients and the AP. The clients exchange single-user packets with the AP in the non-overlapped TWT time slots without contention. We build a power profile that records the power consumption in each stage of a TWT session. The effectiveness of energy saving using TWT is then examined (more detail in Section 6.2.3). For CSMA/CA experiment, we disable the OFDMA at the AP configuration and tear down all the TWT sessions. The clients contend for the channel during the entire experiment.

The structure of the following sections is as follows. In Section 6.2.2 and 6.2.3, we first build the profile of the instant power for OFDMA and TWT. Then, in Section 6.2.4, we will benchmark the overall energy efficiency of these schemes for traffic with various characteristics.

6.2 Results

6.2.1 Baseline power consumption. To separate the communication power consumption of Wi-Fi radio from the overall consumption of the device, we first measure the baseline power of the device. The client (Raspberry Pi and S10) is associated with the AP but does not have any traffic except periodically receiving the beacons from the AP. Both devices are using the default power saving mode (PSM). A Raspberry Pi consumes 1440 mW, while an S10 smartphone costs 377 mW. We consider these values as non-communication power and substrate them from the overall power to obtain the communication power of a Wi-Fi radio.

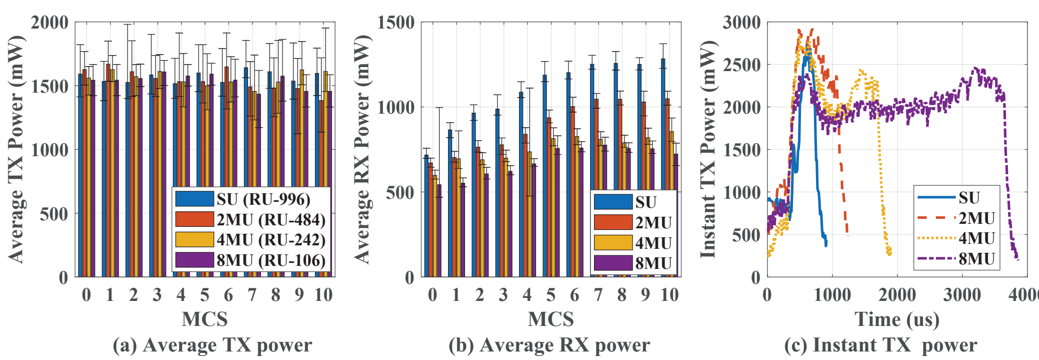

Fig. 11. OFDMA power profile (S10): (a) TX power, (b) RX power, (c) Instant TX power vs. RU size.

6.2.2 OFDMA power profile. The section examines the impact of RU size (allocated bandwidth) on power consumption. We obtain both the transmission and reception power consumption of four representative RU configurations: (1) SU: the entire 80 MHz bandwidth (996 RU) is allocated to one client, (2) 2MU: 996 RU are equally shared by 2 clients. Each client uses around 40 MHz, (3) 4MU: four 242 RU blocks allocated to 4 clients, and (4) 8MU: eight 106 RU blocks allocated to 8 clients. The results of Galaxy S10 are depicted in Fig. 11.

TX Power. Fig. 11(a) compares the average transmission power consumption of four RU sizes. **Critically, we observe the average power consumption of the transmitter does not decrease proportionally to the allocated bandwidth.** On the contrary, the power consumption of S10 is not notably reduced when it is allocated to a smaller bandwidth. For all the combinations of bandwidth and MCS, we observe that the transmitter of S10 consumes a constant 1455 mW. Fig. 11(c) further demonstrates examples of the instant power of the radio transmitting on various bandwidths. Comparing OFDMA settings (e.g., 8MU) with the single-user one (SU), the instant power does not reduce significantly, whereas the transmission time increases proportionally with the number of clients due to the reduced data rate. **As a result, OFDMA could dramatically increase the overall energy to transmit a packet, imposing a significant challenge for power consumption**. We also build the profile for the latest AX210 PCI-E card (Fig. 20 in Appendix). The instant power is slightly reduced for smaller bandwidth but it also does not decrease proportionally.

RX Power. The average reception power consumption is plotted in Fig. 11(b). We found that the power consumed for reception is reduced when operating on a smaller bandwidth. For example, receiving on the entire 80 MHz (996 tones) with MCS 9 consumes 1270 mW, while it is reduced to 800 mW (0.63×) when receiving on the 484 tones (0.5× of the bandwidth). This can help mitigate the power consumption for OFDMA clients. However, the speed of power consumption decreases is still slower than the reduction of data rate (e.g., 0.63× of the power vs. 0.5× of the data rate). In addition, when further reducing RU size to 242 and 106 tones, the decreases become marginal. For example, receiving on 106 tones costs 720 mW, which only reduces 10% compared to 484 tones but the data rate decreases by 75%. Therefore, OFDMA will still increase overall RX energy consumption.

Overall, the power profile implies that OFDMA could lead to an increase in power consumption. The detailed benchmark of overall energy efficiency will be given in Section 6.2.4.

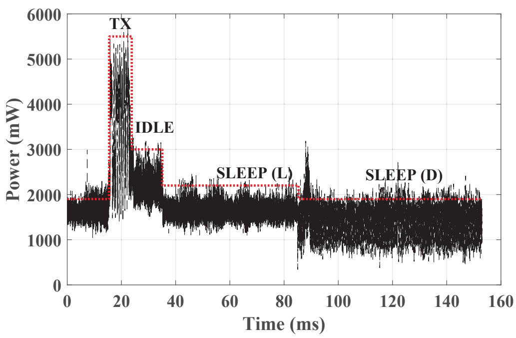

Fig. 12. Example timeline of TWT power states.

	S10	Pi (AX210)
TX (SU MCS9)	1536	2403
TX (2MU MCS9)	1477	2037
TX (8MU MCS9)	1455	1592
RX (SU MCS9)	1250	1881
RX (2MU MCS9)	1029	1436
RX (8MU MCS9)	754	978
IDLE_LISTEN	365	818
SLEEP(LIGHT)	N/A	266
SLEEP(DEEP)	0	0

Table 4. Radio states and powers (mW).

6.2.3 *TWT power profile.*

Target wake time (TWT) allows a client to be active only when it has data to transmit or receive, thus saving energy. Fig. 12 and Table 4 illustrate the power consumed in various TWT states. Note that for better depicting the instant power, Fig. 12 plots the raw power (i.e., includes the baseline power consumption of 1440 mW) which need to be subtracted to yield the values in Table 4. In specific, Wi-Fi radio wakes up at the negotiated time slot and completes the data exchange with AP. It remains in idle mode until the TWT session finishes and the radio enters light sleep mode. After a specific time (51 milliseconds in AX210), a timeout is triggered and the radio goes to deep sleep. As the profile shows, the power consumption in both sleep modes is significantly lower than in active states (TX/RX and idle). Therefore, by allocating the clients to different TWT time slots, we could dramatically reduce power consumption.

6.2.4 *Energy Efficiency.*

In Section 6.2.2 and 6.2.3, we observe that new channel access schemes (i.e., OFDMA and TWT) in Wi-Fi 6 could have a significant impact on instant power consumption. The section presents the overall energy efficiency, i.e., the energy (in Joule) consumed to transmit or receive one bit. We evaluate two representative scenarios: intensive traffic and short transmission while demystifying the impact of various factors e.g., access schemes, and the number of clients.

Intensive traffic. First, we measure the energy efficiency of traffic-intensive applications. To produce intensive traffic, the Wi-Fi is configured with a large frame aggregation (AMPDU size = 10). `iperf3` generates saturated traffic between AP and clients. Three experiments are conducted using OFDMA, TWT, and legacy CSMA/CA settings. Fig. 13 (a) compares TX energy efficiency of Galaxy S10. For illustration purposes, we focus on the results of MCS 9, while the conclusion is consistent for other MCS. **Our key observation is that the energy consumption of OFDMA dramatically increases (approximately linearly) with the number of clients transmitting in parallel.** For example, a single-user transmission consumes 5.1 nJ/bits whereas the consumption rises to 30 nJ/bits for 8MU OFDMA (by 6×). This coincides with our observation in Section 6.2.2 that with

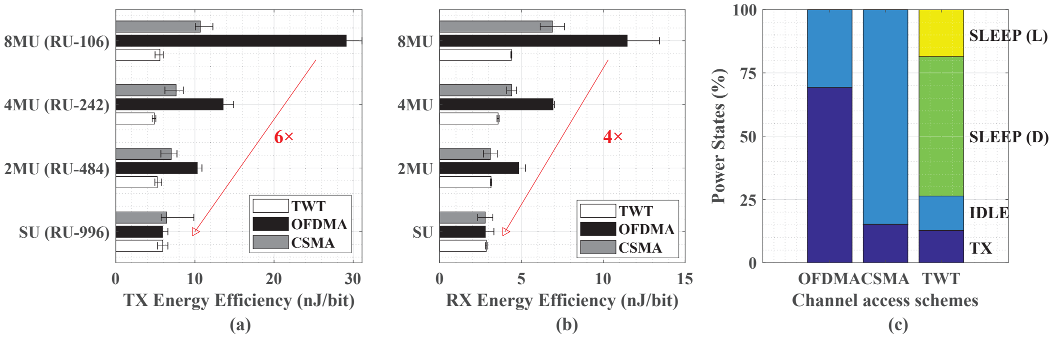

Fig. 13. Energy efficiency of traffic-intensive scenarios: OFDMA vs. TWT vs. CSMA/CA, (a) UL energy efficiency, (b) DL energy efficiency, (c) The percentage of time in each radio states (8MU).

more clients (and thus a smaller RU size for each client), the time to finish transmission increases proportionally, while the instant power does not notably drop. **Furthermore, OFDMA costs significantly more energy than TWT and legacy CSMA/CA.** The consumption of CSMA/CA (grey bar) shows moderate increases with a larger number of clients. The extra cost mainly comes from a growing channel contention among clients. Finally, the consumption of TWT remains stable with the number of clients because each client uses separate time slots and thus does not contend. The results show that TWT is the most energy efficient in the multi-client scenario, while OFDMA imposes challenges on the energy efficiency of Wi-Fi 6 devices. This result confirms our hypothesis that OFDMA leads to the reduction of overall energy efficiency for intensive traffic.

Similar results are obtained for RX shown in Fig. 13 (b). The overall reception power consumption increases with the number of clients. We note that the growth in RX is slightly less significant than TX (e.g., 4× in 8 MU). This is because when operating in a smaller RU size, the instant reception power is reduced more than the transmission power, as shown in Section 6.2.2. To further break down the energy consumption of three channel access schemes, we obtain the percentage of time the radio spent in various power states. The statistical results of the 8-client setting are depicted in Fig. 13 (c). The radio in the OFDMA scheme spends 70% time in active modes (TX and RX). In contrast, using TWT and CSMA/CA, the radio is active only during 13% of the time and TWT further allows the radio to sleep 75% of the time. The remarkably longer active time leads to the notably higher energy consumption of OFDMA.

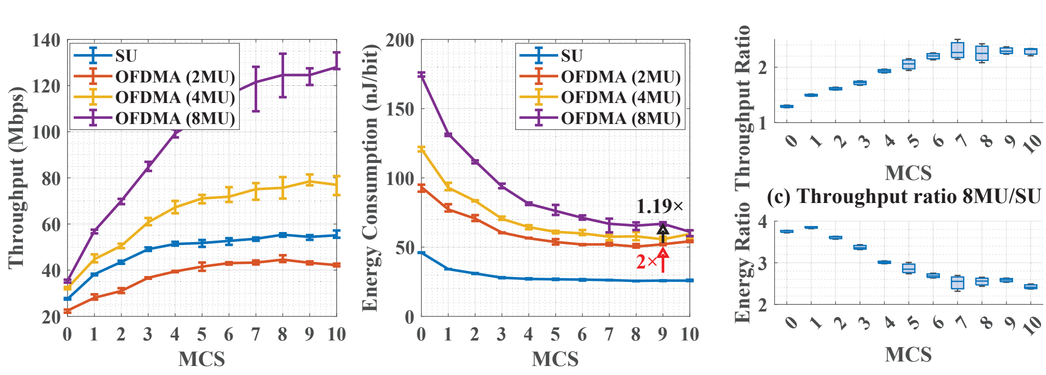

Fig. 14. Energy efficiency of short transmissions: (a) Throughput, (b) Energy efficiency, (c) Impact of MCS on throughput gain (d) Impact of MCS on energy efficiency.

Short transmission. Our throughput benchmark in Section 4 shows that OFDMA effectively reduces the communication overhead of the short transmission (≤ 1.6Kbytes). We examine if the gain in spectrum efficiency can help compensate energy cost of OFDMA. In specific, We create the short transmission scenario by configuring the AMPDU size of the clients to 1. The aggregated RX throughput and energy efficiency of various OFDMA configurations and MCS are depicted in Fig. 14. **Our key observation is that OFDMA also increases the energy consumption of short transmission but the growth ratio is lower than intensive traffic.** As Fig. 14(a) shows, OFDMA improves the overall throughput (except 2MU settings). For example, 8MU provides up to 2.6× gain for MCS 9. On the other hand, energy consumption (depicted in Fig. 14(b)) increases from 25.8 nJ/bit to 66.8 nJ/bit (2.4×). Recall that 8MU incurs a 4× RX consumption in the traffic-intensive scenario. This result demonstrates that the improved spectral efficiency for short transmission indeed mitigates the energy consumption cost of OFDMA. **The second important observation is that increasing the number of clients in the OFDMA transmission only introduces marginal extra consumption costs.** Specifically, as we point out in Fig. 14(b), switching from SU to 2MU OFDMA doubles energy consumption (from 25.8 nJ/bit to 55.6 nJ/bit) but a further increase to 8MU only costs 19% additional energy (from 55.6 nJ/bit to 66.8 nJ/bit). This is very distinct from the intensive traffic scenario where the energy grows approximately linearly with the number of clients. **Finally, we observe that OFDMA is inefficient for both throughput and energy consumption in the low modulation order.** Fig. 14 (c) depicts the throughput gain over various modulation orders indicated by the ratio of the 8MU and SU throughput. Fig. 14 (d) plots the relative energy consumption ratio. In low modulation order (MCS 0), the OFDMA throughput gain is trivial (1.2×) while it consumes as much as 3.7× more energy. This is because the time for payload transmission at a low MCS is very large (similar to intensive traffic). With an increasing MCS, the throughput gain gradually rises and the consumption ratio declines.

6.3 Large-scale network simulation

Our experiment shows that both throughput and energy consumption in OFDMA scheme increase with the number of clients in OFDMA transmissions. Due to the constrain of the current commodity AP, our empirical experiment can only provide the benchmark for 1 to 8 clients. Theoretically, an AP can communicate with up to 37 clients in an OFDMA frame. To estimate the trend in a large-scale network, we conduct a simulation.

6.3.1 Simulation setup. We develop a simulator using MATLAB WLAN Toolbox [6]. The toolbox can generate MAC frames and physical-layer signals. We simulate both short transmission (with 200 bytes payload) and more intensive traffic (with 8 KBytes payload) and obtain the throughput. The power consumption model of Galaxy S10 is used to calculate energy efficiency. More details of the simulator setup and validation are given in Appendix B.

6.3.2 Simulation results. **The results of the large-scale simulation are consistent with our empirical experiment**. The first row of Fig. 15(a-d) shows the simulation results of short uplink transmissions. OFDMA demonstrates a significant gain for aggregated throughput at a cost of increased power consumption. As depicted in Fig. 15(a), the overall throughput of MCS 10 rises from 10 Mbps to 129 Mbps when the number of clients (denoted as MU size) grows from 1 to 37, providing a 13× gain. On the other hand, the energy consumption demonstrated in Fig. 15(b) increases from 141 nJ/bit to 355 nJ/bit (2.5×). The relative OFDMA throughput gain and energy consumption growth compared to SU are further depicted in Fig. 15(c) and (d). We obtain the same trend as our empirical study that the throughput is dramatically improved with an increasing number of OFDMA clients whereas the extra energy consumption paid for serving more clients is marginal. The exception is the low modulation order (e.g., MCS 0) in which OFDMA is both

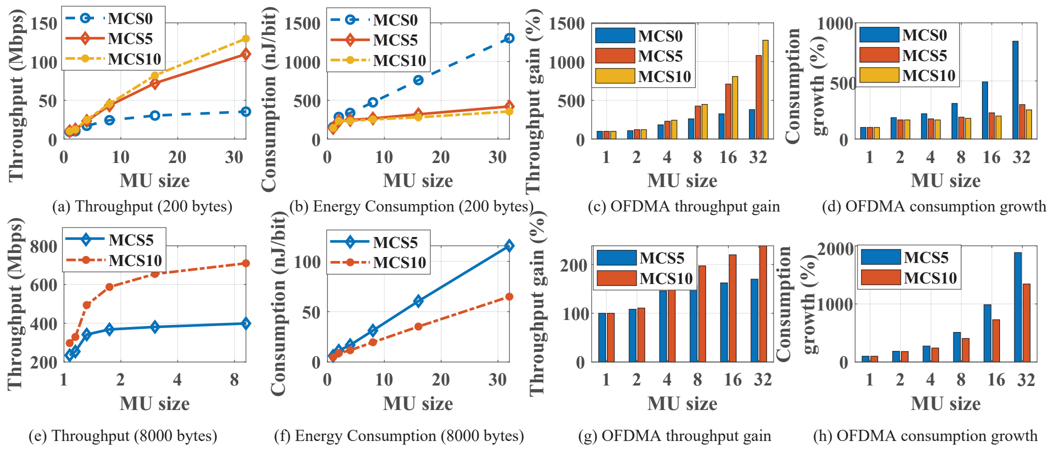

Fig. 15. Throughput and consumption in a large-scale network: (a-d) Results of short transmission: OFDMA improves the aggregated throughput by up to 11× at the cost of 2.5-8.4 × consumption, (e-h) results of intensive traffic: OFDMA improves the aggregated throughput by up to 2.3× at the cost of 18× consumption.

inefficient in throughput and energy. The results show that in general serving a large number of clients with short transmission using OFDMA (like IoT data) causes a moderate increase in energy consumption (up to 2.5×) but can offer a significant advantage to the throughput.

The results for more intensive traffic (the second row of 15 (e-h)) demonstrate that OFDMA incurs severe energy consumption costs on a large-scale network. For example, serving 32 traffic-intensive clients leads to 13× growth of energy consumption at MCS 10 and 18× growth at MCS 5. In contrast, the benefit of throughput is very limited (up to 2× for MCS 10 and 1.5× MCS 5). We also simulate the extremely intensive traffic pattern (e.g., mobile VR) with 16KBytes payload size. OFDMA can increase the energy cost by 30×. Therefore, using OFDMA to serve a dense network of high traffic intensive imposes challenges on the energy consumption of devices.

6.4 Implication to applications

Our results indicate new MAC schemes of Wi-Fi 6 have a significant impact on the energy consumption of the clients. In this section, we discuss the implication to various applications and the tradeoff that a Wi-Fi scheduling algorithm should make between throughput, latency, and energy.

VOIP and IoT. As discussed in Section 4.3, VOIP and low-power IoT applications feature short transmissions and a large number of clients. Scheduling the traffic with OFDMA improves the overall throughput at a cost of up to 2.5× more energy consumption, which could shorten the battery life. Yet, this result is based on first-generation Wi-Fi devices. We believe there are opportunities to bring down power consumption. Furthermore, from the energy consumption perspective, the scheduler could allocate as many clients with short packets as possible into an OFDMA transmission because the extra energy cost for serving more clients is moderate. In addition, using TWT is more energy efficient for IoT devices. However, it does not offer the benefit to the overall throughput.

Streaming. Video calls, mobile VR, and video streaming (e.g., YouTube) generate intensive traffic. Using the OFDMA to serve these applications impose challenges on the battery life, especially when the number of clients is large. From the energy consumption perspective, TWT is a more energy-efficient option for streaming because a client can sleep when other intensive devices are served. Yet, TWT offers no gains to spectral efficiency and it also cannot alleviate the latency due to the competition with external stations. As a result, the optimal scheduler could combine OFDMA and TWT to balance various criteria. For example, we can use TWT to wake up a carefully chosen

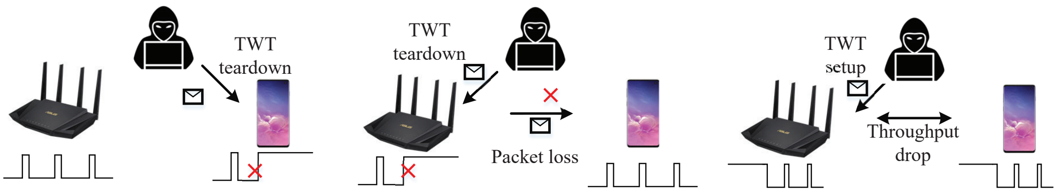

(a) Battery depletion attack (b) Packet loss attack (c) Throughput throttle attack

Fig. 16. Three TWT Attack models.

number of clients in a specific time slot for OFDMA communication, which both improves spectrum efficiency and restricts power consumption. We will apply this idea in the case study (Section 9).

6.5 Takeaway

New channel access schemes have a non-trivial impact on power consumption. While improving the overall throughput, OFDMA can incur severe energy consumption costs. On the other hand, TWT is more energy saving but it offers no throughput gain. Finally, the energy consumption and throughput tradeoff should be considered depending on the characteristics of an application.

7 SECURITY

7.1 Attack model

We demonstrate in Section 6.2.3 that TWT schedules clients into individual time slots, which reduces the energy consumption of the device as well as the channel contention. Although effective, we observe that the TWT negotiation protocol in the commodity Wi-Fi devices is vulnerable to various types of adversarial attacks that could cause severe issues (e.g., energy consumption, spectrum fairness, and network stability). As Fig. 16 depicts, we present three critical attack models.

Battery depletion attack. Low-power IoT devices establish TWT sessions with AP so that it performs a low-duty cycle and only wakes up in the negotiated time slot. However, as depicted in Fig.16(a), an adversary can impersonate the AP and send a TWT teardown message to the client. As a result, IoT devices have to stay awake and consume significantly more energy.

Packet loss attack. In TWT mode, the client wakes up at the negotiated time slot to receive messages from the AP. As shown in Fig.16(b), an adversary can impersonate the client and tear down the TWT session with the AP. The AP assumes that the client will stay awake and can receive messages at any time while the client is unaware and continues to perform the duty cycle. In consequence, downlink packets will drop when the client is in low-power sleep, causing severe packet loss and even network disconnection.

Throughput throttle attack. TWT specifies the duration during which a client can communicate with AP, thus constraining its maximum throughput. An adversary may leverage this to limit the throughput of other clients and obtain an unfair share of the spectrum. We depict the attack in Fig.16(c), where the attacker pretends to be the client and sets up a TWT session with an extremely low duty cycle. As the result, the throughput of the client is throttled.

7.2 Methodology

We implement the adversary attack using commodity Wi-Fi devices and test the vulnerability of our testbed. In specific, we use the monitor mode of a commodity Wi-Fi radio to sniff the TWT setup message among AP and client. These action frames are not encrypted so that the attacker can figure out the critical information of TWT session (e.g., the dialog ID). We use the information to build a spoofed TWT teardown or setup frame and send it using packet injection method. It turns out that the attacks that impersonate the client (i.e., packet loss and throughput throttle) are straightforward to achieve because the AP constantly waits for packets. In contrast, impersonating the AP is more challenging because the client in TWT scheme is in sleep mode for the majority of the time and the attacker needs to figure out when the client is awake. We predict the client's state using the timing information (offset, interval, duration of TWT session) included in the TWT setup message and send multiple attack frames to increase the success probability.

7.3 Results and Defence

Our experiment on the testbed proves that three proposed attack models are feasible on the commodity Wi-Fi 6 devices. Fig. 17 shows an example of throughput throttle attack. The attacker manages to change the traffic pattern and control the throughput of the victim. Interestingly, we observe although the spoofed TWT setup is sent to the AP, both throughput from downlink (from AP) as well as uplink from the client are throttled. This is because the protocol allows the TWT session to be initiated by AP (i.e. sending a TWT setup response without a request from the client). Thus, when the AP receives the malicious TWT setup request, it replies with a setup acceptance response, which naturally misguides the client to perform TWT in the uplink. For battery depletion attack, we manage to produce the forged TWT teardown frame that changes the energy consumption pattern of the client. It can dramatically increase the average consumption (from 50mW to 900mW) and shortens the battery life. The packet loss attack is also verified. The severe downlink packet loss leads to temporary device disconnection. We discuss a few solutions to defend against the proposed attack and the limitations.

Protected TWT. The most straightforward way is to encrypt the TWT action frames such that the attacker cannot forge the setup and teardown messages. This can be done by enforcing IEEE 802.11w action frame protection on the TWT messages. Specifically, an additional message integration code (MIC) is produced using the established key and added to TWT setup and teardown frames, which prevents forging. Although effective, this method has a few limitations. First, the encryption key is established through EAPOL 4-way shake hands, which requires the network to use WPA or WPA2. Therefore, it does not work for the open network (e.g., passwordless Wi-Fi in airports, shopping mall, hotel, etc.). Furthermore, it has been shown that IEEE 802.11w may cause compatibility issues to unsupported devices and is often disabled by the AP [28].

Anomaly detection. The AP and client can collaboratively detect abnormal TWT action frames. For example, the AP can detect a battery depletion attack if it overhears a TWT teardown that impersonates its identity. Similarly, the smartphone might overhead the packet loss attack and report to the AP. Yet, these methods fail when the attack is out of the range of the AP or client. Finally, the client might monitor its traffic pattern to detect potential throughput throttle attacks.

8 APPLICATION: WI-FI SLICING

The section presents a case study of Wi-Fi design by exploring the key insights we obtained from this measurement study - we develop Wi-Fi network slicing [17] using OFDMA and TWT. In specific, network slicing aims at producing multiple logical networks tailored to the different user requirements over a common infrastructure. The key expectations of network slicing are resource isolation (the resources of each slice are isolated) and Quality of Service (every device achieves guaranteed QoS). Fig.18 shows an example of combining OFDMA and TWT to achieve this goal. Specifically, we create a slice for IoT services (e.g., wireless sensors and smart meters) that features a large number of devices with short transmissions and a streaming slicing that produces intensive traffic (e.g., mobile VR, video conference, and Wi-Fi display). To achieve resource isolation, we use TWT to wake up IoT and streaming devices in a group of disjoint service periods. Within the service period, the devices are allocated to RU size using OFDMA to eliminate the contention and improve the quality of service. **We explore the unique throughput, latency, and energy**

consumption characteristics of IoT and streaming slices obtained from the measurement. Specifically, due to the short transmission, IoT devices suffer significantly from communication overhead. To improve the spectral efficiency of IoT slices, we schedule as many IoT devices as possible into the same TWT service period. Since we show that increasing the number of short transmissions into an OFDMA transmission only causes moderate power consumption growth, we can both provide high spectral efficiency and ensure that the battery life of IoT devices is acceptable. On the other hand, streaming devices require both low latency as well as energy efficiency. We need consider a trade-off between the two metrics because increasing the number of intensive traffic into an OFDMA transmission reduces the channel access delay but could severely increase the power consumption. In order to reduce the latency while being energy efficient, we schedule the streaming devices into multiple TWT service periods (as shown in Fig.18) to reduce the contention while limiting the number of clients in each service period to restrict power consumption.

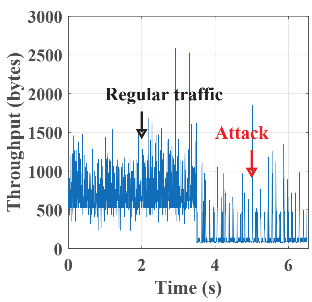

Fig. 17. Throughput throttle.

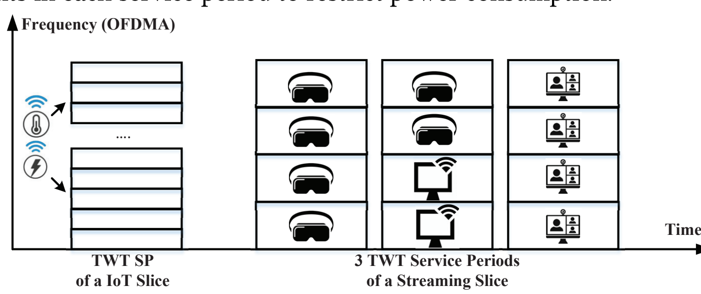

Fig. 18. Wi-Fi slices combine OFDMA and TWT for traffic isolation, high spectral efficiency (IoT), and latency-energy tradeoff (Streaming).

9 DISCUSSION AND FUTURE WORK

Admittedly, the commercial products tailored to W-Fi 6 technology are still in the early stage. Therefore, our measurement is restricted by the current implementation of the available devices. The section discusses the limitations of this work and points out future directions.

OFDMA client scale. First, the maximum OFDMA client number supported by our AP is 8 as opposed to a maximum of 74 clients defined by the standard. Thus, our empirical experiment focuses on 8 client scenarios while we combine these empirical results with simulations to estimate the performance on a larger scale in Section 6.3.

Broadcast TWT and Wi-Fi slicing. Second, Wi-Fi 6 proposes individual and broadcast TWT. However, our extensive experiment shows that current broadcast TWT implementation in commodity Wi-Fi devices suffers from severe compatibility issues, leading to packet loss and network disruption. Therefore, this study only reports the measurement of individual TWT, while we are actively working with the open-source community to address the compatibility issues. In addition, the broadcast TWT is critical to implement Wi-Fi slicing proposed in Section 8 (e.g., group and allocate a various number of clients into time slots). Due to the unavailability of broadcast TWT, we leave the implementation and evaluation of slicing in future work.

Bandwidth and MCS. Third, we mainly investigate the performance with an 80 MHz bandwidth setting due to the limitation in the Ethernet backhaul. However, the conclusion in the study is generically applicable to other bandwidths. Additionally, to control the number of impacting factors, our experiment constraints a few parameters. For example, we force all clients to use the same MCS in an OFDMA transmission. In the future, we plan to investigate the performance when the clients are allocated with diverse modulation orders. Furthermore, we mainly focus on scenarios with good signal quality. The distance between the AP and clients is 3 meters with a line-of-sight path. We will deploy the testbed in more practical environments (e.g., scatter the device in the office).

Hybrid clients. Finally, this study mainly considers homogeneous Wi-Fi networks, i.e., all Wi-Fi clients of an AP use Wi-Fi 6. To further reveal the real-world performance, the coexistence of Wi-Fi 6 and legacy clients should be investigated. In addition, our WiFi slicing design in the ISM bands (e.g., 2.4 and 5 GHz) also needs to carefully take the legacy clients into consideration.

10 RELATED WORK

There is a large body of work on the performance benchmark of previous Wi-Fi versions (i.e., 802.11n [19, 30, 32, 33] and 802.11ac [31, 37], 802.11ad [7, 8]). The impact of bandwidth, frame aggregation, and MIMO on 802.11n throughput is examined in [32]. The authors of [37] studied the impact of the large channel width of 802.11ac on energy efficiency and interference. The power consumption and throughput of 802.11n/ac smartphone are measured in [31]. Per-frame energy consumption components in Wi-Fi devices and power consumption models are studied in [18]. All the measurements conducted on the legacy Wi-Fi edition only examine the performance of *single-user* transmission. In contrast, our study focuses on the *multi-client* transmission enabled by new channel access mechanisms of Wi-Fi 6 (e.g., OFDMA and TWT). A recent 802.11ad measurement [7] compares the energy efficiency of 802.11ad and 802.11ax. However, the study only reports the result of the consumption of single-user transmission. By conducting extensive multi-client experiments, our study reveals that new channel access approaches in Wi-Fi 6 have a significant impact on throughput, latency, and power consumption. Finally, the spoofing attack has been extensively investigated for previous Wi-Fi protocols [34, 35]. Our work is the first to showcase the possibility to conduct TWT-based attacks on real devices by leveraging new Wi-Fi 6 features.

So far, the Wi-Fi 6 performance evaluations are mostly conducted via analytical models [11, 20, 24, 25, 27] and simulations [10, 13, 14, 25, 36]. The authors of [10] simulate OFDMA performance using ns-3 simulator and report a 10% throughput gain and 4 ms latency reduction. Yet, our experiment shows that simulations have several limitations. First, the simulator does not have several critical characteristics in real Wi-Fi devices (e.g., low-level frame aggregation and increased interframe spacing). Second, these simulations do not produce real-world channel contention from external competitors and among Wi-Fi 6 clients. Therefore, the results from simulations cannot fully represent the performance in the real network. Last but not least, we cannot capture the energy consumption in commodity devices using simulation or analytical analysis. The authors of [26] evaluate the empirical throughput and jitter of Wi-Fi 6 network, and show that Wi-Fi 6 achieves more efficient channel utilization than IEEE 802.11ac. However, they only test single-user scenarios and thus did not study the effectiveness of OFDMA to multi-user networks.

To our best knowledge, our paper provides the first comprehensive testbed-based evaluation of Wi-Fi 6 with a focus on multi-user performance and the impact of key features of this new Wi-Fi standard (i.e., OFDMA and TWT). We are also the first to report energy consumption concerns and security vulnerabilities on COTS devices. We envision that our results will provide valuable insights into Wi-Fi system design and inspires more research works for next-generation WLANs.

11 CONCLUSION

To sum up, we present the first comprehensive performance evaluation of real Wi-Fi 6 network. With the proposed novel measurement methods for multi-user scenarios, we benchmark the throughput, and latency gains and then report the challenges of energy consumption and security. We believe that the insights obtained from this study provide guidance for future Wi-Fi system design.

12 ACKNOWLEDGEMENTS

This work was supported by Nokia Bell Labs summer internship program. We sincerely thank the shepherd Prof. Xinyu Zhang and anonymous reviewers for their valuable comments.

REFERENCES

[1] 2022. A labeled dataset with malicious and benign IoT network traffic . https://www.stratosphereips.org/datasets-iot23.
[2] 2022. Gigabit Wireless Router for Security-savvy Travelers. https://www.gl-inet.com/products/gl-ar750s/.
[3] 2022. Tcpreplay - Pcap editing and replaying utilities. https://tcpreplay.appneta.com/.
[4] 2022. The Most Colorful Nature in Real 4K HDR - Incredible Tropical Animals and Relaxing Music - 2022 . https://youtu.be/17ntdUP5-Do.
[5] 2022. VR traffic characterization. https://www.eecis.udel.edu/~salehi/vr.html.
[6] 2022. WLAN Toolbox Simulate, analyze, and test WLAN communications systems. https://www.mathworks.com/products/wlan.html.
[7] Shivang Aggarwal, Moinak Ghoshal, Piyali Banerjee, Dimitrios Koutsonikolas, and Joerg Widmer. 2021. 802.11 ad in smartphones: Energy efficiency, spatial reuse, and impact on applications. In *IEEE INFOCOM 2021-IEEE Conference on Computer Communications*. IEEE, 1–10.
[8] Shivang Aggarwal, Srisai Karthik Neelamraju, Ajit Bhat, and Dimitrios Koutsonikolas. 2022. A Detailed Look at MIMO Performance in 60 GHz WLANs. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 6, 2 (2022), 1–26.
[9] ASUS. 2022. AX3000 Dual Band WiFi 6 (802.11ax) Router supporting MU-MIMO and OFDMA technology, with AiProtection Pro network security powered by Trend Micro™, compatible with ASUS AiMesh WiFi system. https://www.asus.com/us/Networking-IoT-Servers/WiFi-Routers/ASUS-WiFi-Routers/RT-AX58U/.
[10] Stefano Avallone, Pasquale Imputato, Getachew Redieteab, Chittabrata Ghosh, and Sumit Roy. 2021. Will OFDMA improve the performance of 802.11 WiFi networks? *IEEE Wireless Communications* 28, 3 (2021), 100–107.
[11] Dmitry Bankov, Andre Didenko, Evgeny Khorov, and Andrey Lyakhov. 2018. OFDMA uplink scheduling in IEEE 802.11 ax networks. In *2018 IEEE International Conference on Communications (ICC)*. IEEE, 1–6.
[12] Compex. 2022. Qualcomm IPQ6018 Cypress Series Dual Band Dual Concurrent MU-MIMO OFDMA. https://compex.com.sg/qualcomm-ipq6018-cypress-series-dual-band-dual-concurrent//.
[13] Vinicius Da Silva Goncalves and Edward W. Knightly. 2022. An Experimental Study of Triggered Multi-User Uplink Access with Real Application Traffic. In *2022 IEEE/ACM 30th International Symposium on Quality of Service (IWQoS)*. 1–10. https://doi.org/10.1109/IWQoS54832.2022.9812876
[14] Yousri Daldoul, Djamal-Eddine Meddour, and Adlen Ksentini. 2020. Performance evaluation of OFDMA and MU-MIMO in 802.11 ax networks. *Computer Networks* 182 (2020), 107477.
[15] Dell. 2022. Dell™ OptiPlex™ 790 desktop. https://i.dell.com/sites/csdocuments/Shared-Content_data-Sheets_Documents/en/us/optiplex-790-customer-brochure.pdf/.
[16] EKA. 2022. EKA1080p power monitor. http://www.emkia.com/.
[17] Salah Eddine Elayoubi, Sana Ben Jemaa, Zwi Altman, and Ana Galindo-Serrano. 2019. 5G RAN slicing for verticals: Enablers and challenges. *IEEE Communications Magazine* 57, 1 (2019), 28–34.
[18] Andres Garcia-Saavedra, Pablo Serrano, Albert Banchs, and Giuseppe Bianchi. 2012. Energy consumption anatomy of 802.11 devices and its implication on modeling and design. In *Proceedings of the 8th international conference on Emerging networking experiments and technologies*. 169–180.
[19] Daniel Halperin, Ben Greenstein, Anmol Sheth, and David Wetherall. 2010. Demystifying 802.11 n power consumption. In *Proceedings of the 2010 international conference on Power aware computing and systems*. USENIX Association, 1.
[20] Yein Heo, Jonggyu Jang, Yeongjun Kim, and Hyun Jong Yang. 2020. Performance Comparison of SU- and MU-MIMO in 802.11ax: Delay and Throughput. In *2020 International Conference on Information and Communication Technology Convergence (ICTC)*. 879–882. https://doi.org/10.1109/ICTC49870.2020.9289385
[21] Tim Higgins. 2022. Does OFDMA Really Work? https://www.smallnetbuilder.com/wireless/wireless-features/33222-does-ofdma-really-work-part-1.
[22] Intel. 2022. Intel® Wi-Fi 6E AX210. https://www.intel.com/content/www/us/en/products/sku/204836/intel-wifi-6e-ax210-gig/specifications.html/.
[23] Evgeny Khorov, Anton Kiryanov, Andrey Lyakhov, and Giuseppe Bianchi. 2019. A Tutorial on IEEE 802.11ax High Efficiency WLANs. *IEEE Communications Surveys Tutorials* 21, 1 (2019), 197–216. https://doi.org/10.1109/COMST.2018.2871099
[24] Evgeny Khorov, Vyacheslav Loginov, and Andrey Lyakhov. 2016. Several EDCA parameter sets for improving channel access in IEEE 802.11ax networks. In *2016 International Symposium on Wireless Communication Systems (ISWCS)*. 419–423. https://doi.org/10.1109/ISWCS.2016.7600940
[25] Davide Magrin, Stefano Avallone, Sumit Roy, and Michele Zorzi. 2021. Validation of the ns-3 802.11 ax OFDMA Implementation. In *Proceedings of the Workshop on ns-3*. 1–8.
[26] Siraj Muhammad, Jiamiao Zhao, and Hazem H. Refai. 2021. An Empirical Analysis of IEEE 802.11 ax. In *2020 International Conference on Communications, Signal Processing, and their Applications (ICCSPA)*. 1–6. https://doi.org/10.1109/ICCSPA49915.2021.9385748

[27] Gaurang Naik, Sudeep Bhattarai, and Jung-Min Park. 2018. Performance Analysis of Uplink Multi-User OFDMA in IEEE 802.11ax. In *2018 IEEE International Conference on Communications (ICC)*. 1–6. https://doi.org/10.1109/ICC.2018.8422692
[28] Extreme Network. 2022. Client unable to connect to SSID set to 802.11w mandatory/required mode. https://extremeportal.force.com/ExtrArticleDetail?an=000105816.
[29] Maddalena Nurchis and Boris Bellalta. 2019. Target wake time: Scheduled access in IEEE 802.11 ax WLANs. *IEEE Wireless Communications* 26, 2 (2019), 142–150.
[30] Konstantinos Pelechrinis, Theodoros Salonidis, Henrik Lundgren, and Nitin Vaidya. 2010. Experimental characterization of 802.11 n link quality at high rates. In *Proceedings of the fifth ACM international workshop on Wireless network testbeds, experimental evaluation and characterization*. 39–46.
[31] Swetank Kumar Saha, Pratik Deshpande, Pranav P Inamdar, Ramanujan K Sheshadri, and Dimitrios Koutsonikolas. 2015. Power-throughput tradeoffs of 802.11 n/ac in smartphones. In *2015 IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 100–108.
[32] Vivek Shrivastava, Shravan Rayanchu, Jongwoon Yoonj, and Suman Banerjee. 2008. 11 n under the microscope. In *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*. 105–110.
[33] Dionysios Skordoulis, Qiang Ni, Hsiao-Hwa Chen, Adrian P Stephens, Changwen Liu, and Abbas Jamalipour. 2008. IEEE 802.11 n MAC frame aggregation mechanisms for next-generation high-throughput WLANs. *IEEE Wireless Communications* 15, 1 (2008), 40–47.
[34] Mathy Vanhoef. 2021. Fragment and Forge: Breaking {Wi-Fi} Through Frame Aggregation and Fragmentation. In *30th USENIX Security Symposium (USENIX Security 21)*. 161–178.
[35] Mathy Vanhoef, Prasant Adhikari, and Christina Pöpper. 2020. Protecting wi-fi beacons from outsider forgeries. In *Proceedings of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks*. 155–160.
[36] Kaidong Wang and Konstantinos Psounis. 2020. Efficient scheduling and resource allocation in 802.11 ax multi-user transmissions. *Computer Communications* 152 (2020), 171–186.
[37] Yunze Zeng, Parth H Pathak, and Prasant Mohapatra. 2014. A first look at 802.11 ac in action: Energy efficiency and interference characterization. In *2014 IFIP Networking Conference*. IEEE, 1–9.

A GLOSSARY

ACRONYMS

AMPDU Aggregate MAC Protocol Data Unit. 6

COTS commercial-off-the-self. 1
CSMA/CA Carrier-sense multiple access with collision avoidance. 3

MCS Modulation and Coding Scheme. 6, 13
MPDU MAC Protocol Data Unit. 6

OFDMA orthogonal division multiple access. 1

TWT target wake time. 1

B SIMULATION SETUP AND VALIDATION

In order to examine the performance of Wi-Fi 6 network with more than 8 clients, we develop a simulator using MATLAB WLAN Toolbox. The toolbox provides functions that configure OFDMA RU allocation (`wlanHEMUConfig`) and produce various types of Wi-Fi 6 physical-layer signals (including HE UL/DL OFDMA) based on the RU allocation, payload size, and modulation scheme. Specifically, we assume the bandwidths are equally shared by the clients. Simulating the MAC procedure of OFDMA is straightforward because all the communications are scheduled centrally by the AP. Using `wlanWaveformGenerator`, we produce the physical-layer signal of the data frame, interframe spacing, and acknowledgment. Based on the signal, overall airtime is obtained and throughput is calculated using the payload size and airtime. In addition, we figure out the power consumption of the devices by recording the time it spends in each power state (e.g., TX, RX, and

IDLE) and we use the power consumption values of Galaxy S10 to obtain the results. For small RU sizes (e.g., 52 and 26), we assume that the device consumes a constant TX power of 1455 mW.

To validate the effectiveness of the simulator, we simulate the network with 1, 2, 4, and 8 clients and compare the results with our empirical measurement. Fig.19 (a) and (b) depict simulated OFDMA throughput and throughput gain. The results are consistent with our measurements (e.g., Fig.7). For example, the overall throughput is doubled in OFDMA (8MU) setting. Note that the absolute throughput of the simulation is slightly higher than the real-world measurement because the simulations are free from interference and noise. Furthermore, the simulations demonstrate similar power consumption growth caused by OFDMA. For example, Fig.19 (c) shows that for 16 Kbytes Wi-Fi transmissions, OFDMA with 8 clients can lead to a 6× power consumption. Based on the results, we further use this simulation method to obtain the performance of networks with more than 8 clients.

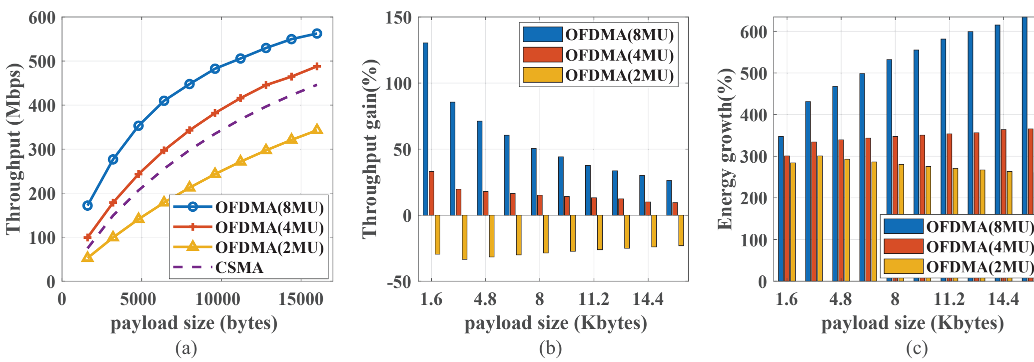

Fig. 19. Simulation results: (a) Throughput, (b) Throughput gain, (c) Energy consumption growth.

C OFDMA POWER PROFILE OF INTEL AX210

The power profile of Intel AX210 PCIE Wi-Fi card is shown in Fig.20.

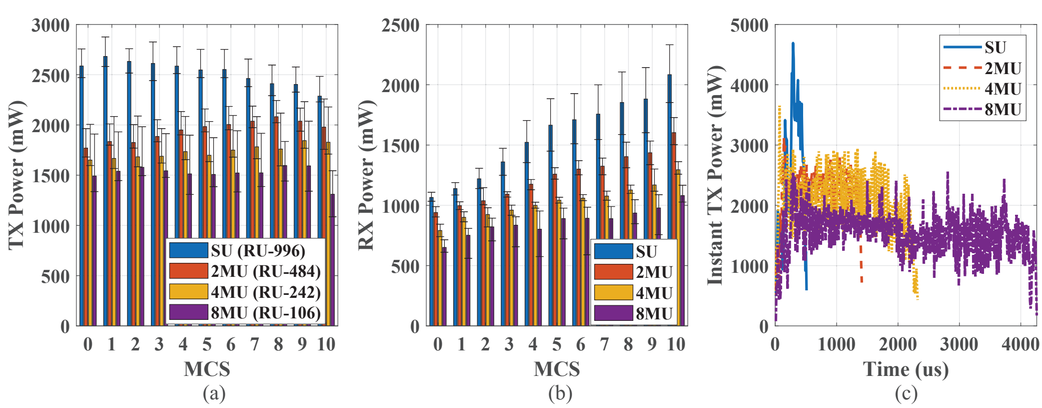

Fig. 20. OFDMA power profile (AX210): (a) TX power, (b) RX power, (c) Instant TX power vs. RU size.

Received October 2022; revised December 2022; accepted January 2023