Wi-Fi: Twenty-Five Years and Counting

Giovanni Geraci, Francesca Meneghello, Francesc Wilhelmi, David Lopez-Perez, Iñaki Val,
Lorenzo Galati Giordano, Carlos Cordeiro, Monisha Ghosh, Edward Knightly, and Boris Bellalta

***Abstract*—Today, Wi-Fi is over 25 years old. Yet, despite sharing the same branding name, today's Wi-Fi boasts entirely new capabilities that were not even on the roadmap 25 years ago. This article aims to provide a holistic and comprehensive technical and historical tutorial on Wi-Fi, beginning with IEEE 802.11b (Wi-Fi 1) and looking forward to IEEE 802.11bn (Wi-Fi 8). This is the first tutorial article to span these eight generations. Rather than a generation-by-generation exposition, we describe the key mechanisms that have advanced Wi-Fi. We begin by discussing spectrum allocation and coexistence, and detailing the IEEE 802.11 standardization cycle. Second, we provide an overview of the physical layer and describe key elements that have enabled data rates to increase by over 1,000×. Third, we describe how Wi-Fi Medium Access Control has been enhanced from the original Distributed Coordination Function to now include capabilities spanning from frame aggregation to wideband spectrum access. Fourth, we describe how Wi-Fi 5 first broke the one-user-at-a-time paradigm and introduced multi-user access. Fifth, given the increasing use of mobile, battery-powered devices, we describe Wi-Fi's energy-saving mechanisms over the generations. Sixth, we discuss how Wi-Fi was enhanced to seamlessly aggregate spectrum across 2.4 GHz, 5 GHz, and 6 GHz bands to improve throughput, reliability, and latency. Finally, we describe how Wi-Fi enables nearby Access Points to coordinate in order to improve performance and efficiency. In the Appendix, we further discuss Wi-Fi developments beyond 802.11bn, including integrated mmWave operations, sensing, security and privacy extensions, and the adoption of AI/ML.**

***Index Terms*—Wi-Fi, WLAN, IEEE 802.11, Wi-Fi Alliance, 802.11bn, unlicensed spectrum, PHY, MAC, multi-user technologies, energy savings, multi-link operation, multi-AP coordination.**

I. Introduction

You have probably downloaded this article over Wi-Fi. On September 15, 1999, the Wireless Ethernet Compatibility Alliance (WECA)—the future Wi-Fi Alliance [1]—unveiled the term Wi-Fi as the consumer-facing brand for the Institute of Electrical and Electronics Engineers (IEEE) 802.11 standard, ensuring that consumers never needed to learn to say "I triple E eight oh two dot eleven." Last year marked the 25th anniversary of this transformative technology [2], a milestone that underscores how deeply Wi-Fi has embedded itself into modern society. For Generation Z, the idea of a world without Wi-Fi is nearly unimaginable, and even millennials struggle to remember a time when a wired Ethernet cable was needed to get online [3]. Today, Wi-Fi is the most popular way to access the Internet across the globe. It is embedded in everything, from smartphones and laptops to televisions, smart appliances, and industrial sensors. With more than 21 billion Wi-Fi-enabled devices in use, and over 63% of global Internet traffic flowing through Wi-Fi, its ubiquity is both a technological achievement and a societal cornerstone [4].

Few technologies have had such far-reaching and sustained impact. Wi-Fi is not just a convenience. It is an essential utility that underpins the fabric of our digital lives. It has revolutionized how we work, learn, communicate, and entertain ourselves. It powers online classrooms in rural areas, enables telemedicine in urban hospitals, supports logistics in global supply chains, and fuels social interaction in every corner of the world. During the COVID-19 pandemic, Wi-Fi was nothing short of indispensable [5]: It enabled a rapid shift to remote work, kept students engaged in virtual classrooms, and allowed businesses to continue operating amid global disruption.

But the story of Wi-Fi is not only one of societal diffusion; it is also a narrative of relentless technological evolution. Since its inception in the late 1990s, when early versions of IEEE 802.11 offered data rates of just 1–2 Mb/s and basic wireless connectivity for laptops and personal digital assistants (PDAs), Wi-Fi has undergone continuous reinvention to keep pace with escalating user demands and emerging applications.

A. Use Case Evolution: From Browsing to Real-Time Critical Applications

Wi-Fi's early use cases centered on best-effort broadband access [6], supporting basic activities such as web browsing, music streaming, and downloading files. As computing and communication increasingly shifted to the cloud and user expectations evolved, Wi-Fi was drawn into more demanding roles [7]. Today, home and enterprise networks require multi-gigabit transmission rates and low-latency performance to support 4K/8K video, real-time collaborative tools, and secure remote access. The rise of hybrid work and cloud-native productivity platforms has further elevated Wi-Fi to the status of a mission-critical utility.

Beyond these mainstream productivity and media demands, Wi-Fi is now being stretched to meet the stringent performance requirements of a growing set of specialized, latency-sensitive, and mission-critical applications:

G. Geraci is with Nokia Standards and Universitat Pompeu Fabra, Spain. F. Meneghello is with Northeastern University, USA, and the University of Padova, Italy. F. Wilhelmi and B. Bellalta are with Universitat Pompeu Fabra, Spain. D. Lopez-Perez is with Universitat Politècnica de València, Spain. I. Val is with MaxLinear, Spain. L. Galati Giordano is with Nokia Bell Labs, Germany. C. Cordeiro is with Intel Corporation, USA. M. Ghosh is with the University of Notre Dame, USA. E. Knightly is with Rice University, USA.

This work was in part supported by the Spanish Research Agency through grants PID2021-123999OB-I00, PID2021-123995NB-I00, PID2024-156488OB-I00, CEX2021-001195-M, CNS2023-145384, and CNS2023-144333; by SGR 00955-2021 AGAUR; by ICREA Academia, by the European Union - Next Generation EU under the Italian National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.2, CUP C93C24004880002, project "CAMELIA"; by the Generalitat Valenciana, Spain, through grant CIDEXG/2022/17; by Horizon Europe SNS-JU through grant UNITY-6G 101192650; by Cisco; by Intel; and by NSF grants 2433923, 2402783, and 2211618.

Last revised: December 20, 2025.

- Augmented reality (AR)/virtual reality (VR) and immersive media applications require high throughput and tight latency bounds (<10 ms) to deliver realistic, lag-free experiences.
- Online gaming, particularly competitive and cloud-based platforms, demands sub-30 ms latency and high reliability, making jitter and packet loss as detrimental as slow speeds.
- Industrial Internet of things (IIoT) and smart manufacturing are increasingly looking to Wi-Fi as a viable alternative to wired Ethernet and 5G ultra-reliable low latency (URLLC). Real-time control loops, time-sensitive networking, and safety-critical functions require bounded latency, high availability, and deterministic behavior.
- Healthcare spaces adopt Wi-Fi for everything from real-time patient monitoring to robotic-assisted surgery preparation. These applications cannot tolerate performance lapses.
- Public infrastructure, including airports, smart cities, and connected public transportation systems, relies on Wi-Fi to deliver not only connectivity to users but also telemetry, surveillance, and machine coordination.

As these applications mature, the bar continues to rise. Beyond throughput and latency, predictability, reliability, and coordination are becoming new pillars of Wi-Fi performance.

B. From Speed to Efficiency: The Generational March

Wi-Fi is based on the IEEE 802.11 standard, which sets the foundations of the Physical layer (PHY) and Medium Access Control (MAC) layer of this technology to be implemented by access point (AP) and station (STA) devices (in the 802.11 specification, referred to as AP and non-AP devices, respectively). The 802.11 standard has significantly evolved in the last two and a half decades by following an iterative and backward-compatible approach, where each generation of Wi-Fi, marked by a new 802.11 amendment, has been introduced to address emerging technological and application needs. Early amendments (1990s-2000s) laid the groundwork for distributed channel access but employed relatively simple radio technologies for data transmissions. Over the years, newer 802.11 amendments gradually enriched the standard with features intended not only for higher capacity, but also for higher efficiency, more flexibility, and enhanced security. With this, Wi-Fi ensured support for new applications and devices with increasing communication requirements and, at the same time, fostered their proliferation.

Among some of the major breakthroughs in the history of Wi-Fi, we highlight: 802.11n (2009), which introduced packet aggregation and multiple-input, multiple-output (MIMO); 802.11ac (2013) with its multi-user MIMO (MU-MIMO) capabilities and wider channels to reach higher peak data rates; 802.11ax (2021), featuring orthogonal frequency-division multiple access (OFDMA) for scheduled uplink access, target wake time (TWT) for energy savings, and spatial reuse (SR) for spectral efficiency; 802.11be (2024), which introduced multi-link operation (MLO). Today, with 802.11bn (expected in 2028) on the horizon, Wi-Fi aims to add increased reliability to its wide portfolio of features.

This generational evolution is more than a series of speed upgrades. It represents a deepening of capabilities to support an increasingly complex and demanding ecosystem of applications and devices. More details on Wi-Fi generations and the standardization process are included in Section II. To aid the reader, Table I presents a summary of the key technological developments in each amendment and the use cases they target.

C. The Road to Wi-Fi 8: Low Latency and High Reliability

To support the next decade of wireless innovation, set by the previously highlighted use cases, Wi-Fi must expand into new domains including time-sensitive applications, highly-reliable control systems, and synchronized multi-device environments. This challenge is being taken up by the 802.11bn amendment [8], which is developing the next-generation standard expected to underpin Wi-Fi 8.

The primary objective of 802.11be was to increase capacity and link throughput and also improve worst-case latency and jitter with at least one mode of operation. While the latter was a novel endeavor compared to previous Wi-Fi generations, target latency and jitter were not quantified, making this only an initial step towards reliability. 802.11bn builds upon 802.11be foundational features, while making low-latency and high reliability first-class citizens. This shift is more than incremental, it represents a paradigm change and an architectural evolution in how Wi-Fi is designed, scheduled, and coordinated [9] [10]. To achieve such targets, key innovations agreed in the 802.11bn draft (D1.1) include [11]:

- Low-density parity check code (LDPC) enhancements: Upgrading the error correction capabilities to support higher reliability under poor channel conditions or for critical data streams (see Section III-A).
- Unequal modulation (UEQM): Allowing different modulation orders per spatial stream, optimizing spectral efficiency in multi-antenna systems with heterogeneous link qualities (see Section III-B).
- Enhanced long range (ELR) PHY protocol data units (PPDUs): Extending coverage and improving link robustness for long-distance transmissions, especially in sparse or high-interference environments (see Section III-F).
- Non-primary channel access (NPCA): Allowing stations and APs to initiate transmissions over secondary or tertiary channels, increasing flexibility in channel utilization and reducing contention (see Section IV-E).
- Dynamic subband operation (DSO): Providing very high flexibility to narrow-band STAs (e.g., using 20 MHz channels) by allowing them to be allocated to any subchannels among the AP's available bandwidth (see Section IV-E).
- Distributed resource units (DRUs): Enhancing power distribution flexibility in power spectral density (PSD)-limited regulatory regimes, enabling improved energy concentration and transmission robustness (see Section V-A).
- Dynamic power save (DPS): Enabling finer-grained power management strategies that adapt to application needs, helping extend the battery life of constrained Inter-

TABLE I: Evolution of IEEE standard amendments and associated Wi-Fi generations, use cases, key features, and peak data rates.

IEEE Standard	Generation	Year	Representative Use Cases	Key Features	Peak Rate
802.11b	—	1999	Basic broadband: email, browsing, file sharing (home/small office)	2.4 GHz, DSSS	11 Mb/s
802.11a	—	1999	Faster enterprise connections, early media streaming	5 GHz, OFDM	54 Mb/s
802.11g	—	2003	Home media, consumer electronics integration	2.4 GHz, OFDM, backward compatible with 802.11b	54 Mb/s
802.11n High Throughput (HT)	Wi-Fi 4	2009	HD video, cloud access, multi-device homes	Up to 40 MHz channels, packet aggregation, MIMO (up to 4×4)	600 Mb/s
802.11ac Very High Throughput (VHT)	Wi-Fi 5	2013	4K video, online gaming, mobile offload, enterprise WLANs	Up to 160 MHz channels, DL MU-MIMO (up to 4×4), 256-QAM	6.9 Gb/s
802.11ax High-Efficiency (HE)	Wi-Fi 6/6E	2021	Dense deployments (stadiums, campuses), IoT, teleconferencing	DL/UL OFDMA and MU-MIMO (up to 8×8), 1024-QAM, spatial reuse, TWT, 6 GHz band	9.6 Gb/s
802.11be Extremely-High Throughput (EHT)	Wi-Fi 7	2024	8K video, real-time collaboration, cloud gaming, AR/VR, IIoT pilot	Up to 320 MHz channels, 4096-QAM, MLO, multi-RU, enhanced QoS	23 Gb/s
802.11bn Ultra-High Reliability (UHR)	Wi-Fi 8	2028+	Robotic surgery prep, industrial automation, holography, ultra-reliable closed control loops	ELR PPDU, distributed-RU, LDPC enhancements, unequal modulation, seamless roaming, dynamic power save, NPCA, DSO, MAPC	23 Gb/s

net of things (IoT) and mobile devices (see Section VI-C).

- Seamless roaming: Introducing mechanisms for fast and non-disruptive handovers across APs to support mobility-sensitive applications such as AR/VR and real-time industrial control (see Section VII-C).
- Multi-access point coordination (MAPC): Moving beyond traditional mesh or roaming architectures, to support coordinated scheduling across APs. This will enable more deterministic performance and spatial reuse in dense deployments and may form the cornerstone for network-wide latency and reliability guarantees (see Section VIII).

D. Other Related Surveys and Tutorials

Wi-Fi has evolved significantly over the past two decades, with each generation introducing key innovations to address growing performance demands and increasingly diverse use cases. In parallel, a number of surveys and tutorials have emerged, each offering different perspectives on the evolution of the IEEE 802.11 standard. In Table II, we report a selection of the most relevant and representative review articles from the literature. Despite these valuable articles, the literature currently lacks a comprehensive, tutorial-style account that traces Wi-Fi's full evolution from 802.11b to the emerging 802.11bn, Wi-Fi 8. Existing works tend to focus on specific generations or technical layers, often omitting broader system-level perspectives and tutorial-level explanations. Crucially, no survey to date covers the final set of 802.11be features and 802.11bn overall. This article fills the gap by providing an integrated, accessible guide to the historical and technical development of Wi-Fi, offering both a foundation for new researchers and a forward-looking reference for ongoing work.

TABLE II: Other relevant articles on Wi-Fi.

802.11 for WLAN (core)	**Selected references**
Surveys and tutorials	[12], [13], [14], [15], [16], [17], [18], [19], [20]
802.11bn (Wi-Fi 8)	[8], [9], [10], [21]
802.11be (Wi-Fi 7)	[22], [17], [18], [23], [24], [25], [26], [19], [27], [21]
802.11ax (Wi-Fi 6)	[28], [14], [15], [29], [30], [16]
Older 802.11 for WLAN standards	[12]
802.11 for other use cases	**Selected references**
Wi-Fi coexistence (Coex SC)	[13], [31], [32], [33]
802.11ah (Wi-Fi HaLow)	[34], [35]
802.11s (Wi-Fi multihop)	[36]
802.11ad/ay (WiGig/mmWave)	[37], [38], [39]
Future of 802.11	**Selected references**
Wi-Fi sensing (TGbf) and positioning (TGbk)	[40], [41], [42], [43], [44], [45], [46]
Integrated mmWave (TGbq)	[10]
Artificial intelligence and machine learning (AIML SC)	[20], [47]

E. Contribution and Outline

This article provides a comprehensive tutorial on the key innovations and challenges across Wi-Fi generations from the first release of the 802.11 to the anticipated 802.11bn (Wi-Fi 8). We aim to serve both researchers and practitioners by highlighting how Wi-Fi has adapted to meet growing demands for throughput, latency, reliability, and multi-device coordination. We deliberately use original illustrations to give the main concepts in Wi-Fi technology a clear visual identity. Through carefully designed figures and diagrams, our tutorial consolidates knowledge across standards, use cases, and research trends, and aims to position readers to understand and contribute to future developments in wireless local area network (WLAN).

The primary contributions of this article are as follows:

- We detail the 802.11 evolution, focusing on PHY and

MAC layer innovations and their corresponding use cases, from the early amendments to the upcoming 802.11bn.
- We provide a tutorial overview of the core technologies introduced in 802.11ax/be (Wi-Fi 6/6E/7), including OFDMA, multi user (MU)-MIMO, and MLO, supported by examples and simplified explanations for accessibility.
- We examine emerging features in 802.11bn such as MAPC and many others, and discuss their expected impact on future WLAN applications.

The remainder of this paper is organized as follows:

- Section II provides an overview of the Wi-Fi spectrum allocation and standardization process, and summarizes the key features introduced in each amendment.
- Section III delves into PHY foundations and innovations, including symbol structures, modulation formats, and frame types.
- Section IV discusses MAC-layer techniques including distributed coordination function (DCF) operation, channel bonding, resource unit (RU)-based access, and power management, thereby setting the stage for advanced multi-user and coordination mechanisms.
- Section V builds on these PHY/MAC foundations and analyzes multi-user techniques such as OFDMA and MU-MIMO, detailing their design rationale and performance impacts.
- Section VI then describes Wi-Fi's energy-saving mechanisms adopted over the generations, and discusses how they interact with the previously introduced multi-user features.
- Section VII builds on the single-link operation of earlier sections and explores MLO, with a focus on link aggregation strategies and their role in reducing latency and boosting throughput.
- Section VIII finally surveys MAPC mechanisms, which extend the single-AP and multi-link concepts to coordinated operation among multiple APs, covering schemes from distributed scheduling to centralized coordination and discussing trade-offs and open challenges.
- Section IX concludes the paper.

Appendix: Building on the advancements envisioned for 802.11bn, the Appendix presents parallel efforts shaping the future of Wi-Fi, including mmWave operations (Appendix A), sensing (Appendix B), security and privacy (Appendix C), and the integration of artificial intelligence (AI) and machine learning (ML) techniques (Appendix D).

A summary of the main acronyms used in this tutorial is provided in Table III.

TABLE III: List of the main acronyms used throughout this article.

Acronym	Definition
AC	access category
ACK	Acknowledgment
A-MSDU	Aggregate MAC Service Data Unit
A-MPDU	Aggregate MAC Protocol Data Unit
AP	access point
BA	Block Acknowledgment
BSS	basic service set
CA	collision avoidance
Co-BF	coordinated beamforming
CCA	clear channel assessment
CFR	channel frequency response
Co-SR	coordinated SR
CSMA	carrier sense multiple access
CW	contention window
DCF	distributed coordination function
DIFS	distributed interframe space
DL	downlink
DRU	distributed resource unit
DSO	dynamic subband operation
DSSS	direct sequence spread spectrum
EDCA	enhanced distributed channel access
EHT	Extremely-High Throughput
EIFS	extended interframe space
ELR	enhanced long range
ICF	Initial Control frame
IEEE	Institute of Electrical and Electronics Engineers
ISI	inter-symbol interference
LTF	Long Training field
MAC	Medium Access Control
MAPC	multi-access point coordination
MCS	modulation and coding scheme
MIMO	multiple-input, multiple-output
MLD	multi-link device
MLO	multi-link operation
MPDU	MAC protocol data unit
MU	multi user
NAV	network allocation vector
NPCA	non-primary channel access
OBSS	overlapping basic service set
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
PCH	primary channel
PHY	Physical layer
PPDU	PHY protocol data unit
RF	radio frequency
RTS	request to send
RU	resource unit
SCH	secondary channel
SIG	Signal field
SIFS	short interframe space
SNR	signal-to-noise ratio
SP	service period
SR	spatial reuse
STA	station
STF	Short Training field
TB	trigger-based
TWT	target wake time
TXOP	transmission opportunity
UEQM	unequal modulation
UHR	Ultra-High Reliability
UL	uplink
WLAN	wireless local area network

II. Wi-Fi Spectrum Allocation, Standardization and Certification

As a cornerstone of today's digital ecosystem, Wi-Fi operates in unlicensed spectrum, enabling ease of deployment and free and widespread access across the globe. Among other things, Wi-Fi's widespread success can be attributed to the collaborative efforts of key industry stakeholders, industry forums such as Wi-Fi Alliance (WFA), and standards organizations, most notably the IEEE, which have continuously developed and refined the 802.11 family of standards and associated WFA's certification programs. These standards specify the PHY and MAC protocols that underpin Wi-Fi, ensuring interoperability and robust performance across a diverse range of devices and environments. This section provides an overview of the Wi-Fi standardization cycle, from the latest

advances in unlicensed spectrum allocation and regulation to the development of IEEE 802.11 standards and the subsequent Wi-Fi certification process led by WFA.

A. Spectrum Allocation and Coexistence

Spectrum allocation has been a key driver of Wi-Fi's success, enabling it to operate in unlicensed bands, primarily the 2.4 GHz and 5 GHz bands, which are allocated worldwide for industrial, scientific, and medical (ISM) applications. There are only three non-overlapping 20 MHz channels in 2.4 GHz and a single non-overlapping 40 MHz channel spanning the available bandwidth of approximately 80 MHz, severely limiting high-throughput applications. The 5 GHz bands offer wider bandwidths and more channels designated in the U.S. into four Unlicensed National Information Infrastructure (U-NII) bands as follows: U-NII-1 (5.15 - 5.25 GHz), U-NII-2A (5.25 - 5.35 GHz), U-NII-2C (5.47 - 5.725 GHz), U-NII-3 (5.725 - 5.85 GHz) and U-NII-4 (5.85 - 5.895 GHz). The U-NII-2A and U-NII-2C bands have to implement dynamic frequency selection (DFS) to protect airport radars that operate in those bands, and hence are sparsely used.

As Wi-Fi has evolved to meet growing demands for high-speed applications such as cloud storage, video streaming, and real-time communications, these bands have become increasingly congested. Even though 802.11ax brought improvements in spectrum efficiency, high-density scenarios still face significant challenges due to spectrum congestion in 5 GHz [48].

The introduction of Wi-Fi 6E (based on 802.11ax) in the unlicensed 6 GHz band marked a major advancement, offering up to 1.2 GHz of additional spectrum in some countries. This new band provides significant relief from congestion, allowing for up to seven 160 MHz channels, as compared to only two in the 5 GHz band. Wi-Fi 6E also introduced tri-band APs that simultaneously support the 2.4 GHz, 5 GHz, and 6 GHz bands, enabling devices to operate in the frequency band with the best available quality at a given time. The latest 802.11be amendment, with MLO and 320 MHz channels, further enhances this by allowing devices to connect over two or more bands simultaneously, improving overall capacity, reducing latency, and boosting reliability. The 6 GHz band is partitioned into four U-NII bands as follows: U-NII-5 from 5.925–6.425 GHz, U-NII-6 from 6.425–6.525 GHz, U-NII-7 from 6.525–6.875 GHz and U-NII-8 from 6.875–7.125 GHz. Most regulatory regions around the globe have opened either all these bands for unlicensed operation or only U-NII-5 [49].

1) The 6 GHz Unlicensed Band Allocation: The 6 GHz band is a shared but unlicensed band. The primary incumbents in U-NII-5 and U-NII-7 are fixed links while Broadcast Auxiliary Services (BAS) are the main incumbent in U-NII-6 and U-NII-8. In addition, satellite uplink and downlink allocations exist across the band. Allocating the band on a low-power unlicensed basis instead of high-power, exclusively licensed basis permits these incumbents to continue operating without interference. For Wi-Fi, the 6 GHz band is particularly valuable because it is free from contention with and interference by legacy Wi-Fi devices operating in the 2.4 GHz and 5 GHz bands. This spectrum expansion is critical for accommodating the exponential increase in demand for Wi-Fi capacity. Following the US Federal Communication Commission (FCC) decision in 2020 to open the 6 GHz band for unlicensed operation, allocation of the 6 GHz band has become a focal point of regulatory efforts worldwide, particularly during the World Radiocommunication Conference 2023 (WRC-23). One of the most debated issues is the future of the upper 6 GHz band (6425–7125 MHz). Regulatory bodies across the three International Telecommunication Union (ITU) regions are evaluating how to balance its use between Wi-Fi and future 6G applications:

- *Region 1:* Covering Europe, the Middle East, Africa, and the CIS, with regulators such as CEPT, ASMG, ATU, and RCC.
- *Region 2:* The Americas, regulated by CITEL.
- *Region 3:* The Asia-Pacific region, regulated by APT.

Early studies suggest the upper 6 GHz band shows promise for extending indoor cellular coverage, but coexistence with Wi-Fi remains a challenge, particularly regarding interference. Coexistence of indoor and outdoor Wi-Fi in a dense stadium deployment has been studied in [50], demonstrating that coexistence even within Wi-Fi systems may be a challenge. While reducing the power output of cellular systems might mitigate interference with Wi-Fi, this would compromise the indoor penetration and overall effectiveness of cellular networks. In addition to this WLAN-cellular spectrum sharing challenge, the 6 GHz band already supports several incumbent services such as satellite communications, fixed links, and passive Earth observation systems. Effective coordination between these services and new Wi-Fi deployments are critical to successful spectrum sharing in this band [51], [52].

2) 6 GHz Band Operating Rules: Strict regulatory frameworks govern the use of the 6 GHz spectrum to ensure coexistence with incumbent services operating in this band. These rules vary between different countries. Depending on location (indoor or outdoor), devices operating in the 6 GHz band must adhere to specific power emission limits to avoid interference. To protect incumbent services and manage potential interference, the 6 GHz band has been divided into categories of Wi-Fi devices, each with distinct operational rules and power limits:

- *Low power indoor (LPI) devices* are restricted to indoor use and are subject to lower power limits to minimize interference risks. These devices do not require automatic frequency coordination (AFC) but must meet stringent PSD limits, use fixed antennas, and cannot operate on battery power. In the USA, LPI devices are limited to a PSD of 5 dBm/MHz and maximum equivalent isotropic radiated power (EIRP) of 30 dBm, which is achieved over 320 MHz channels. However, in the EU and UK the EIRP limits are 23 dBm and 24 dBm respectively, irrespective of channel bandwidth.
- *Very low power (VLP) devices*, designed for portable and mobile applications, can be used indoors and outdoors with minimal power output, making them suitable for applications like wearables and IoT sensors. They can operate in all U-NII bands, with a PSD limit of -5 dBm/MHz.

- *Client devices* can operate indoors and outdoors across all U-NII bands but must comply with power limits that are 6 dB lower than their associated APs.
- *Standard power devices* are permitted to operate both indoors and outdoors in the U-NII 5 and U-NII 7 bands. However, they are required to use AFC—discussed in the sequel—to dynamically adjust power levels and frequency usage, preventing interference with incumbent services. Additionally, their transmissions are restricted to angles below 30 degrees to avoid interfering with services deployed at higher altitudes, such as satellite links. Standard power has been fully authorized in North America and other countries are considering similar rules [53].

3) Automated Frequency Coordination (AFC): AFC plays a crucial role in the regulatory framework for the 6 GHz band, ensuring that standard power Wi-Fi devices can operate without disrupting incumbent services. This system works by dynamically allocating frequencies and power levels based on the location of the Wi-Fi device and nearby incumbent systems. Before transmitting, a standard power device queries the AFC system to determine which frequencies are safe to use at its location and what power levels are permissible. The AFC system takes into account the geographical coordinates, elevation, and antenna characteristics of the device, as well as the presence of other users in the band. In the USA, standard power devices operating under an AFC do not need to implement a contention-based mechanism unlike LPI devices: this allows the use of the band by standard power systems that do not use the Wi-Fi protocol. This centralized approach ensures that Wi-Fi devices stay within regulatory boundaries, while still making full use of the available 6 GHz spectrum. By managing interference in real-time, AFC allows Wi-Fi devices to operate alongside existing services without causing disruption.

B. Wi-Fi Amendments: From Drafting to Certification

The IEEE 802.11 standard forms the backbone of Wi-Fi, but Wi-Fi technology goes far beyond and involves a complex set of standardization procedures in which different players have their own role. Fig. 1 summarizes the procedures and stakeholders involved in the elaboration of Wi-Fi specifications, from the creation of 802.11 standards to their certification by the WFA. However, for Wi-Fi to be part of the broader picture that is the Internet, its interworking with the protocol suite (TCP/IP) is enabled thanks to the efforts of organizations like Internet Engineering Task Force (IETF) and protocols like dynamic host configuration protocol (DHCP), Domain Name System (DNS), or Internet Protocol (IP), which are deeply integrated into Wi-Fi devices like APs.

1) IEEE 802.11 Standardization: The IEEE 802.11 standard defines the MAC and PHY layers of Wi-Fi, but does not go above in the TCP/IP stack. At the core of 802.11 is the concept of basic service set (BSS), which represents a basic network unit managed by an AP and to which STAs (or non-AP STAs, according to 802.11 terminology) can connect and access the Internet thanks to built-in routing capabilities on the AP devices, e.g., DHCP and network address translation

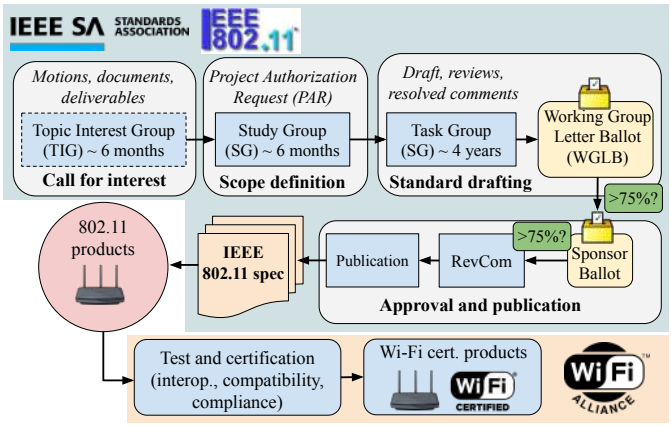

Fig. 1: Overview of the Wi-Fi standardization and certification process. The figure shows the main steps within IEEE 802.11, from the initial call for interest and Topic Interest Group, through Study and Task Groups, balloting, and final approval and publication of the IEEE 802.11 specification, and then how the published standard feeds into Wi-Fi Alliance interoperability testing and certification of commercial Wi-Fi products.

(NAT). Based on the BSS concept, multiple solutions and architectures are derived. While mesh and device-to-device modes also exist within Wi-Fi, in this paper, we will mainly focus on AP-STA communication within BSSs.

The 802.11 working group (WG) oversees all the standardization activities within the group, which occur at different levels through the task groups (TGs), study groups (SGs), Standing Committees (SCs), ad-hoc committees (AHCs), and topic interest groups (TIGs). While the development cycle is specific to each 802.11 amendment, it can be roughly divided into three phases:

- *Scope definition:* After identifying new market demands and the need for a new amendment to the 802.11 standard—typically as a result of multiple discussions within the 802.11 WG or from the activities of a TIG—a project authorization request (PAR) is prepared to narrow down the goals that the upcoming amendment will need to fulfil. The PAR, typically prepared by an appointed SG within six months, serves as the green light to initiate the development of the amendment.
- *Standard amendment drafting:* On approving a PAR, a TG is appointed to develop the drafting of the specification, for which four years are granted. Drafting an amendment to the standard involves delving into the technical requirements and solutions to achieve the goals specified in the PAR. For that, TG participants discuss features and contributions until a stable draft of the amendment is achieved. At that stage, the draft is put to vote among all members of the 802.11 WG. This procedure is known as working group letter ballot (WGLB), and it requires a minimum of 75% positive votes to move forward. During this process, comment resolution typically entails several back and forth interactions between the TG and the WG.
- *Standard approval and publication:* When the WGLB is approved, a second voting level (the Sponsor Ballot) needs to be passed with a minimum support of the 75%. In this case, the vote is extended to the IEEE Standards Association (IEEE-SA) and the resulting text is forwarded

to the IEEE Standards Review Committee (RevCom) for final approval. Once obtained, the final specification is published, so that 802.11 vendors and manufacturers can develop their products accordingly.

2) Wi-Fi® Certification: The Wi-Fi certification process helps guarantee that certified devices (e.g., network equipment, computers, smartphones, appliances) implement the latest 802.11 features, security protocols, and more, and ensure a good user experience. The WFA plays a vital role in ensuring device compatibility and interoperability by certifying products that adhere to 802.11 standard amendments. WFA's history began in 1999, when a group of companies founded the WFA—at that time called the WECA. In 2000, the term "Wi-Fi" was coined, and since then the WFA controls the use of the "Wi-Fi Certified" logo, a registered trademark. This logo is only stamped on devices that pass the certification tests defined by the WFA.

The WFA defines multiple certification programs and tracks (*FlexTrack*, *QuickTrack*, and *Derivative*), which focus not only on physical transmissions and data formats, but also on other functionalities such as those related to security or quality of service. The WFA provides certification programs directly aligned with 802.11 amendments (e.g., Wi-Fi CERTIFIED 7™ [54]) and specific to other functionalities such as WPA3 [55] or Wi-Fi direct [56]. Apart from the certification tests performed by authorized test laboratories, the WFA organizes interoperability events (or plugfests), where Wi-Fi vendors, developers, and manufacturers meet to conduct tests that demonstrate robust performance of multiple Wi-Fi implementations under real-world conditions, interoperability with other certified devices, and backward compatibility.

C. Overview of IEEE 802.11 Amendments

The IEEE 802.11 standard amendments form the backbone of Wi-Fi technology, with each iteration introducing new features and improvements to address the evolving needs of wireless communications. The intricate interaction between an AP, its STAs, and neighboring BSSs is key to understanding Wi-Fi's operational framework, the challenges it faces, and the features developed to address those challenges.

1) From Legacy 802.11 to 802.11ax: The original IEEE 802.11 standard (1997), following the FCC's 1985 opening of the ISM band for WLAN usage, laid the groundwork for data transmission over unlicensed spectrum bands. Back in the 1990s and early 2000s, Wi-Fi could offer up to 2 Mb/s to devices operating in the 2.4 GHz band, but its meteoric pace of development has led to reaching up to 9.6 Gb/s just two decades in the 802.11ax amendment (Wi-Fi 6/6E). This dramatic increase has been driven by the adoption of new bands (5 GHz and 6 GHz) and technologies such as orthogonal frequency-division multiplexing (OFDM) and MIMO. The evolution of Wi-Fi from its inception to Wi-Fi 6/6E can be summarized as follows:

- Wi-Fi's journey began in 1997 with the original 802.11 standard, which was based on carrier sense multiple access (CSMA)/collision avoidance (CA) and single-carrier modulations such as direct sequence spread spectrum (DSSS) or frequency-hopping spread spectrum (FHSS). In addition, the foundations in terms of association and authentication were laid at this point.
- In 1999, 802.11b amended the initial specification from 1997 to enhance data rates to up to 11 Mb/s, in part due to the adoption of complementary code keying (CCK) modulation. It was at this time that Wi-Fi products started to be commercially attractive, and thus its massive adoption began to provide residential wireless Internet access, wireless networking in offices, and public hotspots. The growth in popularity of Wi-Fi concides with the creation of the WFA, the growth of broadband Internet usage, and the need for wireless mobility by the increased accessibility to laptops by the bulk of the population.
- That same year, 802.11a was introduced to operate in the 5 GHz band. This, together with the adoption of OFDM, allowed reaching peak data rates of 54 Mb/s. In 2001, the FCC permitted the use of OFDM in the 2.4 GHz band, leading to the 802.11g amendment in 2003. The 802.11g amendment became very important thanks to its backward compatibility with 802.11b.
- The release of 802.11n (Wi-Fi 4 [57]) in 2009 marked another significant leap in the evolution of Wi-Fi [58]. The adoption of features such as MIMO spatial multiplexing and channel bonding enabled peak data rates of up to 600 Mb/s, and spatial diversity led to improved reliability. 802.11n expanded Wi-Fi beyond casual web browsing and email to support high definition (HD) video streaming and richer web applications.
- In 2013, 802.11ac (Wi-Fi 5 [59]) built upon the advancements of 802.11n and introduced even larger channel bandwidths (up to 160 MHz), higher modulations (256-quadrature amplitude modulation (QAM)), and MU-MIMO in the downlink, enabling multiple devices to receive data simultaneously [60]. The 802.11ac reached multigigabit speeds, cementing Wi-Fi as the de facto standard for home entertainment, cloud services, and mobile offload.
- In 2021, 802.11ax (Wi-Fi 6 [61]) [62] marked a shift from raw speed to efficiency. Recognizing the challenge of high-density environments such as stadiums, airports, and apartment complexes, it introduced features such as OFDMA, uplink MU-MIMO, and spatial reuse (BSS coloring), significantly improving spectral efficiency, latency, and fairness. Wi-Fi 6E [63], an extension of Wi-Fi 6, unlocked 1.2 GHz of bandwidth in the 6 GHz band, providing additional spectrum to reduce congestion and enable cleaner, high-throughput links.

Fig. 2 provides an overview of the evolution of Wi-Fi amendments in terms of their achievable peak data rates (see Section III-E for their computation).

2) 802.11be Extremely High Throughput: Commercialized in 2024, the 802.11be amendment (Wi-Fi 7) introduces even more advanced features to do justice to its full name, *EHT*. The performance goals of the 802.11be are achieved through new MAC and PHY modes of operation, as well as the use of the 6 GHz band.

The 802.11be's PHY defines support for wider bandwidths

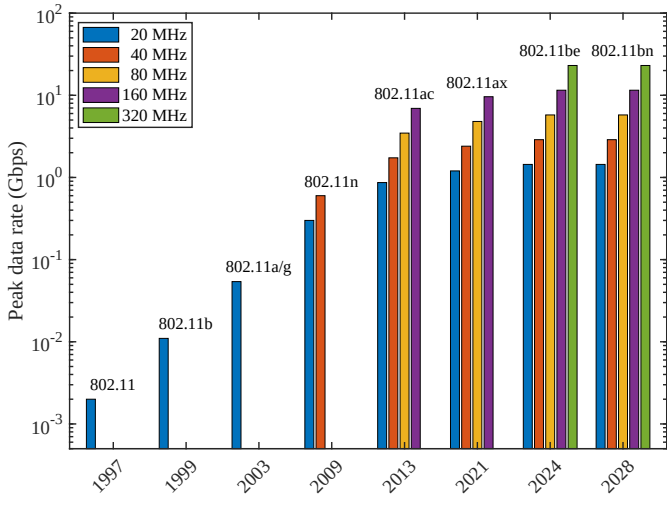

Fig. 2: Peak data rate and bandwidth utilization across IEEE 802.11 standard amendments. The bars show the evolution of the theoretical peak PHY data rate (Gb/s, logarithmic scale) for different channel bandwidths (20–320 MHz, color-coded) as new 802.11 amendments are introduced over time (from 802.11 to 802.11bn).

(up to 320 MHz channels) and higher modulation schemes (up to 4096-QAM). To exemplify the potential of 802.11be's PHY, a link using a 320 MHz channel, 8×8 MIMO, a short guard interval of 0.8 μs, and 4096-QAM with 5/6 coding can achieve a data rate of 23 Gb/s. A more conservative setup, e.g., a 160 MHz channel with 2×2 MIMO and 1024-QAM with 3/4 coding, would still lead to a 2.1 Gb/s data rate.

As for the 802.11be's MAC, the main breakthrough has been achieved by MLO, which we will discuss in depth in Section VII. In a nutshell, MLO further increases throughput by using multiple physical links in parallel through a single association. Apart from that, the 802.11be includes appealing features such as support for multiple RUs in OFDMA.

Although 802.11be makes important strides toward lower latency and higher reliability, these aspects were not fully quantified and rather deferred to 802.11bn, which will prioritize reliability in addition to throughput.

3) 802.11bn Ultra High Reliability: As Wi-Fi continues to evolve, emerging use cases and applications require not only increased throughput and reduced latency, as addressed by Wi-Fi 7, but also a higher level of reliability. These evolving demands are shaping the development of the 802.11bn, with a strong focus on UHR. Some of the most critical emerging use cases for 2030 are driven by the advances in immersive communications, digital twins, and real-time control systems, among others (see Table I). Each of these applications requires low-latency, highly reliable communication. To meet these requirements, Wi-Fi is moving toward more deterministic performance.

At the time of writing, the UHR Task Group (TGbn) has released a draft of the 802.11bn amendment (D1.1) [11]. As included in the PAR, 802.11bn will target a 25% improvement in data rates, even at lower signal-to-interference-plus-noise ratio (SINR) levels, along with a 25% reduction in tail latency and packet loss rates, particularly in environments with mobility and overlapping basic service set (OBSS). Additional advancements in power-saving features and peer-to-peer operation are expected, with 802.11bn extending its operation across all sub-7 GHz bands.

The TGbn, which was formally established in November 2023, will define the technical objectives and protocol functionalities of Wi-Fi 8 by 2028, when the final amendment is expected to be published. In the meantime, the first draft (D1.0) is planned to be released in July 2025, whereas the final draft is expected to be submitted for Sponsor Ballot in May 2027.[1]

To guarantee low-latency, high-reliable communication in unlicensed bands, one of the key areas of development for the 802.11bn will be MAPC, extensively discussed in Section VIII, which will allow multiple APs to cooperate in managing interference, coordinating channel access, and improving overall spectrum efficiency. This approach is expected to significantly improve performance, especially in dense network environments where many devices compete for limited spectrum. MAPC will entail a true paradigm shift in Wi-Fi, which has historically relied on the inherently non-deterministic CSMA/CA, to comply with the regulation on the unlicensed spectrum that mandates the usage of listen-before-talk (LBT) [64].

4) Other Ongoing Standardization Activities: While 802.11bn will focus primarily on improving Wi-Fi performance in the sub-7 GHz bands, the future of Wi-Fi also includes expansion into higher frequency ranges. The 60 GHz band, in particular, presents new opportunities for ultra-high-throughput and low-latency applications. The 802.11bq Task Group (TGbq) is tasked with exploring operational expansion into the 60 GHz band, building upon the work done in previous amendments like 802.11be and 802.11bn. The 60 GHz band offers significant advantages for short-range, high-bandwidth applications, but also introduces challenges such as propagation loss and sensitivity to obstacles, as discussed in Appendix A. Solving these challenges will be key to enabling the next generation of Wi-Fi applications, such as uncompressed video streaming and high-speed data transfer for industrial use.

As discussed in Appendix B, Wi-Fi sensing, standardized in 2025 under the 802.11bf amendment ("Wireless LAN Sensing"), leverages Wi-Fi signals to enable applications such as motion detection, human activity recognition, and proximity sensing. By tracking changes in channel, devices can infer environmental factors without requiring the target to carry a Wi-Fi device. Wi-Fi sensing has promising applications in residential, enterprise, and industrial settings, offering capabilities like security monitoring, gesture control, and smart device management. The 802.11bf amendment unifies sensing procedures across devices, paving the way for interoperable and scalable sensing solutions using existing Wi-Fi infrastructure. As a complementary capability, the IEEE 802.11-2016 standard revision incorporated the Fine Timing Measurement (FTM) protocol to enable precise indoor localization through Wi-Fi. FTM estimates the distance between devices by computing the round trip time (RTT) of sounding frame transmissions. Building on this, in 2022, the 802.11az

[1] https://www.ieee802.org/11/Reports/802.11_Timelines.htm#tgbn

amendment ("Enhancements for Positioning") standardized a procedure to enable an STA to identify its position relative to a set of APs by combining FTM measurements collected from each of them. The update also improved the accuracy from 1-2 meters of 802.11-2016 to sub-meter precision. In 2025, this further evolved into the IEEE 802.11bk amendment ("320 MHz Positioning") which extends the FTM positioning mechanisms to make use of 320 MHz wide channels available with 802.11be.

In terms of privacy and security, the TGbh and TGbi are currently working on randomized MAC addresses (so that tracking devices can be more difficult for potential attackers) and enhanced data privacy (EDP) (including mechanisms to protect transmissions and management frames), respectively. The 802.11bh amendment has already passed the Sponsor Ballot and is about to be published officially, while 802.11bi is currently at D1.0. Both 802.11bh and 802.11bi are further discussed in Appendix C.

Another critical development in Wi-Fi standardization is the integration of AI and ML features, aimed at enhancing network management and operational efficiency and discussed in Appendix D. As Wi-Fi networks grow increasingly complex, AI/ML will play a pivotal role in real-time decision-making and optimization. The efforts to incorporate AI/ML were initiated by the IEEE 802.11 AIML TIG and are now being followed up by the IEEE 802.11 AIML SC, which is exploring how these technologies can be leveraged in novel use cases that allow automating network management tasks, optimizing resource allocation, or predicting traffic patterns to improve overall network performance, among others.

III. Wi-Fi Physical Layer

Wi-Fi's PHY underpins the data transmission process over wireless channels. At the transmitter device, the PHY receives data from the MAC in a PHY service data unit (PSDU), decides the parameters for the PHY transmission (e.g., modulation and coding scheme (MCS)), and then builds a PPDU by adding to the PSDU a PHY preamble, which specifies the selected parameters and other control fields for frequency offset correction and time synchronization at the receiver. The sequence of data and control bits in the PPDU is then processed to obtain a radio frequency (RF) signal to be irradiated through the antennas. Note that while the parameters for the data field (PSDU) are decided at the PHY, the transmission parameters for the control fields in the preamble are fixed and specified in the standard. At the receiver, the PHY uses the control information in the preamble to detect the start of the PPDU and decode it into a PSDU, which is forwarded to the MAC. This section provides an overview of the physical bit transmission pipeline, deepening the explanation of the modulation and coding strategies together with the spatial and frequency multiplexing techniques. The PPDUs formats are also detailed at the end of this section.

A. Physical Layer Block Diagram

Fig. 3 provides an overview of the processing steps to generate the Wi-Fi RF signals to be radiated through the different antennas available at the transmitter side, starting from the PPDU bits. In particular, the figure summarizes the processing steps applied to the PPDU data field (see Section III-F). The other PPDU fields, e.g., Short Training field (STF) and Long Training field (LTF), are processed following a similar pipeline where some of the blocks are removed.

The procedure is as follows. At first, the data bits are scrambled using an additive scrambler, which applies an exclusive OR (XOR) of the bits with a scrambling sequence. This operation reduces the probability of long sequences of bits equal to zeros or ones, which helps time synchronization at the receiver. The receiver should use a matched descrambler to decode data. Redundancy is then added to the scrambled bits for forward error correction (FEC). Two different FEC coding techniques can be used for this purpose, namely binary convolutional code (BCC) and LDPC [65]. The amount of redundancy is a function of the coding rate, which is specified through the MCS parameter, as detailed in Section III-B. Note that padding should be added before the scrambler (pre-FEC PHY) and after the FEC encoder (post-FEC PHY).

Depending on the coding strategy adopted (BCC or LDPC), an additional processing block is included in the diagram in Fig. 3. Specifically, BCC requires an interleaver block, which consists of a two-step permutation of the coded bits to avoid long sequences of noisy bits on the BCC decoder and thus improve transmission robustness to burst errors. The first permutation maps adjacent coded bits into non-adjacent OFDM subcarriers (see Section III-C), while the second maps adjacent coded bits alternately onto less and more significant bits of the constellation. Instead, a tone mapper is used for LDPC-coded bits after being modulated. The complex modulated symbols are permuted such that each two consecutive symbols will be transmitted on two OFDM subcarriers that are separated by at least D_{TM} OFDM data subcarriers, where D_{TM} is an integer parameter specified in the standard for each bandwidth: 4, 6, 8, 9 for $20\,$MHz, $40\,$MHz, $80\,$MHz, and $160\,$MHz respectively.

While the first version of 802.11 considered the transmission of a single modulated symbol per time resource, multiplexing techniques are key components of current Wi-Fi networks to improve spectrum efficiency. Hence, the coded bits are associated with N_{SS} different spatial streams for MIMO transmission in a round-robin fashion, creating groups of N_{BPSC} bits, where $2^{N_{\mathrm{BPSC}}}$ is the modulation order. The bits on each of the N_{SS} spatial streams are modulated independently, using the modulation scheme defined by the MCS parameter (see Section III-B). Hence, the modulated symbols are forwarded to the OFDM-MIMO modulator for frequency and spatial multiplexing. The complex modulated symbols in each spatial stream are processed in groups of N_{SD} symbols via OFDM, where N_{SD} depends on the bandwidth and the OFDM subcarrier spacing, as it will be detailed in Section III-C. Additional known OFDM symbols called pilots are added within the set of N_{SD} modulated data symbols in specific positions for receiver-transmitter synchronization purposes to make the data detection robust against frequency offsets and phase noise. The indices of the pilot OFDM subcarriers are defined by the 802.11 standard for the different operational bandwidths. Cyclic shift diversity (CSD) is used to prevent

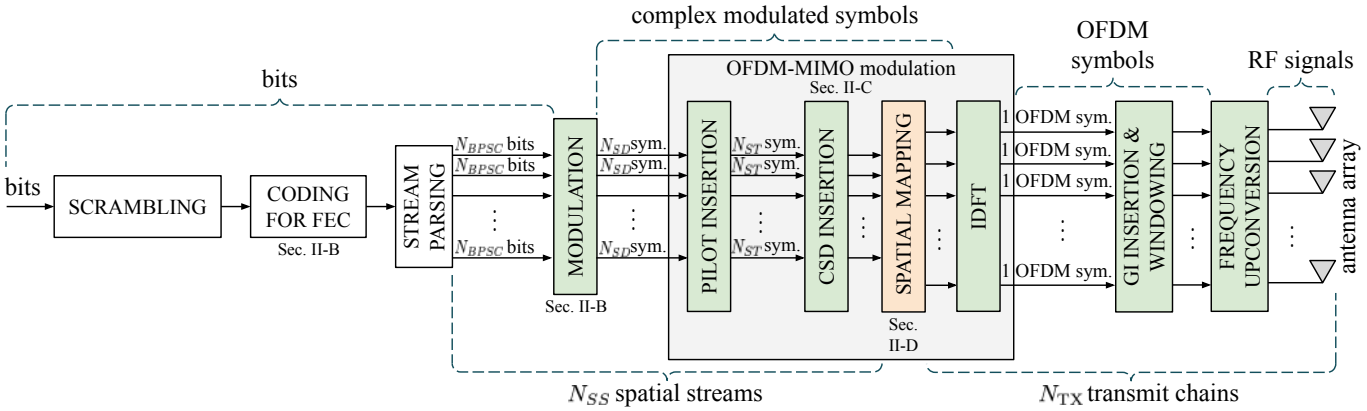

Fig. 3: Main signal processing blocks used by 802.11-compliant devices to convert bits into RF signals. The green blocks refer to processing steps applied to each spatial stream or transmit chain independently, e.g., N_{SS} modulation blocks are used, one for each of the N_{SS} spatial streams. The spatial mapping (yellow block) simultaneously processes all the data streams. Depending on the coding strategy adopted (BCC or LDPC), a few additional processing blocks are included. The notation 'sym.' stands for 'symbol(s)'.

unintentional beamforming—which occurs when correlated signal are transmitted over multiple antennas—by applying slightly different cyclic shift to each of the space-time streams. This technique was introduced in the 802.11n amendment and is also essential for backward compatibility with single-antenna receivers.

Spatial mapping is applied next, as it will be explained in Section III-D. Following this, an inverse fast Fourier transform (IFFT) operation allows converting the set of modulated symbols plus pilots into an OFDM symbol, as it will be described in Section III-C. A Guard interval (GI) is added at the beginning of each OFDM symbol to reduce inter-symbol interference (ISI) during transmission, and a time-windowing function is applied to set the boundary of the transmitted signal. The specific time-windowing function to be applied is not specified in the standard and is therefore a choice of the manufactures. Finally, the RF signals to be transmitted over the different antennas are obtained through a frequency up-conversion at the specific operating frequency defined by the selected Wi-Fi channel.

Latest 802.11bn Enhancements: Building upon the transmitter block logic described above, recent developments under the forthcoming 802.11bn amendment introduce significant enhancements aimed at improving performance and reliability. One of the key updates focuses on FEC through the inclusion of a new, longer block-length LDPC code of 3888 bits (2×1944). This new code complements the existing options of 648, 1296, and 1944 bits, and is designed to enhance error correction, particularly in high-reliability scenarios like extended reality (XR) applications and dense wireless environments. The longer code provides greater coding gain, improving performance across different MCS levels (see Section III-B). It follows a quasi-cyclic (QC) structure derived from the base 1944-bit code and supports code rates of 1/2, 2/3, 3/4, and 5/6. The 3888-bit code integrates smoothly into the updated PPDU encoding table while staying compatible with legacy codeword lengths for smaller payloads. Thanks to its modular structure, this design allows for parallel encoding and decoding, enabling higher throughput with minimal added complexity. Performance evaluations show gains of 0.5 to 1.0 dB, especially at higher MCS levels and when multiple spatial streams are used. That said, using 3888-bit codes can be challenging for short payloads, where puncturing (see Section IV-E) may hurt performance. Successful implementation depends on the ability to parallelize decoding. In the 5 GHz and 6 GHz bands, this is often feasible due to the presence of multiple decoders. However, in the 2.4 GHz band, where channels are narrower, this capability may be lacking. As a result, it may be best to restrict support for 3888-bit LDPC codes to the 5 GHz and 6 GHz bands to avoid added complexity.

B. Modulation and Coding Schemes

The MCS index is a positive integer number that defines the parameters for modulation and coding. With modulation, we refer to the techniques adopted in wireless systems to convert digital data (bits) into analog signals by changing the shape of a reference signal, called carrier, within a set of possible modifications. The resulting modulated signal carries the bit stream. Devices adopting the 802.11 standard support phase-shift keying (PSK) and QAM schemes, where the former entails changing the phase of the carrier to represent data while the latter does it by also changing the carrier's amplitude. The decoder detects the specific changes to retrieve the transmitted bits. Each modification can be identified as a point in the complex plane (I/Q) and, in turn, is referred to as constellation point. The modulation order identifies the number of constellation points in the modulation scheme ($2^{N_{\mathrm{BPSC}}}$), which in turn defines the number of bits that can be carried by each signal modulated differently (N_{BPSC}). For this, the stream of coded bits is divided into groups of N_{BPSC} bits, each of which is mapped into a constellation point generating a modulated symbol. For example, a 256-QAM modulated symbol is obtained by combining 16 bits.

The second information provided by the MCS is the coding rate, which defines the redundancy level introduced by BCC or LDPC coding to enable FEC at the receiver. For example, a 1/2 coding rate means that, for each information bit, there is another bit for FEC purposes. This rate should be used under less favorable channel conditions where more redundancy is

required, while coding rates of 3/4 or 5/6 can be adopted when the signal-to-noise ratio (SNR) is sufficiently high.

The MCS to be used for data transmission is defined by the transmitter by evaluating the channel conditions. Low MCS values, i.e., low modulation order and high coding redundancy, lead to more robust transmissions at the expense of reduced throughput and should be preferred in low SNR conditions. High MCSs require increasingly higher SNR to maintain robustness against noise and interference but increase data rates by modulating more bits together per symbol and reducing coding redundancy, benefiting bandwidth-intensive applications such as 4K video streaming, virtual reality, and large file transfers. For this, in addition to the MCSs [0, . . . , 9] available in 802.11ac, 802.11ax introduced 1024-QAM with two possible coding rates (3/4 for MCS 10 and 5/6 for MCS 11), improving link efficiency in high-density deployments [26]. The 802.11be amendment took this further by introducing 4096-QAM, tied to MCS 12 and 13 [66]. However, 4096-QAM has stricter error vector magnitude (EVM) requirements, about 3 dB higher than 1024-QAM, necessitating exceptionally clean signals for successful demodulation: the required SNR is around 40 dB, which may only be achievable through techniques like beamforming.

802.11be also introduced dual carrier modulation (DCM) to improve robustness to fading in signal reception by transmitting each modulated symbol into two different OFDM subcarriers [66]. DCM is used in conjunction with BPSK and coding rate 1/2 and defines two new MCSs, i.e., MCS 14 and MCS 15, both of which can only be used for single spatial stream transmissions. The difference between these two new MCSs is that the first requires the use of duplication (DUP) mode, where the data is also duplicated in the lower and upper parts of the channel (e.g., lower and upper 484 subcarriers for 80 MHz), thus doubling the duplication performed through DCM and, in turn, further increasing robustness. DUP mode is only applicable in the 6 GHz band with LDPC coding. Note, however, that these two additional MCSs have not been certified by the Wi-Fi Alliance as part of Wi-Fi 7, thus limiting their adoption into commercially available products.

Latest 802.11bn Enhancements: In earlier 802.11 devices employing multi-stream MIMO transmissions (see Section III-D), all spatial streams were configured with the same modulation order and a unified coding rate. However, when beamforming is used, the wireless channel conditions can vary significantly across different spatial streams, leading to unequal SNR. To address this, 802.11bn introduces UEQM, a feature that allows assigning different modulation orders to each stream within a single transmission. Despite this variability, a common LDPC code is applied to all streams to maintain a consistent error correction. By adapting modulation to the real-time conditions of each stream, UEQM can improve both spectral efficiency and link robustness, making it especially useful in high-interference or spatially heterogeneous environments. Complementing this, 802.11bn also adds several new MCS levels to allow for finer-grained link adaptation. Examples include QPSK with a 2/3 coding rate and 5/6 16-QAM, offering greater flexibility to optimize performance in diverse deployment scenarios. However, these enhancements also increase the complexity of rate adaptation. With more combinations of MCS and number of spatial streams (N_{SS}) available, the transmitter must navigate a larger decision space. To manage this effectively, improved feedback mechanisms may be needed, guiding the transmitter in selecting the most suitable MCS and transmission configuration under current channel conditions.

C. Orthogonal Frequency Division Multiplexing (OFDM)

OFDM was introduced in 802.11a and allows multiplexing modulated symbols in the frequency domain. To do this, N_{ST} partially overlapping orthogonal subcarriers are obtained from the available bandwidth B, and used to simultaneously transmit N_{SD} modulated symbols. The remaining $N_{\mathrm{ST}}-N_{\mathrm{SD}}$ subcarriers are occupied by N_{SP} pilot symbols and additional unused subcarriers [65]. The unused ones include the DC subcarrier (center subcarrier) and the guard band subcarriers at the band edges. Moreover other null (unused) subcarriers located near the DC provide protection from transmit center frequency leakage, digital-to-analog converter (DAC) and analog-to-digital converter (ADC) offsets at the receiver, and interference from neighboring RUs in OFDMA transmissions (see Section V-A) [67]. Their number depends on the operating bandwidth and the subcarrier spacing. For example, in 802.11ax, there are 3 null subcarriers around the DC for 20 MHz transmissions, and 5 for 40 MHz and 80 MHz transmissions. Note that OFDM subcarriers are also referred to as *OFDM sub-channels* or *tones* in the literature.

The spectrum width of each subcarrier is defined by the subcarrier spacing $\Delta_{\mathrm{F}} = B/N_{\mathrm{ST}}$, which is a parameter specified in 802.11. The spacing has been historically set to $\Delta_{\mathrm{F}} = 312.5$ KHz, and has been reduced to $\Delta_{\mathrm{F}} = 78.125$ KHz in 802.11ax. For a fixed channel bandwidth, this entails that the number of modulated symbols that can be simultaneously transmitted over different subcarriers increases by four times, as depicted in Fig. 4.

OFDM transmissions can be implemented by upconverting each modulated symbol to the frequency of the k-th subcarrier ($k \in \{0, \ldots, N_{\mathrm{ST}}-1\}$). However, to improve the efficiency of the procedure, the signal is usually obtained by forwarding the N_{ST} modulated symbols—referred to as OFDM samples—though an IFFT block. The set of N_{ST} samples at the output of the IFFT block is referred to as an OFDM symbol, which is the fundamental unit of a Wi-Fi transmission. Importantly, a cyclic prefix (CP) is added at the beginning of each OFDM symbol by repeating the last portion of the same symbol. This acts as a GI between symbols, helping to reduce ISI even in environments with significant delay spread. The operations are performed in the base-band domain, and the signal is then upconverted to the center carrier frequency f_c. Being a_k the k-th modulated symbol (OFDM sample), the m-th OFDM symbol is obtained as

$$x_m(t) = \sum_{k=-N_{\mathrm{ST}}/2}^{N_{\mathrm{ST}}/2-1} a_{m,k}\, e^{j2\pi(f_c + k\Delta_{\mathrm{F}})t}. \tag{1}$$

OFDM effectively mitigates the frequency-selective fading responsible for ISI, i.e., the phenomenon for which delayed copies of a transmitted signal carrying previous symbols

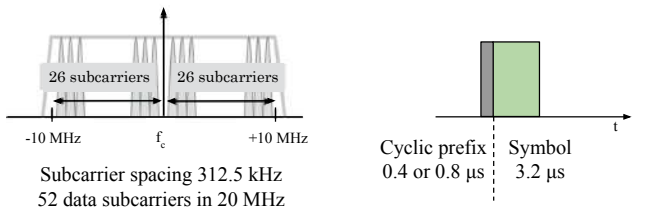

(a) PHY features in IEEE 802.11ac and earlier amendments.

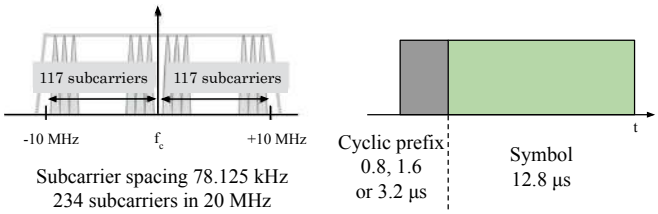

(b) PHY features in IEEE 802.11ax and beyond.

Fig. 4: OFDM subcarrier spacing with an example of subcarrier distribution on a 20 MHz channel (left), and corresponding OFDM symbol duration (right). The top (a) illustrates IEEE 802.11ac and earlier amendments with 52 data subcarriers at 312.5 kHz spacing and a 3.2 µs symbol (with 0.4 or 0.8 µs cyclic prefix). The bottom (b) shows IEEE 802.11ax and beyond with 234 subcarriers at 78.125 kHz spacing and a 12.8 µs symbol and longer cyclic prefix options.

overlap at the receiver with the signal carrying the current symbol. By design, the OFDM symbol duration is $T = 1/\Delta_{\mathrm{F}}$, resulting in an OFDM symbol time that is N_{ST} times longer than that of a single-carrier system. By appropriately choosing N_{ST}, the symbol duration can be made significantly larger than the typical delay spread of the channel, thereby reducing ISI and improving the overall system's robustness. The change in the subcarrier spacing introduced by 802.11ax goes in this direction. Indeed, it results in a symbol time of 12.8 µs compared to the 3.2 µs in previous amendments, thus enhancing performance in environments with high delay spreads, such as outdoor scenarios. Moreover, longer symbol duration reduces cyclic prefix overhead relative to the symbol length, increasing efficiency, especially indoors. However, longer symbols require more precise carrier frequency offset (CFO) correction to avoid signal degradation.

For what concerns the cyclic prefix duration, the legacy value is 0.8 µs, ideal for indoor environments where delay spread is minimal, as it minimizes overhead. In addition to this, 802.11ax introduced two more options: 1.6 µs and 3.2 µs. The 1.6 µs cyclic prefix strikes a balance between efficiency and robustness, making it suitable for outdoor use and uplink MU-MIMO/OFDMA transmissions. The 3.2 µs cyclic prefix is specific for extreme delay spread environments, such as outdoor uplink MU-MIMO/OFDMA, where longer GIs prevent ISI by ensuring that reflections of one OFDM symbol do not overlap with the subsequent one.

Latest 802.11bn Enhancements: Rate adaptation selects the best transmission parameters based on current channel and noise conditions. However, it may react too slowly to effectively handle rapidly time-varying interference. To address this limitation, 802.11bn is introducing receiver-based interference mitigation (IM) techniques. IM enhances interference estimation and suppression by embedding additional pilot tones within the data portion of the PPDU. These pilots provide insight into the characteristics of the interference observed during reception. When the number of receiver antennas exceeds the number of spatial streams, IM pilots can support noise covariance estimation and receiver-side beamforming, allowing the receiver to suppress portions (or even the entirety) of the interference. A trade-off must be made in pilot tone allocation: denser pilot insertion improves interference tracking but reduces data throughput. Implementing IM also raises configuration challenges. Fixed pilot patterns are simple but inflexible. Flexible configurations offer a limited set of predefined layouts, balancing adaptability and complexity. Fully configurable modes provide maximum flexibility but at the cost of high signaling overhead and implementation difficulty. In practice, fixed or flexible configurations are preferred due to their reduced complexity. It is also important to note that IM may be less effective in downlink MU-MIMO scenarios, where receiver constraints can limit its interference suppression capabilities.

D. Multi-Antenna Systems

The availability of multiple antennas at the AP, and possibly, at the STAs allows improving the robustness to fading (*diversity gain*), the received SNR (*beamforming gain*), and the capacity (*multiplexing gain*) of a Wi-Fi system by manipulating the data streams before and/or after their transmission over the MIMO channel (*precoding*) [68]. The key idea behind these processing techniques is that, if the antennas in a MIMO system are spaced by at least half a wavelength ($\lambda/2$), they provide space diversity, i.e., they generate independent fading channels, which can be leveraged to enhance the robustness of the system to noise and increase the capacity [67].

The diversity gain is linked to the fact that signals transmitted over such independent fading channels are unlikely to experience deep fades simultaneously. In turn, robustness to fading can be achieved by transmitting the same signal over the different antennas at the transmitter and combining the collected signals at the receiver, obtaining a signal where the fading is reduced. This technique is useful in scenarios with low SNR [69].

Beamforming gain is achieved by modifying the complex modulated symbols before transmission to compensate for the channel impairments and increase the SNR at the receiver(s) antennas. Note that when a technique leads to an increase in the SNR plus an increase in robustness and/or in capacity, it is referred to as beamforming in 802.11.

Finally, the independent fading channels generated by the multiple antenna pairs can be used for multiplexing, i.e., to transmit multiple data streams simultaneously using the same time-frequency resource. The number of data streams that can be simultaneously transmitted is constrained by the channel rank and the condition number. The rank indicates the number of linearly independent rows (or columns) of the $N_{\mathrm{RX}} \times N_{\mathrm{TX}}$-dimensional channel matrix, where N_{RX} and N_{TX} are respectively the total numbers of antennas at the receiver and transmitter sides. The condition number relates to their quality. From a physical standpoint, the rank depends on the number of uncorrelated transmission paths available between

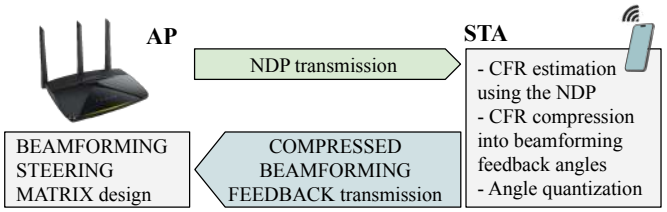

Fig. 5: Explicit sounding procedure for beamforming steering matrix calculation at the AP (beamformer). The AP first transmits a null data packet (NDP) to the STA, which estimates the channel frequency response (CFR), compresses it into beamforming feedback angles with quantization, and returns this information as a compressed beamforming feedback frame, enabling the AP to design the steering matrix for subsequent beamformed transmissions.

the transmitter and the receiver. In single user (SU)-MIMO, the multi-path propagation in non-line-of-sight (NLOS) scenarios offers the degrees of freedom needed for multiplexing, while line-of-sight (LOS)-dominated settings reflect in high channel matrices' condition numbers, making it challenging to transmit multiple streams. In addition to this, the 802.11 standard defines the maximum number of spatial streams that can be simultaneously transmitted: 802.11be and 802.11bn-compliant devices shall support up to 8 streams [66].

The mapping and spatial processing between the stream(s) to be transmitted (N_{SS}) and the transmitted antennas (N_{TX}) is represented by the yellow box in Fig. 3 and is performed by applying specific spatial matrices to the modulated and coded symbols obtained from the input bits, as detailed in Section III-B. The 802.11 standard provides some examples of the spatial matrices that can be used, but the implementation is not restricted to them. The simplest approach consists of applying a direct mapping, which can be an identity matrix mapping or a diagonal matrix where the diagonal elements are the cyclic shifts in the time domain. When the number of data streams to be transmitted is less than the transmitting antennas, a spatial expansion matrix can be used to replicate the streams over different antennas. The columns of the spatial expansion matrix should be orthogonal among them, and the matrix can also be different for the different OFDM subcarriers. If MIMO channel information is available at the beamformer, a *beamforming steering matrix* designed based on such knowledge can be applied to improve the decoding at the receiver and transmit multiple streams simultaneously. This approach is referred to as transmit beamforming (TxBF) in 802.11, and the transmitter and receiver are referred to as the beamformer and beamformee, respectively. When multiple streams are to be transmitted, they are simultaneously output through all the beamformer antennas using the combination (precoding) factors defined by the beamforming steering matrix. This technique is identified as SU-MIMO. The beamforming steering matrix may be devised to minimize the interference among data streams and improve the decoding quality at the different beamformee antennas. However, note that other types of precoders are not precluded, as the beamforming steering matrix is implementation-dependent.

The beamforming procedure is detailed in the following. We will focus on TxBF/SU-MIMO, while the extension to the MU-MIMO case is detailed in Section V-B. At first, a preliminary *channel sounding* phase allows the AP (beamformer) to obtain the channel frequency response (CFR) over the different antennas and OFDM subcarriers. As introduced before, this channel knowledge at the beamformer is needed to shape the beamforming steering matrix. Channel sounding is continuously performed—at a suggested rate of 10 ms [70]—to provide the AP with the most updated channel information to generate proper precoding. Theoretically, two approaches can be followed to acquire the channel information, namely, implicit sounding and explicit sounding. Implicit sounding leverages the channel reciprocity between uplink and downlink channels, and is more convenient from a spectrum efficiency perspective as it requires less control data transmissions with respect to explicit sounding. However, it requires precise hardware calibration to avoid mismatches that could compromise channel reciprocity and degrade performance. For this reason, explicit sounding, which entails obtaining the estimate of the downlink channel from the STA, is the only sounding mechanism standardized in 802.11 so far. This approach directly provides the AP with the downlink channel information, and does not require relying on channel reciprocity, thus avoiding calibration. The procedure is summarized in Fig. 5, and proceeds as follow [71], [72].

- The AP transmits a sounding frame, named Null Data Packet (NDP), to allow the STA measuring the response of the downlink channel (green arrow in Fig. 5). This is preceded by a Null Data Packet Announcement (NDPA) frame, which informs the STA about the start of the sounding. To enable channel estimation at the STA, the NDP contains a number of LTFs equal to the number of streams to be transmitted (N_{SS}). These LTFs are included as different subsequent fields in the PPDU and, in turn, are transmitted sequentially in time. Specifically, each LTF is transmitted through the N_{TX} AP antennas through CSD (see Section III-A) and applying an orthogonal mapping over the antennas and LTFs. This orthogonal mapping allows the STA to obtain a number of independent equations sufficient for the estimation of the $N_{\mathrm{ss}} \times N_{\mathrm{TX}}$-dimensional CFR matrix for each OFDM subcarrier.
- The STA uses the known LTFs in the NDP to estimate the CFR, which describes how the signal propagates from the transmitter to the receiver for each pair of transmitter and receiver antennas over the different OFDM subcarriers. In turn, the CFR, denoted by $\mathbf{H}$, is an $N_{\mathrm{RX}} \times N_{\mathrm{TX}} \times (N_{\mathrm{SD}} + N_{\mathrm{SP}})$-dimensional complex-valued matrix, where N_{RX} and N_{TX} are the numbers of receiver and transmitter antennas, respectively, while $N_{\mathrm{SD}} + N_{\mathrm{SP}}$ is the number of OFDM subcarriers including data and pilots, as described in Section III-C.
- This information should then be fed back to the AP. To do it quickly and efficiently, the $N_{\mathrm{RX}} \times N_{\mathrm{TX}}$ CFR matrices for each OFDM subcarrier k are first compressed through singular value decomposition (SVD), and only the right unitary matrices are considered for the feedback, as it has been shown that the performance of using them for beamforming is comparable to using the complete CFR.

In particular, only the sub-matrices consisting of the first N_{SS} columns of the right unitary matrices, referred to as $\mathbf{V}_k$, $\forall k$, are retained. As a further compression, each $\mathbf{V}_k$ matrix is decomposed into rotation matrices, which are fully defined by a set of angular values denoted by ϕ and ψ. In turn, feeding back such angles is sufficient for the bemaformer to retrieve the $\mathbf{V}_k$ matrix. The number of angles depends on N_{TX} and N_{SS} [73]. As a final step, the angles are quantized with a number of bits specified in the standard, and fed back inside a compressed channel feedback frame (blue arrow in Fig. 5) [74]. To further reduce overhead, the 802.11 standard defines grouping strategies to transmit the feedback for a subset $\hat{N}_{\mathrm{SD}} = N_{\mathrm{SD}}/N_g$ of the OFDM subcarriers, where N_g is the grouping factor, instead of feeding back the angles computed for all of them. Moreover, the compressed channel feedback frame is transmitted unencrypted to reduce latency, which would let the channel estimate be outdated when reaching the AP [71].

- Upon receiving the compressed beamforming feedback frame, the AP retrieves the angles and uses them to reconstruct the $\mathbf{V}_k$ matrices for all the subcarriers, which are used to define the beamforming steering matrix.

The procedure just described is known as non-trigger-based (TB) sounding and can only be used for SU-MIMO. Another channel reporting procedure, namely TB sounding, will be detailed in Section V. Such a procedure is mostly used for MU-MIMO feedback but can also be used to simultaneously collect multiple SU-MIMO feedback frames.

In addition to the beamforming feedback matrix, 802.11ax has introduced the so-called channel quality indicator (CQI) feedback, entailing the transmission to the beamformer of the singular values obtained from the SVD decomposition of the CFR matrix. This data is useful for MCS selection. Once the channel feedback is obtained, the beamformer derives the beamforming steering matrix through an implementation-dependent procedure. Possible strategies are zero-forcing (ZF) and minimum mean-square-error (MMSE) precoding. Through ZF, the steering matrix $\mathbf{Q}_k$ is obtained by applying channel inversion: $\mathbf{Q}_k = \tilde{\mathbf{V}}_\mathbf{k}\left(\tilde{\mathbf{V}}_\mathbf{k}^\dagger \tilde{\mathbf{V}}_\mathbf{k}\right)^{-1}$. The MMSE method includes an additional regularization term that improves reception performance in low SNR scenarios [75]. An example of applying the beamforming steering matrix is depicted in Fig. 6, where the STA is requesting two data streams. After data transmission, the beamformee may use MMSE equalization to retrieve and decode the received streams. The decoding matrix is usually obtained using the CFR estimated on the LTF included in the transmitted frame, and allows separating the different streams simultaneously transmitted [76], [77].

Overall, while explicit sounding is the best strategy to provide the beamformer with precise information to shape the steering matrix, it comes with significant overhead due to the feedback transmission, which increases with the number of transmit antennas, spatial streams, and bandwidth [8]. For example, to support 16 data streams (a number considered for 802.11be but not included in the amendment), with 16 antennas at the beamformer and 320 MHz bandwidth, the

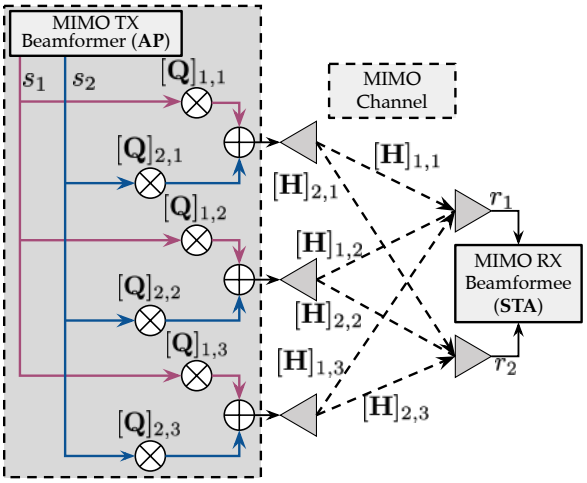

Fig. 6: Example of SU-MIMO beamforming with one STA (MIMO RX Beamformee, on the right) requesting two data streams s_1 and s_2. The AP (MIMO TX Beamformer, on the left) obtains the $\mathbf{Q}$ precoding weight matrix using the CFR of the MIMO channel $\mathbf{H}$. $[\mathbf{A}]_{x,y}$ indicates the entry of the matrix in row x and column y. The two data streams s_1 and s_2 are multiplied by the corresponding antenna weights before transmission. r_1 and r_2 represent the signals collected at the two receiving antennas.

feedback size—applying the most aggressive compression and quantization supported by the 802.11 standard—would be about 22.4 kB (120 ϕ angles quantized with 4 bits each and 120 ψ angles quantized with 2 bits each, for each of the $\hat{N}_{\mathrm{SD}} = 249$ reported subcarriers, i.e., $N_{\mathrm{SD}} = 3984$ data subcarriers with a grouping of $N_g = 16$ [78]) corresponding to an airtime overhead of about 7.5 ms (considering a typical feedback transmission rate of 24 Mb/s for high reliability), which is a large amount considering the typical sounding rate of 10 ms. As such, an open research avenue focuses on reducing the overhead of the explicit sounding process. This includes innovations in compressed feedback matrices and optimized channel sounding protocols. As for the former, researchers are proposing to use deep learning strategies in place of the standard feedback compression algorithm to further reduce the airtime overhead linked with feedback transmission [79]–[82] (see Appendix D). Regarding the optimization of the channel sounding procedure, a possible approach is to adapt the sounding interval to the variability of the channel as initially proposed in [83], [84]. Another strategy consists of using some side information to reduce the sounding rate [85]. None of these solutions is currently implemented in the 802.11 standard, but their integration will be key to increase the number of antennas of a MIMO system and in turn its multiplexing gain.

E. Maximum Achievable PHY Data Rate

Taking as an example the most recent 802.11be amendment, the maximum data rate achievable at the physical layer for the specific combination of MCS, OFDM and MIMO parameters introduced before is obtained as

$$\text{Max. PHY data rate} = \frac{N_{\mathrm{SD}} \cdot N_{\mathrm{CBPS}} \cdot R \cdot N_{\mathrm{SS}}}{T_{\mathrm{SYM}}} \qquad (2)$$

where:

- $N_{\mathrm{SD}} \in \{234, 468, 936, 1960, 3920\}$ is the number of data subcarriers for channel bandwidths {20 MHz, 40 MHz, 80 MHz, 160 MHz, 320 MHz}, respectively.
- $N_{\mathrm{CBPS}} \in \{1, 2, 4, 6, 8, 10, 12\}$ is the number of coded bits per OFDM symbol for modulations {BPSK, QPSK, 16QAM, 64QAM, 256QAM, 1024QAM, 4096QAM}, respectively.
- $R \in \{\frac{1}{2}, \frac{2}{3}, \frac{3}{4}, \frac{5}{6}\}$ is the coding rate.
- $N_{\mathrm{SS}} \in \{1, 2, 4, 8\}$ is the number of spatial streams.
- $T_{\mathrm{SYM}} = 12.8\,\mu\mathrm{s} + T_{\mathrm{GI}}$ is the OFDM symbol duration, including a GI $T_{\mathrm{GI}} \in \{0.8\,\mathrm{us}, 1.6\,\mathrm{us}, 3.2\,\mathrm{us}\}$.

While Eq. (2) refers to the peak data rate on a single link, MLO described in Section VII can further increase the data rate in 802.11be by using several physical links in parallel.

F. Physical Layer Protocol Data Units (PPDUs)

The PPDU represents the single unit of information to be transmitted over the wireless channel and consists of two main parts: the PHY preamble and the data payload. The preamble is critical for facilitating successful communication, as it contains information needed at the receiver to perform automatic gain control, timing synchronization, channel estimation, decoding, and demodulation. The data payload includes one or more PSDUs (as described later, in Section IV-B), where a PSDU entails the Application data and the protocol control information (PCI) (header) for the Transport, Network and MAC layers functionalities. Various PPDU formats have been defined in the different versions of the 802.11 standard, each optimized for specific transmission conditions and user scenarios. In particular, the design of the preamble and field structures has evolved across Wi-Fi generations to accommodate new features, while maintaining backward compatibility with earlier amendments.

1) PPDU Formats: Two PPDU formats are defined in 802.11be, building upon two formats originally defined in 802.11ax. Additionally, as part of the 802.11bn standardization process, an additional PPDU format is being considered for long-range communications. The three formats are depicted in Fig. 7, and described next. Note that, as part of the payload, after the Data field, a Packet Extension (PE) field is added to take the channel occupied and provide the receiver with more time to process the received data.

- *Multi-user PPDU (EHT/UHR MU)*: Used for both single-user and multi-user downlink transmissions and non-TB single-user uplink transmissions. Multi-user transmissions can be implemented through either OFDMA or MU-MIMO, depending on the RU allocation (see Section V). The EHT MU PPDU includes an EHT-Signal field (SIG) field, which provides details for decoding this specific PPDU format.
- *Trigger-based PPDU (EHT/UHR TB)*: Used for uplink transmissions in response to a trigger frame sent by the AP. This format is crucial for coordinating uplink OFDMA and MU-MIMO (see Section V), reducing contention and increasing efficiency when multiple STAs transmit simultaneously. Additionally, the EHT-STF is twice as long as in the EHT MU PPDU version. This allows improving performance, particularly in environments with multiple simultaneous uplink transmissions. The EHT-SIG field is not present as the information has already been communicated to the STA inside the trigger [86].
- *ELR PPDU (UHR ELR)*: Introduced in 802.11bn to improve the range of uplink transmissions and overcome the link budget imbalance between downlink and uplink channels, which occurs because APs often transmits at higher power levels than STAs [87], [88]. The ELR PPDU is configured for single-user transmission over a 20 MHz channel, supporting MCS 0 and 1 with four repetitions of frequency domain duplications across a 52-tone RU. It operates in the 2.4 GHz band for both uplink and downlink, and in the 5 GHz and 6 GHz bands for uplink only. This configuration mitigates uplink signal limitations while supporting data rates higher than legacy DSSS. Hence, ELR improves uplink range and reliability, and is thus particularly useful in scenarios where Wi-Fi coverage extends beyond typical indoor environments, such as outdoor campuses, industrial sites, and low-data-rate applications like IoT and telemetry. To allow for better detection at the receiver, the transmitter of an ELR PPDU is expected to synchronize its transmit clock to that of the AP. This can be done by compensating the frequency offset relative to packets received from the AP.

2) Preamble Design: A critical objective of 802.11be and now 802.11bn is maintaining backward compatibility with earlier Wi-Fi amendments, while operating across the 2.4 GHz, 5 GHz, and 6 GHz bands. The preamble plays a key role in facilitating this compatibility, providing key transmission parameters such as the MCS and the number of spatial streams. For this reason, their preambles include legacy (L) fields, namely L-STF, L-LTF, and L-SIG, which enable some common initial processing steps on the signal(s) collected at the receiver. To start with, the L-STF allows detecting the starting of the PPDU, and performing automatic gain control (AGC) and coarse carrier recovery. After that, the L-LTF enables fine carrier and timing recovery and channel estimation for the equalization of the subsequent SIG fields in the preamble. Once equalized, the L-SIG together with its replication, referred to as repeated legacy SIG (RL-SIG), provides the information to recognize if the PPDU is in one of the 802.11be formats (MU or TB). If so, it also indicates the specific format. Note that before the IEEE 802.11n amendment, only the L-STF and L-LTF filed were indicated as 'PHY preamble', while the 'L-SIG was referred to as 'PHY header'. Instead, the current standard terminology makes less use of 'header' for the PHY layer fields and refer to all the fields added to the PSDU as 'preamble'.

Following the legacy fields in the preamble, 802.11be introduced a universal SIG (U-SIG) field to simplify the detection of different PHY versions, and ensure compatibility with both older and future devices. The U-SIG spans two OFDM symbols, and contains both version-independent and version-dependent bits [86]. Among the version-dependent bits, the U-SIG indicates whether coordinated SR (Co-SR) or coordinated

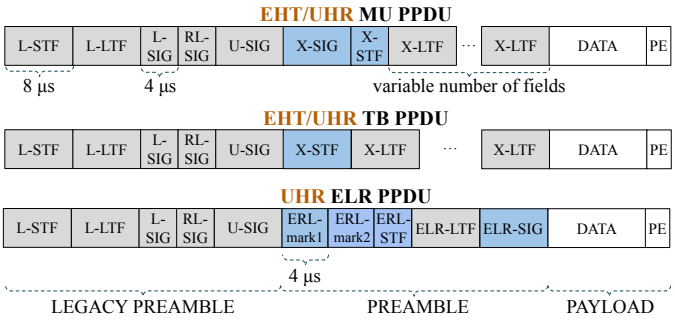

Fig. 7: PPDU formats in IEEE 802.11be/bn. The fields that differ in the EHT MU and EHT TB PPDU formats are highlighted in blue. The width of the gray and blue boxes is associated with their time duration. 'X' stands for either 'EHT' or 'UHR'. The length of the 'DATA' field is variable.

beamforming (Co-BF) are used (see Section VIII-B).

The subsequent EHT/UHR-SIG field is only present in EHT/UHR MU PPDUs, and provides information regarding the use of OFDMA or MU-MIMO as the multi-user technology, including the frequency and spatial resources allocated to each user, and their assigned MCS (see Section V).

Finally, both the MU and TB PPDU preambles include 802.11 version-specific short and long training fields: EHT/UHR-STF and EHT/UHR-LTF. The former is used to refine the AGC, carrier, and timing. The latter is used to obtain a channel estimate to equalize the Data field. The number of EHT/UHR-LTFs is at least equal to the number of spatial streams per RU.

The UHR ELR PPDU maintains the same structure of the two EHT-PPDU formats for the legacy part of the preamble, including the U-SIG field. In addition, two ELR-Mark symbols are added for ELR mode classification [89]. ELR-STF and ELR-LTF are included as in 802.11be. After the ELR-LTF, a ELR-SIG is included to indicate the transmission direction, the MCS and coding strategy, and the length of the payload.

IV. Wi-Fi Medium Access Control

Wi-Fi generations have consistently introduced substantial enhancements at the MAC layer to improve network efficiency, reduce latency, and ensure fair access to the wireless medium, especially in dense environments. However, to understand how Wi-Fi works and why its various enhancements are designed the way they are, it is essential to master the core components of its channel access protocol and understand their broader implications.

This section begins with an overview of DCF, the foundational MAC protocol on which all subsequent enhancements are built. One such enhancement is enhanced distributed channel access (EDCA), which serves as the default access mechanism in current 802.11 devices.[2] We also introduce prioritized enhanced distributed channel access (P-EDCA), a new feature in 802.11bn that addresses a key limitation of EDCA: its inability to effectively manage multiple high-priority accesses within a single BSS and multiple OBSS. Next, we examine packet aggregation and more deterministic access mechanisms, such as AP scheduled transmissions, which extend EDCA by granting the AP greater control over uplink (UL) and peer-to-peer (P2P) scheduling, and restricted TWT (R-TWT), a feature for latency control. Finally, we present wideband access techniques—including channel bonding, preamble puncturing, and two additional features introduced in 802.11bn, namely NPCA and DSO—and SR for coexistence among OBSSs.

[2]Note that we do not cover the point coordination function (PCF), now considered obsolete, and the hybrid coordination function controlled channel access (HCCA), as these mechanisms have seen limited adoption in commercial deployments.

A. Enhanced Distributed Channel Access

The DCF protocol has been a fundamental part of 802.11 since its inception, enabling devices to share the unlicensed spectrum fairly and efficiently. It implements CSMA with CA through a random backoff countdown and an automatic repeat request protocol (ARQ) to handle packet errors and retransmissions.

1) Distributed Coordination Function: The basic principle of CSMA/CA is simple: before transmitting, a device listens to the channel to determine if it is free. If the channel is busy, the device waits until the channel becomes free to transmit. Devices utilize two methods to check whether the medium is free: (i) physical carrier sensing, which detects energy above -62 dBm to determine channel occupancy; and (ii) virtual carrier sensing, in which the device first decodes the PHY header when the detected energy is above -82 dBm, and then decodes the MAC header to read the Duration/ID field that carries the medium occupancy information. For pre-HE devices, the Duration field in the MAC header is utilized to update the network allocation vector (NAV) counter, which reduces the need for continuous physical carrier sensing and helps conserve power. Importantly, this update can only occur after the complete MAC protocol data unit (MPDU) has been received and the Frame Check Sequence (FCS) has been validated. In HE devices and subsequent amendments, the HE-SIG-A preamble introduces the transmission opportunity (TXOP) Duration field, which specifies the remaining time in the current TXOP. This mechanism enables expedited NAV updates, as only the PPDU preamble needs to be decoded. A similar approach is employed in EHT and UHR devices, where the U-SIG preamble field provides equivalent functionality.

However, relying solely on checking whether the channel is free before transmitting does not prevent multiple devices, concurrently sensing the channel, from transmitting simultaneously, which can result in a collision. To minimize such collisions, 802.11 employs a *collision avoidance* strategy based on the binary exponential backoff (BEB) mechanism. With BEB, after detecting that the medium is idle, a device waits for a random backoff time, uniformly selected between 0 and a contention window (CW). The CW is initially set to a minimum value ($\mathrm{CW}_{\min}$) to enable fast channel access, and it doubles after each failed transmission attempt—up to a maximum limit ($\mathrm{CW}_{\max}$)—to improve network stability and throughput under high contention. Importantly, during the backoff process, a device pauses its countdown whenever it detects that the channel is busy, resuming only after the channel has been idle again for at least a distributed interframe space (DIFS). As a result, the actual backoff duration is

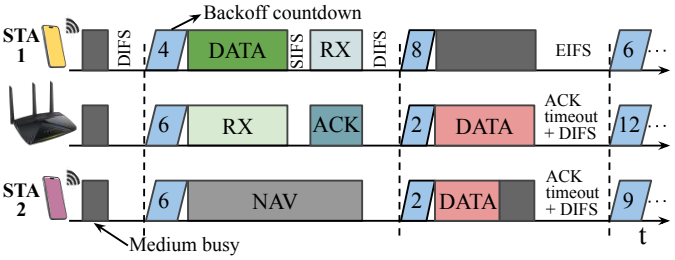

Fig. 8: Distributed Coordination Function (DCF) operation for two contending STAs and one AP. The timeline shows DIFS/SIFS/EIFS intervals, DATA and ACK exchanges, NAV periods, ACK timeouts, and intervals when the medium is sensed busy (grey segments); the numbers in the blue boxes represent the backoff countdown values at each station.

uncertain, not only due to the initial random value, but also because of potential interruptions caused by transmissions from other devices.

When a transmission fails due to channel errors or collisions, current 802.11 systems rely on ARQ to enable retransmissions and ensure reliable communication. Retransmissions are triggered when the expected Acknowledgment (ACK) is not received by the transmitter of the MPDU after an ACK timeout. The transmitter of the MPDU then repeats the process of sending the packet and waiting for the ACK until it is successfully received. This mechanism allows 802.11 devices to recover from errors through multiple retransmission attempts, providing robustness against adverse channel conditions and collisions caused by simultaneous transmissions from multiple devices. However, ARQ retransmissions often necessitate conservative link adaptation, typically involving a reduction in the MCS to maintain reliable communication (see Section III-B).

An example of the DCF operation is illustrated in Fig. 8, where one AP and two STAs share the same channel. We observe that the first transmission is successful (only STA 1 accesses the medium), while the second results in a collision (both the AP and STA 2 access the medium).

The interval between receiving an MPDU and transmitting the corresponding ACK is denoted as short interframe space (SIFS) and represents the time required to switch the transceiver circuitry between transmit and receive modes. The SIFS is shorter than the DIFS to ensure that immediate transmissions such as ACKs have priority over other stations contending for the medium. Devices that detect the channel as busy but are unable to decode the frame header, and therefore cannot update their NAV, must wait for an extended interframe space (EIFS) instead of a DIFS before resuming their backoff countdown. The EIFS duration roughly corresponds to the time required to account for the potential ACK timeout plus a DIFS, protecting ongoing transmissions that the device could not properly decode.

Analyzing the performance of the DCF under various scenarios has been a key research focus since the inception of Wi-Fi, primarily due to the complex interactions between traffic load, the number of contending devices, and the CSMA/CA mechanism. A seminal contribution in this domain is the work of Bianchi [90], which has served as the foundation for most 802.11 performance analyses over the past 25 years.

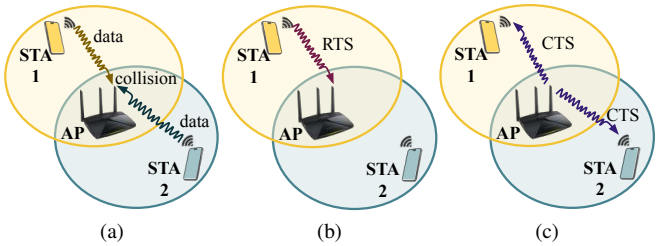

Fig. 9: RTS/CTS mechanism: (a) STA 1 and STA 2 cannot hear each other, AP hears both; (b) STA 1 sends RTS to AP; (c) CTS message prevents a collision between STA 1 and STA 2.

2) RTS/CTS: While DCF effectively manages medium access in most scenarios, it suffers in the presence of hidden nodes. For example, as shown in Fig. 9a, if STA 1 and STA 2 cannot detect each other's transmissions but both communicate with the same AP, collisions may occur at the AP. These collisions cannot be resolved using DCF, as the devices are unaware of each other's activity. To address this issue, 802.11 employs the request to send (RTS) and clear to send (CTS) mechanism. When initiating a transmission, an STA or AP can send an RTS message, which is shorter and more robust as it is transmitted using a low MCS than a typical data frame. The recipient responds with a CTS message, signaling that the channel is reserved by the RTS sender. The RTS and CTS messages include the duration of the subsequent data transmission, allowing all devices receiving the RTS, the CTS, or both, to set up their NAVs. This virtual carrier sensing mechanism ensures that other devices in the vicinity are informed that the channel is in use, thereby preventing them to further try to access the medium until current transmission ends, as illustrated in Fig. 9b and Fig. 9c.

3) Traffic Differentiation and Prioritization: In the IEEE 802.11e amendment, published in 2005 with the goal of providing traffic differentiation and prioritization capabilities, EDCA was introduced to replace DCF, becoming the default channel access mechanism in 802.11 networks [91]. The EDCA architecture introduced several new features:

- *Four access categories (ACs)* are defined: voice (VO), video (VI), best effort (BE), and background (BK), ordered from highest to lowest priority. Each AC is represented by an independent buffer and EDCA function.
- *Traffic classification:* Packets arriving from the network layer are classified into different ACs based on information in the IP headers, e.g., the Differentiated Services Code Point (DSCP) field.
- *Parallel EDCA functions:* Channel access contention is handled in parallel by all ACs with queued packets. Distinct arbitrary inter-frame space (AIFS), $\mathrm{CW}_{\min}$, and $\mathrm{CW}_{\max}$ values are defined for each AC. The AIFS parameter in EDCA generalizes the DIFS used in DCF; in fact, DIFS is equivalent to the AIFS used for the video access category in default configurations. A maximum TXOP duration is also specified per AC, limiting the duration of channel occupancy. Higher-priority traffic (e.g., voice and video) benefits from shorter inter-frame spaces and smaller CWs, enabling reduced access delays and improved performance for time-sensitive applications.

- *Internal collision resolution:* If two or more ACs complete their backoff at the same time, an internal collision occurs. In such cases, an internal resolution mechanism selects the highest priority AC involved in the collision to proceed with transmission.

While EDCA provides effective traffic differentiation when high-priority traffic is sparse, it struggles under high traffic load. For instance, when multiple uplink flows use the voice or video ACs, the resulting competition can lead to frequent collisions and performance degradation due to the use of low $\mathrm{CW}_{\min}$ and $\mathrm{CW}_{\max}$ values. In such scenarios, these flows may perform better if assigned to lower-priority access categories.

Latest 802.11bn Enhancements: P-EDCA is an enhancement of the EDCA mechanism that reduces the access delay distribution tail for VO access category traffic [11], [92]. The use of P-EDCA by an 802.11bn STA shall be balanced to reduce the impact on STAs that do not support this mechanism, focusing on the fairness guarantees. Consequently, P-EDCA is selectively activated for STAs that have repeatedly failed to transmit VO AC traffic using the legacy EDCA procedure. The basic operating principle of P-EDCA involves transmitting a Defer Signal (CTS frame) after the P-EDCA STAs verify that the channel has remained idle, bypassing the legacy backoff procedure. This mechanism grants P-EDCA STAs a priority advantage over other STAs in accessing the channel. As a result, any ongoing backoff processes from non-P-EDCA STAs are interrupted, allowing the P-EDCA contention process to initiate without competition. Following the Defer Signal, the P-EDCA STA initiates a new random backoff using the default AIFS and CW parameters for the VO access category. Once the STA wins the contention, it must protect the channel using an RTS/CTS exchange before transmitting any data frame, following the legacy procedure.

B. Packet Aggregation

Packet aggregation is a fundamental technique in 802.11 networks to improve airtime efficiency by bundling multiple data packets into a single transmission [93]. Originally introduced in 802.11n, it has become essential for supporting high-throughput applications such as large file transfers and high-quality video streaming. Indeed, without packet aggregation, each data packet requires its own PHY preamble and MAC header, resulting in considerable protocol overhead. In addition, control frames such as RTS/CTS and ACKs, interframe spaces, and the initial backoff duration must all be considered in the total airtime. By assembling multiple packets into a single transmission, packet aggregation minimizes these overheads, as they are included only once per transmission. This not only reduces contention and the likelihood of collisions (as fewer transmission attempts are needed) but also improves overall system efficiency, leading to higher network throughput and lower latency. 802.11 employs two primary types of packet aggregation, namely Aggregate MAC Protocol Data Unit (A-MPDU) and Aggregate MAC Service Data Unit (A-MSDU), illustrated in Fig. 10 and detailed next.

A-MPDU aggregation combines multiple MPDUs into a single PPDU, separating them through MPDU delimiters.

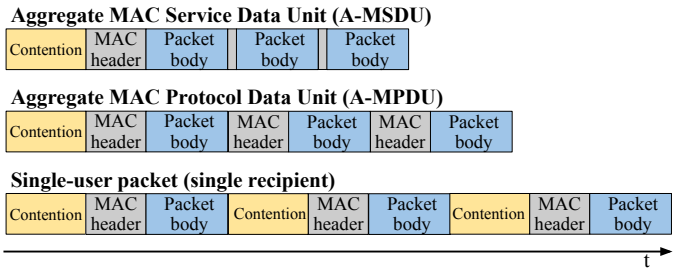

Fig. 10: A-MSDU and A-MPDU packet aggregation strategies versus single-packet transmission, highlighting how aggregation reduces contention overhead and improves MAC efficiency. Top: A-MSDU, where multiple packet bodies are encapsulated under a single MAC header after one contention period. Middle: A-MPDU, where several complete MAC Protocol Data Units (each with its own header) are aggregated into one transmission opportunity. Bottom: separate single-user packets requiring repeated contention.

Each MPDU retains its own MAC header and trailer, but the aggregated transmission uses a single PHY preamble and interframe space, minimizing PHY overhead. The receiver processes the aggregated MPDUs and sends a Block Acknowledgment (BA) to confirm reception, which reduces the number of required ACKs and the associated airtime overhead. A-MPDU aggregation is particularly effective in scenarios involving large data transmissions, such as video streaming or bulk data transfers.

A-MSDU aggregation operates at a higher level than A-MPDU by combining multiple MAC service data units (MSDUs) into a single MPDU. Unlike A-MPDU, A-MSDU aggregation shares a single MAC header and trailer for all the aggregated MSDUs, which further reduces overhead. All aggregated MSDUs share one MAC header and trailer, resulting in lower overhead compared to A-MPDU. A-MSDU aggregation is efficient for smaller packet sizes, where reducing MAC overhead significantly impacts performance. However, A-MSDU is less resilient to errors: if a single MSDU in the aggregated frame is corrupted, the entire MPDU must be retransmitted.

In recent 802.11 amendments, packet aggregation continues to evolve to support the growing demands of modern wireless networks, particularly when combined with the use of wideband channels, MU-MIMO, and OFDMA technologies. Each 802.11 amendment has progressively increased the number of MPDUs that can be aggregated. In 802.11be, this number has been expanded to 1024 MPDUs, representing a substantial enhancement in throughput and transmission efficiency.

C. AP Scheduled Transmissions

In dense and highly congested environments, the contention among devices to access the medium results in significant inefficiencies, including frequent collisions and transmission delays, particularly when numerous devices attempt to communicate simultaneously. AP-triggered transmissions, introduced in 802.11ax, change this dynamic by allowing the AP to take control of when and how STAs transmit. Instead of devices initiating communication on their own based on the regular backoff mechanism, the AP manages and schedules UL and P2P transmissions, reducing contention and improving overall network performance.

In this AP-scheduled mode, trigger-based transmissions allow the AP to enforce traffic specification (TSPEC) and stream classification service (SCS) requirements directly—bypassing EDCA contention—and perform per-flow aware scheduling within the BSS. While this provides IntServ-like control at the MAC level, it remains a local (intra-BSS) mechanism rather than a full end-to-end IntServ architecture.

1) Trigger-based UL Transmission: This is a key feature introduced in 802.11ax and enhanced in 802.11be with SCS and quality of service (QoS) management to improve UL transmission efficiency [94]. Trigger-based UL allows the AP to schedule UL transmissions from multiple STAs simultaneously using OFDMA (see Section V-A). To do so, the AP sends a trigger frame that allocates RUs to specific STAs, which transmit their data simultaneously using the assigned RUs immediately after receiving the trigger frame. The AP answers by sending a Multi-STA BA frame.

2) SCS and QoS Management: Triggered access optimization has been introduced in 802.11be to improve quality of experience (QoE) for real-time applications such as gaming, industrial control systems, and other delay-sensitive use cases that follow a bidirectional and periodic traffic pattern. These applications require predictable latency and guaranteed minimum data rates, which can be challenging in contention-based networks like Wi-Fi. The core idea behind triggered UL access optimization is the establishment of a desired UL access interval. This interval provides deterministic scheduling of UL transmissions, effectively bounding the delay while ensuring a minimum data rate. The mechanism minimizes jitter and delay by aligning transmission opportunities with traffic flow requirements, which is critical for time-sensitive applications.

In 802.11be, STAs can communicate their QoS expectations to the AP. This is achieved through the QoS Characteristics element, a standardized structure that conveys the key QoS parameters for a traffic flow. The QoS Characteristics element includes the following: parameters that define the QoS requirements of specific traffic flows (UL and downlink (DL)), support for SCS which enables classification of traffic streams for optimized handling, and integration with R-TWT for additional scheduling flexibility and latency control (see Section IV-D).

3) Triggered P2P Transmission: The concept of TXOP sharing was first introduced in 802.11ax, where multiple STAs were allowed to transmit simultaneously within a shared TXOP during multi-user transmissions such as in OFDMA. 802.11be extends this concept further with triggered P2P transmission, where the AP can share its TXOP with STAs, allowing them to use the allocated time for direct STA-to-STA communication, using the TXOP Sharing (Mode 2) mechanism. This mechanism supports direct STA-to-STA data exchanges, which are essential for applications such as wireless file transfers between devices (e.g., phone-to-printer communication), video streaming (e.g., laptop-to-monitor wireless display), and VR applications requiring low-latency peer communication.

In triggered P2P transmission, the AP can use favorable contention parameters (e.g., small CW values) to increase its chances of winning a TXOP. The AP transmits a multi-user RTS (MU-RTS) TXOP Sharing Trigger frame Mode 2, allocating a portion of the TXOP for P2P transmission. Other devices defer their transmissions, which improves spectrum utilization and reduces contention within the BSS. Moreover, direct communication between two STAs enhances airtime efficiency, as the transmitted data does not need to be duplicated when it passes through the AP. For example, in the scenario illustrated in Fig. 11, the AP gains a TXOP and allocates a portion of it to STA 1. STA 1 responds with a CTS frame and then uses the allocated time to send data directly to STA 2. STA 2, in turn, sends an ACK to STA 1. If the shared TXOP period ends within the remaining time, the AP can reclaim the rest of the TXOP duration to send its data.

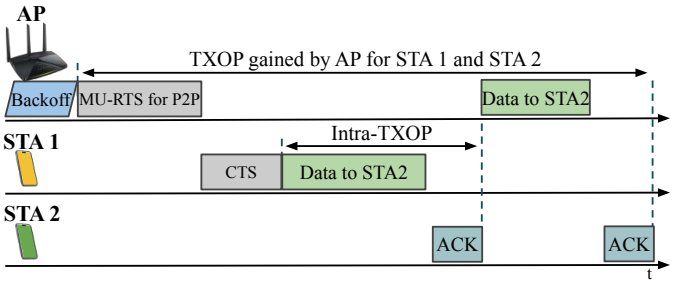

Fig. 11: Illustration of triggered P2P: TXOP Sharing Mode 2. The AP first wins the channel after backoff and gains a TXOP for STA 1 and STA 2 using an MU-RTS/CTS exchange, then, within the same intra-TXOP, triggers a direct data transmission from STA 1 to STA 2, which is completed with individual ACKs from STA 2 and the AP.

Latest 802.11bn Enhancements: 802.11bn aims to extend the concept of P2P to a P2P group, composed of two or more STAs. Similarly, the AP shares the obtained TXOP with a previously created P2P group, and these STAs will communicate with each other within the shared TXOP time length. The idea behind this mechanism is to offload the AP scheduling task for each P2P communication link, easing the AP workload, and maximizing the communication efficiency. Within the shared TXOP, various self-organizing mechanisms can be employed to coordinate channel access among the designated P2P group members, as detailed next.

- *Contention-based approach:* Each device that has a non-empty buffer for P2P traffic selects a random backoff counter, similar to the standard CSMA/CA process, and once a device wins the contention, it starts transmitting its data to its intended peer within the group.
- *Rule-based approach:* It predefines an order or pattern for devices within the P2P group to access the TXOP allocated by the AP. This approach eliminates the need for contention and focuses on deterministic scheduling.
- *Hybrid approach:* In some scenarios, a hybrid approach combining contention and rule-based methods may be employed. For example, devices with higher-priority traffic (e.g., latency-sensitive packets) could be given a shorter backoff window in a contention-based approach, ensuring they win more often. Alternatively, the P2P group could follow a rule-based order initially, switching to a contention-based approach only when there are changes in group membership or unexpected traffic patterns.

While triggered P2P transmission improves spectrum uti-

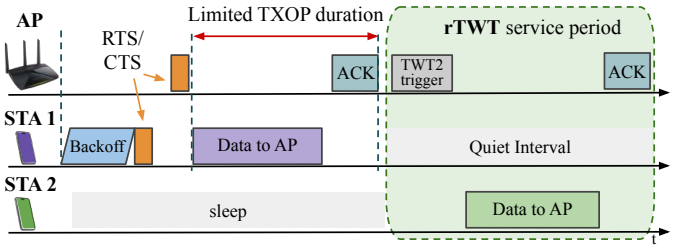

Fig. 12: Illustration of R-TWT operation. The timeline shows an AP first granting a limited TXOP to STA 1 via an RTS/CTS exchange, followed by an R-TWT service period during which STA 2 wakes up from sleep and transmits scheduled uplink data to the AP, while STA 1 remains silent during the quiet interval.

lization and reduces collisions, it requires the AP to have visibility into the buffer state of the participating STAs or rely on additional information to make efficient decisions about TXOP sharing. By enabling direct communication between STAs, this mechanism enhances network efficiency, reduces transmission delays, and supports latency-sensitive peer-to-peer applications.

D. Restricted Target Wake Time (R-TWT)

802.11be introduced R-TWT to support latency-sensitive applications by providing greater control over transmission timing. R-TWT builds upon TWT (described in Section VI-B), which was originally designed to reduce energy consumption.

To introduce determinism in transmissions, R-TWT addresses a key limitation of traditional TWT, namely, the lack of coordination among STAs regarding scheduled wake times. Indeed, in conventional TWT, other STAs might be unaware of an ongoing agreement, leading to potential collisions or inefficiencies. R-TWT overcomes this by ensuring that only STAs with an established R-TWT membership may access the channel during a designated restricted service period.

While R-TWT retains the foundational mechanisms of TWT (explained in detail in Section VI-B), it introduces several key enhancements for deterministic latency:

- During schedule negotiation, the AP and associated STAs establish an R-TWT agreement. This agreement defines the start times, durations, and periodicity of the service periods (SPs). It also includes a list of Traffic Identifiers (TIDs), allowing prioritization of traffic types within the SPs.
- To enforce the restricted access during an SP, the AP initiates explicit channel protection via control signaling, such as CTS-to-Self or Quiet Element frames, at the beginning of each SP. This signaling silences both compliant and non-compliant TWT/R-TWT STAs, ensuring predictable and interference-free transmissions.

Fig. 12 illustrates the operation of R-TWT, where STA 2 adheres to an R-TWT schedule (rTWT2), while STA 1 follows a conventional TWT schedule (TWT1) for energy savings. This coexistence demonstrates the flexibility of R-TWT, which can operate in parallel with other TWT-based mechanisms to optimize both latency and power consumption.

Although R-TWT enhances latency control, it introduces certain trade-offs that may impact spectrum efficiency. Specifically, to meet the strict timing of service periods, ongoing transmissions—such as the one from STA 1 in Fig. 12—may need to terminate early, reducing the number of packets that can be aggregated. In other cases, the requirement that all TXOPs finish before the start of an R-TWT service period may prevent new transmissions from being initiated, resulting in underutilized channel time. Moreover, a key limitation of R-TWT is that its operation is confined to the BSS, meaning it cannot regulate transmissions from STAs associated with neighboring OBSSs. To address this issue, multi-AP coordinated R-TWT is being specified (see Section VI-B), enabling the propagation of R-TWT schedules to neighboring OBSSs.

We note that R-TWT prioritizes deterministic behavior over maximum throughput. Despite its limitations, R-TWT enables predictable and bounded latency, substantially improving the suitability of 802.11be for delay-sensitive applications, such as real-time communication, AR/VR, and industrial automation.

E. Wideband Operations

Accessing the spectrum efficiently is key to achieving higher transmission rates, resulting in increased throughput and lower latency [95]. This property is fundamental to modern 802.11 amendments, enabling them to meet the demands of bandwidth-intensive applications such as 4K video streaming, virtual reality, and other high-data-rate services [96].

1) Channel Bonding: In 802.11ax, AP and STAs can dynamically select the channel width for each frame transmission, bonding 20 MHz channels into wider configurations such as 40, 80, or 160 MHz [95]. This adaptability allows the system to maximize throughput while accounting for real-time channel conditions. For example, an STA that gains access to the medium on a primary 20 MHz channel can expand the bandwidth by adding secondary channels, provided they are idle. The extension from 20 MHz to wider channels follows a hierarchical process as shown in Fig. 13:

- 40 MHz: Adds one 20 MHz secondary channel.
- 80 MHz: Adds two more secondary channels if available.
- 160 MHz: Combines four 40 MHz channels or eight 20 MHz channels.

This process relies on mechanisms such as virtual carrier sense and backoff on the primary channel, followed by quick clear channel assessment (CCA) checks (e.g., during a PCF interframe space (PIFS) period) on secondary channels just before transmission. Note that, for the sake of compatibility with legacy devices, channel access is performed on the primary channel only, and then extended to other channels if channel bonding is implemented.

802.11ax supports channel bonding of up to 160 MHz, achieved by aggregating eight 20 MHz channels, while 802.11be further expands this capability to 320 MHz channels in the 6 GHz band. The 6 GHz spectrum, initially accessible only to Wi-Fi 6E and Wi-Fi 7 devices, offers a pristine, low-interference environment, free from legacy devices and radar restrictions. This allows Wi-Fi 7 to unlock the full potential of channel bonding for next-generation applications. Specifically, the 802.11be-2024 amendment defines 320 MHz channels by doubling the 160 MHz tone plan used in 802.11ax [97]. To maximize efficiency in various spectrum regulatory regimes,

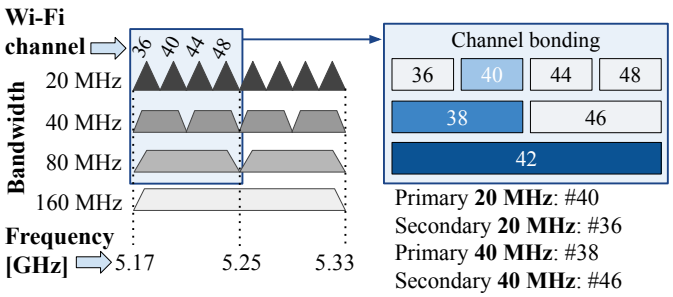

Fig. 13: Illustration of channel bonding in the 5 GHz band. Left: four adjacent 20 MHz channels (36, 40, 44, 48) can be combined into wider 40, 80, and 160 MHz bandwidths. Right: corresponding channel-bonding configuration, with the selected primary and secondary 20 MHz channels (#40 and #36) and the resulting primary and secondary 40 MHz channels (#38 and #46).

some of which may not have allocated the entire 1.2 GHz for unlicensed use, overlapping 320 MHz channels are defined, divided into two configurations: 320 MHz-1 and 320 MHz-2.

2) Preamble Puncturing: While channel bonding provides substantial performance benefits, it also introduces challenges, particularly in environments with heavy spectrum utilization. In the 5 GHz band, several factors constrain the availability of contiguous channels for bonding:

- *Radar interference and DFS requirements:* Some portions of the 5 GHz spectrum overlap with radar systems, requiring 802.11 devices to vacate these channels when radar signals are detected. DFS mandates periodic scanning for radar activity, which can disrupt channel bonding and limit the availability of wider channels.
- *Non-contiguous spectrum:* The 5 GHz band is fragmented, leaving limited room for wider channels. For example, in the USA, only five non-overlapping 80 MHz channels and one 160 MHz channel are available when DFS is in use. This scarcity makes it challenging to fully exploit channel bonding in dense deployments or large outdoor environments.
- *Transmit power regulations:* Power limitations in portions of the 5 GHz band reduce the usability of certain channels for high-power deployments intended to cover large areas with many users, further restricting effective channel bonding.
- *Legacy 802.11 devices:* Older devices operating on 20 MHz or 40 MHz channels can coexist with modern devices but often limit the ability of newer systems to fully utilize wide channels without interference or contention.

In scenarios where spectrum fragmentation is significant, such as in environments with legacy 802.11 networks operating on 20 MHz channels, it becomes increasingly challenging to find contiguous 80 or 160 MHz channels. For example, in a network capable of using an 80 MHz channel, if one of the secondary 20 MHz channels is frequently occupied, the network is restricted from utilizing the entire 80 MHz channel, even if other secondary channels remain idle most of the time. This leads to underutilization of available spectrum, reducing overall network efficiency.

To address this limitation, 802.11ax introduced preamble puncturing, which allows an AP to transmit over non-contiguous channels by skipping busy portions of secondary channels. By relaxing the requirement for contiguous secondary channels, preamble puncturing improves spectrum utilization, particularly in dense deployments where wide channels (e.g., 80 or 160 MHz) may not be fully available [48]. Note that preamble puncturing operates in conjunction with OFDMA, as non-contiguous subchannels must be allocated to different STAs. An example of the performance advantages of preamble puncturing is illustrated in Fig. 14. Without this feature, when secondary channels are frequently occupied, the AP is restricted to the primary 20 MHz channel, resulting in inefficient use of the spectrum and reduced transmission rates. With preamble puncturing, even if some secondary channels are occupied, the AP can transmit over the remaining idle channels, such as utilizing a 40 MHz secondary channel.

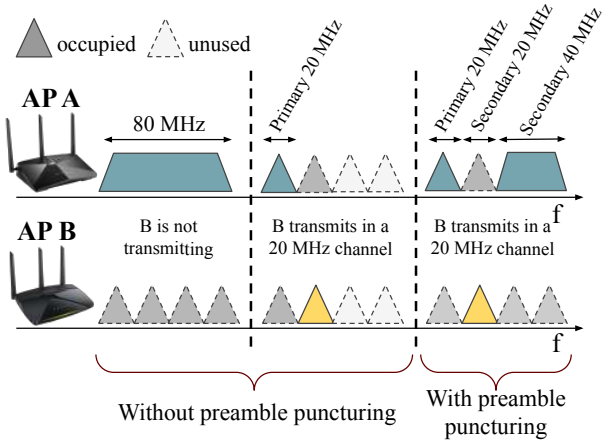

Fig. 14: Example of the benefits of preamble puncturing. Two neighboring APs share an 80 MHz channel where AP A uses the full bandwidth and AP B may transmit on a single 20 MHz subchannel. Without preamble puncturing (left), when AP B becomes active in one 20 MHz portion, AP A can only utilize its primary 20 MHz channel. With preamble puncturing (right), AP A can mute ("puncture") the overlapped 20 MHz subchannel and continue transmitting over the remaining secondary 40 MHz, improving channel utilization and spatial reuse.

While effective when multiple active STAs are present, the benefits of preamble puncturing diminish when the number of active STAs is low. To overcome this limitation, 802.11be introduces the ability to allocate multiple RUs to a single STA, as further discussed in Section V.

3) Non-Primary Channel Access (NPCA): Traditional 802.11 channel access mechanisms rely on a primary channel (PCH) for contention-based backoff and frame transmission. Even if secondary channels are available, transmissions can only occur when the primary channel is idle. This limitation reduces spectrum efficiency, particularly in dense deployments with OBSS, where the primary channel may often be occupied.

To address this challenge, 802.11bn is exploring NPCA, a mechanism that allows STAs to leverage idle secondary channels (SCHs), also referred to as NPCA primary channel, for frame transmission when the PCH is busy. When a Basic NAV is set on the PCH due to an ongoing OBSS PPDU transmission, the STA can switch to an announced NPCA primary

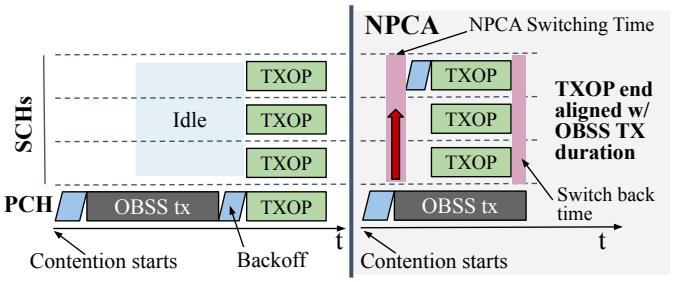

Fig. 15: NPCA enables the use of the secondary channels, reducing channel access delay and improving the overall system performance. Left: baseline operation where an overlapping BSS (OBSS) transmission on the primary channel (PCH) forces the device to remain idle on its secondary channels (SCHs) until it can win a TXOP. Right: NPCA, where the device switches to idle SCHs during the NPCA switching time, initiates a TXOP aligned with the OBSS transmission duration, and then switches back, thus exploiting otherwise unused spectrum.

channel to perform backoff and initiate frame transmission. This approach allows STAs to utilize available resources on SCHs, avoiding idle time and improving overall throughput, as exemplified in Fig. 15.

For NPCA to function effectively, the AP must announce a set of parameters to the participating STAs. These include the following:

- *NPCA Primary Channel Announcement:* The AP announces one NPCA primary channel for backoff, including a puncturing bitmap to disable specific subchannels when necessary.
- *Medium Synchronization:* Synchronization parameters such as energy detection (ED) thresholds and maximum TXOP durations are defined to ensure coordinated channel access among STAs.
- *Channel Switch Timing:* Switching between the PCH and SCH requires reconfiguring the device's RF front-end (e.g., phase-locked loop (PLL) retuning, filter adjustment, AGC stabilization), which introduces a hardware-dependent delay. This effect is particularly clear in mixed-capability deployments. For instance, an AP may support 160 MHz channels, while some associated STAs only 80 MHz. When NPCA triggers a switch to a channel segment outside the STA's native 80 MHz range, the STA must perform a full RF retuning to a new center frequency, incurring a longer analog reconfiguration delay than the AP. To guarantee that all STAs complete this transition before any of them transmits on the new channel, the AP announces a channel switching time large enough to accommodate the slowest device. Consequently, NPCA adopts the maximum channel switching time among the associated STAs to maintain proper synchronization and avoid collisions.

Note that NPCA may work at PPDU or TXOP level, while the second option relies on the initial received TXOP length, creating unexpected situations if the effective TXOP length changes. The first option of using one OBSS PPDU ensures predictable behavior at the expense of reducing transmission efficiency.

NPCA offers several benefits, especially in dense or congested environments. As shown in [98], it can substantially enhance spectral efficiency and throughput, while also reducing access delays under favorable conditions for supporting BSSs. Additionally, NPCA helps alleviate the OBSS performance anomaly, where low-rate transmissions from overlapping BSSs can degrade the performance of all nearby devices.

4) Dynamic Spectrum Operation (DSO): While APs commonly support wide bandwidths such as 160 MHz or 320 MHz, many STAs can only operate on narrower channels of 20, 40, or 80 MHz. This mismatch limits the AP's ability to fully utilize its wideband capabilities, since each transmission must accommodate the bandwidth constraints of the associated STA. Furthermore, all STAs are required to use the same 20 MHz primary channel defined by the AP.

For example, consider three STAs, labeled A, B, and C, associated with a 160 MHz-capable AP. Suppose STA A supports only 20 MHz, STA B supports 40 MHz, and STA C supports 80 MHz. Despite the AP's 160 MHz capability, in SU mode, it can only use its primary 80 MHz channel—and only when transmitting to STA C. When transmitting to STA A or STA B, the AP must restrict the transmission to 20 MHz and 40 MHz, respectively. The use of OFDMA can increase bandwidth utilization efficiency by allowing simultaneous transmissions to all three STAs within the primary 80 MHz channel (see Section V-A). In such a case, the AP could allocate 20 MHz to STA A (its primary channel), 20 MHz to STA B (its secondary 20 MHz), and 40 MHz to STA C (its secondary 40 MHz). However, even with OFDMA, 50% of the AP's total bandwidth remains unused, and the wideband capabilities of STA B and STA C remain underutilized.

DSO offers a compelling solution to address this limitation. The concept is simple: remove the constraint of using the 20 MHz primary channel for data transmissions. By lifting this requirement, an STA can be allocated to any portion of the AP's supported bandwidth, enabling more flexible and efficient use of the channel. To implement DSO, two mechanisms are necessary. First, the AP and each STA must exchange their capabilities information and agree (typically during the association phase) on one or several DSO subchannels, all with maximum width equal to the STA bandwidth, to be used for future transmissions. Second, a signaling mechanism must be established to inform an STA that it will be allowed to receive or transmit data on a specific DSO subchannel (out of the pre-agreed ones) during a given TXOP. This can be achieved by extending existing Buffer Status Report Poll (BSRP) trigger frames to carry subchannel allocation information. Moreover, the network has to account for the required switching times to move to the assigned subchannel and back to the primary 20 MHz channel at the end of the transmission, so that default BSS operation can resume seamlessly.

The example involving STAs A, B, and C is illustrated in Fig. 16. If the AP transmits using SU transmissions, it must perform three separate transmissions, one for each STA, each limited to the maximum bandwidth supported by the respective STA. With OFDMA, a single transmission is possible, but each STA receives only a portion of the available bandwidth, resulting in longer transmission durations. In contrast, by combining DSO with OFDMA, the AP can fully exploit its

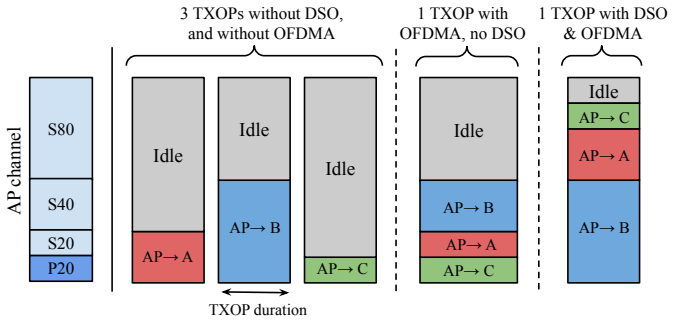

Fig. 16: Benefits of DSO: Combined with OFDMA, it improves spectrum utilization and reduces the TXOP duration. Note that the figure illustrates only the data exchange. In the case of OFDMA without DSO, all stations' channel allocations overlap with the P20 primary channel. With DSO, however, stations can be allocated across the AP's full 160 MHz channel, removing the primary-channel overlap constraint and enabling them to use their full channel width for data exchange.

wideband capability while simultaneously serving all STAs using their maximum supported bandwidths, minimizing the portion of the spectrum that remains unused.

F. Coexistence and Spatial Reuse

802.11 devices belonging to independent BSSs need to coexist when operating on the same frequency channels. Ensuring smooth coexistence among BSSs is crucial to provide an acceptable performance to users, especially in dense deployments. 802.11 includes some coexistence mechanisms that aim to increase the efficiency and improve the performance of OBSS deployments as detailed next.

1) BSS Coloring: This is a feature that was first introduced in 802.11ah for energy-saving purposes and then evolved in 802.11ax to improve efficiency in dense deployments. BSS coloring allows APs and STAs to quickly differentiate between *intra-BSS* (from their own BSS) and *inter-BSS* (from neighboring BSSs) transmissions. To do so, each BSS is assigned a unique 6-bit identifier (*color*) which is informed in the PHY header's signal preamble, part of the HE-SIG-A in the PHY header. The AP is in charge of handling the BSS coloring (e.g., a color change is announced using BSS Color Change Announcement messages), but STAs can also participate by sending *BSS color collision reports* to the AP when they notice that another BSS is using the same color. Fig. 17 illustrates the potential of BSS coloring to enable simultaneous transmissions from multiple BSSs. As described next, BSS coloring is the foundation for specific mechanisms like OBSS packet detect (PD)-based SR (described below), which aim to make a more efficient use of the medium. In addition, by allowing the quick identification of the source of a transmission (read in the PHY headers rather than in the MAC ones), BSS coloring is an appealing feature for energy saving.

2) Basic vs. Intra-BSS NAV: When a device detects the start of a PPDU, it might activate a basic NAV upon certain conditions being met (see Section IV-A). Before 802.11ax, even if their signals were weak enough not to cause harmful interference, inter-BSS transmissions could provoke the activation of the NAV in other devices. Such a conservative approach is detrimental to efficiency, and that is why 802.11ax

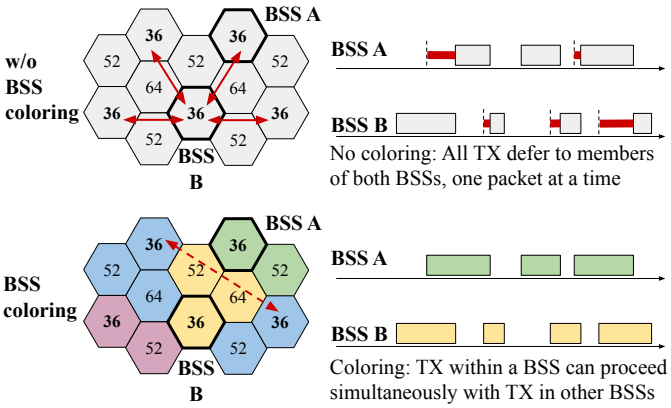

Fig. 17: BSS coloring improves spectrum utilization by enabling simultaneous transmissions in BSSs with different colors. Top: case without coloring, where all transmissions in overlapping BSSs A and B defer to each other and occur one at a time. Bottom: color-marked cells and timelines, where transmissions within each BSS can proceed concurrently with those in neighboring BSSs that use a different color.

introduced a differentiation between intra-BSS and inter-BSS NAVs (or basic NAV). Thanks to BSS coloring, a device can differentiate between intra-BSS and inter-BSS transmissions and, as a result, maintain two different NAVs. For a device to mark the channel as idle, thus being able to initiate a transmission, both the basic and intra-BSS NAVs must be inactive.

3) Spatial Reuse: Given the importance of efficiently using the spectrum in dense scenarios, 802.11ax introduced a mechanism called OBSS PD-based SR [99]. The procedure is depicted in Fig. 18a and works as follows:

1) A device (AP B) detects a PPDU transmission (from AP A).
2) It determines the origin of the transmission by reading the BSS color from the packet headers.
3) Based on the color from AP A (BSS Color 1), which is different from its own color (BSS Color 2), AP B determines it is an inter-BSS transmission.
4) An OBSS PD value is applied to check whether the detected transmission can be ignored or not. The OBSS PD threshold is a value between -82 dBm and -62 dBm that is defined for 20 MHz transmissions and increased by 3 dB as the channel width is doubled.
5) The transmission can be ignored, so the PHY.CCARESET primitive is activated, which resets the PHY CCA.
6) After exhausting the backoff, AP B initiates a transmission. In exchange for relaxing the CCA requirements and using a higher OBSS PD value, a transmit power limitation is applied during the detected SR TXOP.

In OBSS PD-based SR, the maximum transmit power allowed by the OBSS PD-based SR operation depends on the selected OBSS PD, and is defined as

$$\text{TX\_PWR}_{\max} = \text{TX\_PWR}_{\text{ref}} - (\text{OBSS PD} - \text{OBSS PD}_{\min}), \quad (3)$$

where $\text{OBSS PD}_{\min} = -62$ dBm and $\text{TX\_PWR}_{\text{ref}}$ is a reference power (e.g., 21 dBm or 25 dBm, depending on the device's capabilities).

Later, 802.11be extended the SR operation with a second mechanism, parametrized spatial reuse (PSR) [100], which had already been discussed during the standardization of 802.11ax. The mechanism is illustrated in Fig. 18b. It is conceived as an UL variant of SR, in which an AP shares with devices from other BSSs the TXOP where it will receive UL triggered transmissions from its associated STAs. In particular:

1) An AP (AP A) wins a TXOP.
2) Thanks to PSR, it can share the TXOP with other devices from different BSSs. It does so by sending a Trigger frame with PSR information (including the maximum transmit power that can be used during the TXOP).
3) The Trigger frame is detected by other PSR devices (by AP B in Fig. 18b) as a parameterized spatial reuse reception (PSRR) PPDU
4) AP B identifies a PSR opportunity, which is similar to the SR opportunity defined above, and initiates a transmission.

In PSR, a transmit power limitation is imposed to ensure that inter-BSS interference does not disrupt the original PPDU transmission. Specifically, the STA that identifies the PSR opportunity (AP B in Fig. 18b) must apply the following transmit power limitation when transmitting a parameterized spatial reuse transmission (PSRT):

$$\text{TX\_PWR}_{\text{PSR}} - 10 \times \log_{10} N^{20\text{MHz}} \leq \text{PSR}_{\text{min}} - \text{RPL}_{\text{PSRR}}^{20\text{MHz}}, \quad (4)$$

where $\text{TX\_PWR}_{\text{PSR}}$ is the transmit power during a PSR opportunity, $N^{20\text{MHz}}$ is the number of 20 MHz channels used, and $\text{RPL}_{\text{PSRR}}^{20\text{MHz}}$ is the normalized received signal power on at least one of the involved 20 MHz channels in which the PSRR and PSRT are transmitted. As for PSR_{min}, it is the minimum value that the AP holding the TXOP (AP A in Fig. 18b) indicates (using the "SR" field of a PSRR PPDU) to ensure that other inter-BSS transmissions do not affect its own transmission. Its value is computed as

$$\text{PSR\_INPUT} = \text{TX\_PWR}^{(\text{AP})} + \gamma^{(\text{AP})}, \quad (5)$$

where $\gamma^{(\text{AP})}$ defines the acceptable receiver interference level at the AP, in dBm. It is recommended to set $\gamma^{(\text{AP})}$ as the expected signal power for the highest MCS minus the minimum SNR value that yields at most a 10% packet error rate (PER), minus an additional safety margin of up to 5 dB.

V. Wi-Fi Multi-user Technologies

The latest 802.11 amendments have been introducing transformative multi-user technologies designed to enhance network performance, particularly in dense deployments with numerous connected STAs. These technologies improve spectral efficiency, reduce latency, and ensure fair resource allocation among STAs. The two cornerstone innovations in this realm are OFDMA and MU-MIMO, which exploit the frequency and spatial dimensions, respectively, to enable simultaneous transmissions to multiple STAs. These technologies are closely integrated in 802.11, with MU-MIMO transmissions coordinated within the OFDMA framework. Specifically, multiple data streams can be spatially multiplexed within each frequency sub-channel (OFDMA RU). This allows assigning the

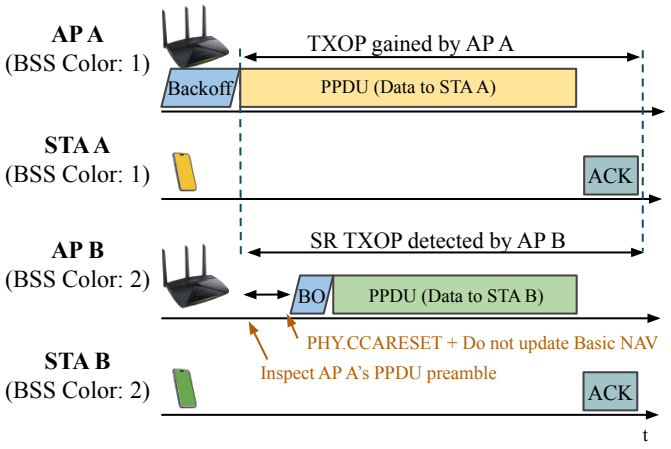

(a) OBSS PD-based SR.

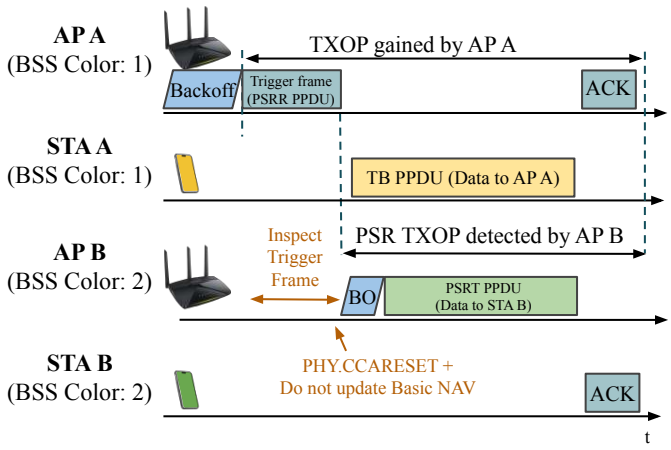

(b) PSR.

Fig. 18: Spatial reuse operation. (a) OBSS PD-based SR: AP B detects an OBSS PD-based SR opportunity upon inspecting AP A's transmission and initiates a subsequent transmission to STA B. (b) PSR: AP A wins the TXOP and sends a trigger frame to allow AP B to transmit during the same TXOP, leading to simultaneous transmissions in both BSSs.

same RU to multiple STAs which then use MU-MIMO for simultaneous transmission. The allocation is announced in the EHT/UHR-SIG of the PPDU preamble (see Section III-F2) for downlink multi-user transmissions and in the Basic Trigger frame for uplink multi-user transmissions (detailed in Section V-A3). In this way, the AP can dynamically allocate radio resources to STAs based on their traffic requirements and channel conditions, ensuring efficient spectrum utilization and improved network performance.

A. Orthogonal Frequency Division Multiple Access (OFDMA)

OFDMA, introduced in 802.11ax, enables efficient utilization of the frequency spectrum by dividing the RF channel where the network is operating into smaller sub-channels called RUs, each of which combines multiple adjacent OFDM subcarriers [101]. The RUs can be independently assigned to different STAs, which tune their radios to the allocated RUs. This enables the AP to transmit data to multiple STAs simultaneously through different RUs within a single transmission opportunity, based on the STAs communication needs (*downlink OFDMA*). Moreover, *uplink OFDMA* enables multiple STAs

to transmit data simultaneously to the AP over their assigned RUs [102]. The STA-RU association is decided by the AP, which triggers this transmission mode.

Downlink and uplink OFDMA allow increasing the efficiency in the spectrum utilization. Indeed, as introduced in Section IV-C, the EDCA channel access protocol becomes increasingly inefficient as the number of STAs grows, leading to high delays in obtaining a TXOP and, in turn, reduced network performance. Uplink OFDMA addresses this challenge by aggregating multiple transmissions within a single TXOP, preventing the sharp decline in capacity observed with CSMA/CA in dense environments. In [26], the authors show that the improvement of using OFDMA is clearly visible starting from a network with four STAs. Moreover, the high flexibility in the RUs assignment for downlink and uplink OFDMA enables a more balanced distribution of RUs of different bandwidths to different STAs, which can be done based on their specific traffic requirements. This allows the network to address varying traffic demands efficiently. For instance, STAs transmitting a high amount of data, such as video, are assigned larger RUs to accommodate their higher bandwidth needs. On the other hand, IoT sensors sending small, infrequent packets are allocated smaller RUs to save spectrum resources and maintain low power consumption, thus extending battery life. This way, OFDMA facilitates smaller and more frequent transmissions, which lowers latency and jitter, being particularly valuable for time-sensitive applications like voice-over-Wi-Fi and video streaming. In the following, we describe the RUs allocation process and the way OFDMA is implemented in the downlink and uplink directions.

1) Resource Unit (RU) Allocation: Proper RU allocation is crucial to enable full exploitation of the flexibility and benefits offered by OFDMA transmissions [103]. The smallest RU in OFDMA consists of 26 adjacent subcarriers (tones), which allows up to 9 STAs to share a 20 MHz channel in the downlink or uplink direction. Note that the central 26-tone RU is obtained by combining the 13 subcarriers before and after the DC subcarriers. Indeed, it is important to note that not all subcarriers in an RU can be used for data transmission. As detailed in Section III-C, some subcarriers are reserved for guard bands to minimize interference, while others function as DC or pilot tones to assist with signal synchronization and demodulation, ensuring reliable communication. Moreover, some null subcarriers are located near the DC or guard tones to reduce transmit center frequency leakage, receiver DC offset, and interference from neighboring RUs. Adjacent 26-tone RUs can be grouped to form 52-tone, 106-tone, and 242-tone RUs on 20 MHz channels. Larger RUs are available increasing the bandwidth and grouping more contiguous 26-tone RUs. Specifically, 40 MHz channels supports RUs with up to 484 tones while one and two 996-tone RUs can be formed in 80 MHz and 160 MHz channels [66]. This hierarchical structure of RU sizes enables precise resource allocation tailored to the traffic characteristics of each STA.

2) Downlink OFDMA: In the downlink case, the AP simultaneously transmits data to multiple STAs using a single transmission. This bundling mechanism minimizes contention time and significantly improves spectrum efficiency. Downlink OFDMA transmissions are set as follows:

- The AP transmits an MU-RTS trigger frame to the STAs that are expected to receive downlink traffic. The participating STAs simultaneously respond with a CTS frame.
- The AP bundles frames intended for each STA and modulates them on the STA-assigned RUs.
- An MU-PPDU (see Section III-F1) is formed by using a single preamble spanning all frequencies (full-band transmission), followed by individually addressed frames within each RU. The STA–RU allocation is detailed in the SIG field of the preamble. To ensure uniform transmission length, padding is used on shorter frames to match the length of the longest frame. Moreover, smaller RUs can be allocated to shorter frames to reduce padding overhead.
- An MU-Block Acknowledgment Request (BAR) trigger frame may be transmitted by the AP at the end of the multi-user transmission to request acknowledgment from multiple STAs simultaneously, reducing overhead.
- A BA is simultaneously transmitted by the STAs to the AP providing information about the correctness of the received data. The BA can also be untriggered when the UL parameters are included in the data transmission.

Note that the actual number of supported STAs can be further increased by integrating spatial multiplexing through MU-MIMO. In this case, sounding should be performed before actual data transmission, as discussed in Section V-B.

3) Uplink OFDMA: While downlink OFDMA is relatively straightforward due to centralized control by the AP, uplink OFDMA requires more complex synchronization mechanisms. Indeed, precise synchronization is required to ensure that all STAs' preamble symbols arrive at the AP simultaneously and, in turn, can be properly decoded. This synchronization is facilitated through careful scheduling performed by the AP, which specifies the exact timing, RU allocation, and transmission parameters for each STA. An uplink OFDMA transmission proceeds as follows:

- The AP transmits an MU-RTS trigger frame to the STAs that are expected to transmit uplink traffic. The STAs simultaneously respond with a CTS frame.
- The AP transmits a BSRP Trigger frame in broadcast to solicit the STAs indicated in the Info List field of the Trigger frame to report the amount of buffered uplink data inside a dedicated Buffer Status Report frame, a SIFS after the trigger. This information helps the AP to optimize RU scheduling.
- Then, the AP transmits a Basic Trigger frame to trigger the coordinated uplink transmission, indicating the STA-RU association inside the User Info List field of the trigger frame. This frame provides coordination for: (i) *time synchronization*, ensuring all STAs start their transmissions simultaneously to avoid overlapping symbols (timing tolerance is 0.4 μs) [66], (ii) *frequency alignment*, maintaining consistent frequency offsets across STAs to prevent interference between adjacent RUs, and (iii) *amplitude calibration*, adjusting signal strength to ensure

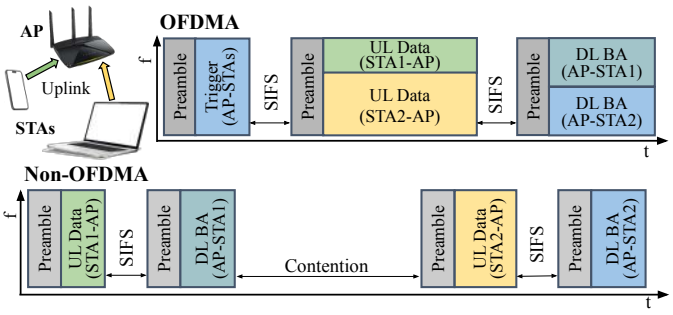

Fig. 19: Example of uplink OFDMA (top) and non-OFDMA (bottom). In the OFDMA case, the AP sends a trigger frame and then receives simultaneous uplink data from two STAs on different frequency resource units, followed by simultaneous downlink block acknowledgments (DL BAs) after SIFS intervals. In the non-OFDMA case, each STA contends independently and transmits its uplink data and receives its DL BA in separate transmission opportunities, incurring additional contention overhead.

uniform reception across all RUs.

- All STAs indexed in the User Info List field of the Trigger frame, transmit their data in the uplink direction simultaneously after an SIFS from the reception of the trigger using a TB PPDU (see Section III-F1).
- The AP uses a Multi-STA BA to inform the different STAs about which frames have been received correctly.

The upper part of Fig. 19 summarizes the last part of this packet exchange considering two STAs that simultaneously transmit data to the AP using UL-OFDMA. The bottom part of the figure shows how the same two-STA uplink transmission would be carried out using CSMA/CA where data streams from the two STAs are transmitted sequentially after independent channel contention. In OFDMA, the AP manages the transmissions in a deterministic way, significantly reducing contention overhead and consequently reducing communication latency for scheduled transmissions. Moreover, as discussed at the beginning of this section, OFDMA enables a more efficient utilization of the spectrum in case of unequal transmission requirements: the AP can assign small RUs to STAs having a small amount of data to transmit (STA 1 in Fig. 19), while reserving more space for STAs with a higher amount of data to be transmitted (STA 2 in Fig. 19).

With 802.11ax, another strategy to identify the STAs with packets to transmit has been introduced [66]. This procedure substitutes the MU-RTS and CTS frame exchange and entails the AP to broadcast an NDP Feedback Report Poll, which includes information about the STAs that are allowed to respond and their assigned RUs. Hence, STAs with uplink traffic respond to the AP with an NDP Feedback Report and wait for the BSRP.

4) Multi-RU Support: One key limitation of OFDMA in 802.11ax is that each STA can only be assigned a single RU. This constraint can lead to underutilization of available bandwidth, particularly in scenarios with few STAs and when preamble puncturing is employed to enable the simultaneous use of non-contiguous portions of the radio spectrum as detailed in Section IV-E. Considering the preamble puncturing example in the rightmost part of Fig. 14, if the AP has data for only one STA, it can only assign the primary 20 MHz channel to that STA, leaving the remaining 40 MHz bandwidth (the majority of the spectrum) unused. While this restriction in 802.11ax helps avoid potential regulatory issues, such as protecting incumbents in the 6 GHz band, it significantly reduces transmission efficiency and throughput in scenarios with low STA density.

The 802.11be standard amendment addresses this limitation by introducing multiple resource unit (MRU) support, which allows the AP to assign multiple RUs to a single STA, thus significantly improving spectrum utilization. MRU support also enhances performance when leveraging it for diversity gain: by transmitting the same data across multiple RUs, the AP can increase the likelihood of successful packet delivery in environments with high interference or variable channel conditions.

RUs are grouped into two categories based on their size: *small-size RUs*, including 26-tone, 52-tone, and 106-tone RUs, and *large-size RUs*, comprising 242-tone or larger RUs. A small-size RU can only be combined with other contiguous small-size RUs to form a small-size MRU, while large-size MRUs can be created by combining different large-size RUs, which may be non-contiguous. Specific combinations of RUs are predefined to ensure compatibility and minimize overhead. Small-size MRUs include 52+26-tone and 106+26-tone configurations, obtained by combining a 52-tone or 106-tone RU with an adjacent 26-tone RU within the same 20 MHz channel, as depicted in Fig. 20a. Large-size MRUs include 484+242-tone, 996+484-tone, 2×996+484-tone, 3×996-tone, and 3×996+484-tone combinations, formed by grouping RUs within the same 80 MHz frequency channel. The RUs that can be combined to create small or large-size MRUs are defined in the 802.11be amendment based on configurations that offer the highest performance gains. Large-size MRUs can also be used for non-OFDMA transmissions to multiple STAs, in which case an additional RU combination, i.e., 996+484+242-tone MRU, may be employed. Examples of 484+242-tone MRUs are shown in Fig. 20b.

Considering the same preamble puncturing example used above, rather than restricting an STA to the primary 20 MHz channel, the AP can employ MRU and allocate to the STA a secondary 40 MHz channel, thus creating a 484+242-tone large size MRU and effectively tripling spectrum utilization.

5) Distributed RU in 802.11bn: OFDMA implementations up to 802.11be use RUs composed of continuous subcarriers. This constraint is expected to be relaxed in the 802.11bn amendment, with the introduction of DRUs. A DRU consists of non-adjacent tones which are evenly distributed along the operational bandwidth (excluding the central DC and null subcarriers). DRU-based OFDMA will addresses the limitations that continuous subcarrier placement has in shared spectrum scenarios, such as in the 6 GHz band, where devices should operate either in Standard Power or Low Power Indoor modes (see Section II-A2). Specifically, the Standard Power mode uses a central frequency coordinator that decides on which channel and at what power the AP should operate. This requires precise localization of the AP, which may be difficult to achieve. On the other hand, the Low Power Indoor mode sets a maximum on the EIRP per MHz, e.g., -1 dBm/MHz.

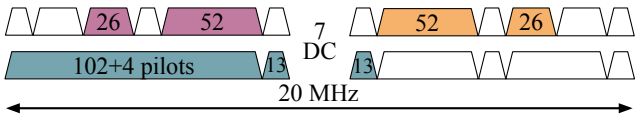

(a) Small-size MRUs: 52+26-tone and 106+26-tone.

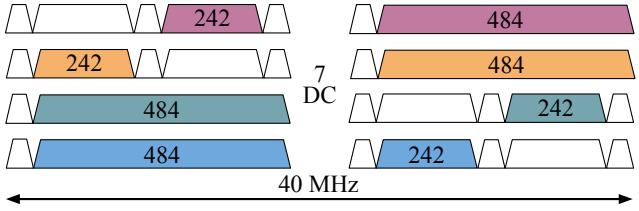

(b) Large-size MRUs: 484+242-tone.

Fig. 20: Examples of different MRU combinations. Top (a): small-size MRUs within a 20 MHz channel, combining 26- and 52-tone (or 106- and 26-tone) resource units around the DC subcarrier and pilot tones. Bottom (b): large-size MRUs over 40 MHz with 484- and 242-tone RUs in various placements across the channel bandwidth.

Considering the 78.125 kHz sub-spacing used since 802.11ax, the cumulative transmit power used on the 12.8 subcarriers hosted on each MHz should meet the EIRP limit. This reduces power efficiency and leads to degradation of the OFDMA transmission. DRUs are expected to overcome this limitation. The main idea is to reduce the number of tones per MHz, concentrating power over fewer tones and thereby increasing the effective transmission power per tone while remaining within regulatory constraints [66]. By evenly distributing tones along the operational bandwidth, DRUs will support a comparable number of units per channel bandwidth as standard OFDMA without degrading spectral efficiency, thus ensuring that they can handle diverse traffic demands effectively. The regular distribution of tones improves channel smoothing, which enhances the performance of DRU-based transmissions through a better estimation of the wireless channel. DRUs also tend to be more robust in environments with high channel fading. The hierarchical tone mapping structure and pilot tone placement used in standard OFDMA is preserved, ensuring compatibility and ease of implementation. The combination of these features could allow DRUs to improve power and spectral efficiency without complicating implementation or sacrificing performance compared to standard OFDMA operations. Note that DRUs are currently defined for use in uplink OFDMA only. The trigger frame used to solicit multi-user uplink transmissions has been updated to include signaling for DRU.

B. Multi-user MIMO

As Wi-Fi has evolved, APs have incorporated more antennas and advanced spatial multiplexing techniques to support the increasing number of connected STAs. MU-MIMO builds on these advancements, allowing the transmission of data streams to multiple STAs simultaneously. This enhances spectral efficiency, reduces contention for accessing the channel, and improves overall network performance [76]. In the following, we review the use and implementation of MU-MIMO in the UL and DL directions.

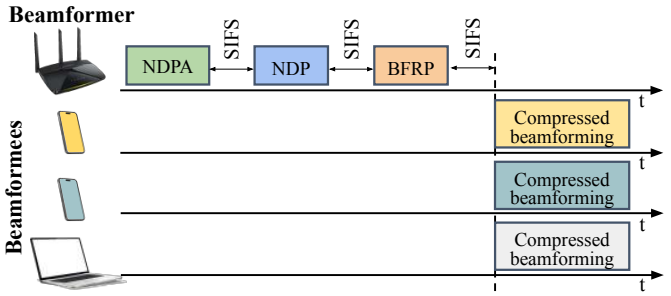

Fig. 21: Example of a TB channel sounding round with three STAs. The beamformer (AP) first transmits an NDPA frame, followed after SIFS by the corresponding NDP and a beamforming report poll (BFRP). In response, the three beamformees send their compressed beamforming feedback frames, so that the AP acquires channel state information from all STAs within a single sounding exchange.

1) Downlink MU-MIMO and MU-MIMO-OFDMA: The 802.11n amendment only supported SU-MIMO, where the AP transmits multiple streams to a single multi-antenna STA as detailed in Section III-D. Its successor, 802.11ac, integrated DL-MU-MIMO in the standard, allowing leveraging at full the benefits of multi-antenna systems. Indeed, SU-MIMO only allows multiplexing a few streams, as most user devices (STAs) have limited antenna arrays (one or two antennas for smartphones and laptops). In MU-MIMO, the AP leverages the lack of correlation among the channels of the different STAs to serve them simultaneously with multiple data streams. Specifically, the rank of the combined channel matrix, indicating the number of simultaneous data streams that can be transmitted, is equal to the minimum between the number of STAs' antennas and the antennas available at the AP. Moreover, as introduced in Section III-D for SU-MIMO, 802.11 regulates the maximum number of simultaneous transmissions. The 802.11ac amendment dictates that APs can transmit up to 8 spatial streams to 4 STAs simultaneously. In 802.11ax, the number of STAs that can be simultaneously served increased to 8 while the total number of streams remained 8. The same limits apply to 802.11be and are expected to remain unchanged in 802.11bn.

Downlink MU-MIMO follows a similar procedure as the one described in Section III-D for SU-MIMO. However, in this case, the AP should collect the channel information from all connected STAs to properly precode the streams and reduce their cross-interference, thus enabling accurate decoding. In turn, the channel sounding procedure is carried out for all the STAs involved in the MU-MIMO transmission as detailed in the sequel. Note that MU-MIMO can be integrated with OFDMA, where multiple STAs are served through MU-MIMO over the same RU. In this case, channel state information is reported only for the intended portion of the bandwidth following the indication in the NDPA frame.

2) Multiuser Channel Sounding: To obtain channel feedback for MU-MIMO transmissions, the MU-MIMO beamformer (AP) shall not use the non-TB sounding procedure described in Section III-D. MU-MIMO sounding is initiated by the AP through the transmission of a single *broadcast* NDP. In this case, the NDPA that comes before the NDP includes a list of the association identifiers (AIDs) corresponding to the involved MU-MIMO STAs. The AID is an integer value in the

range $[1, 2007]$ that is assigned to each STA by the AP as part of the association process upon reception of the Association Request frame from the STA. The AID is selected among the available AID pool—with the constraint that the AID should be unique within a BSS—and transmitted back to the STA in the Association Response frame. Moreover, in case of MU-MIMO-OFDMA transmissions, the NDPA indicates the RUs or MRUs for which channel estimate should be fed back using the Partial BW Info subfield. This information is also included in the BW and Puncturing Channel Information fields, which are part of the U-SIG field of the sounding NDP.

Upon NDP reception, each STA estimates the CFR and compresses it following the procedure described in Section III-D. In the MU-MIMO case, the beamforming feedback report also includes a MU-MIMO Exclusive Beamforming Report field which carries delta SNR values indicating the difference between the subcarrier and the average SNRs. This information can be used by the beamformer to compute the steering matrix.

At this point, the STAs have to feed back their compressed channel estimate. Pre-802.11ax APs perform this collection of beamforming feedback frames from different STAs sequentially: the STA associated with the first AID in the AID list transmits back its feedback after a SIFS from the NDP reception; hence, all the other STAs are triggered to feed back their information by the AP through the transmission of Beamforming Report Poll frames indexing each of them. The 802.11ax amendment introduced a new TB sounding procedure which increase spectrum efficiency [72] and is illustrated in Fig. 21. Instead of transmitting the different feedback frames separately, they are simultaneously transmitted to the AP using an uplink multi-user strategy specified in a new frame named Beamforming Feedback Report Poll (BFRP) which is transmitted an SIFS after the NDP. If the number of STAs that should report the feedback is higher than the maximum allowed by the multi-user transmission, multiple rounds of BFRP and feedback transmissions are allowed. This TB sounding may also be used for SU-MIMO feedback if the beamformee supports it.

Upon reception of each compressed beamforming feedback frame, the AP retrieves the angles and uses them to reconstruct the $\mathbf{V}_{i,k}$ matrices for all the subcarriers k for each STA i (as detailed in Section III-D for SU-MIMO). These matrices—which are a proxy for the CFR—are used by the AP to decide which STAs can be scheduled together in the MU-MIMO transmission to maximize the QoS and spectrum efficiency. The idea is to avoid combining in the same MU-MIMO group STAs that will create interference to each other. A possible strategy consists of computing the inter-STA correlation and selecting the STAs for which this is minimum [104]. However, STA grouping is an implementation-dependent feature and, in turn, each vendor is free to decide the best strategy to adopt. Having obtained the set of STAs in the same group, the $\mathbf{V}_{i,k}$ matrices for subcarrier k and all the STAs are combined into a single $\mathbf{V}_k$ matrix over the spatial stream (column) dimension thus obtaining a $N_{\mathrm{TX}} \times N_{\mathrm{SS}} \times N_{\mathrm{SD}} + N_{\mathrm{SP}}$ matrix where $N_{SS} = \sum_i N_{\mathrm{SS},i}$ and $N_{\mathrm{SS},i}$ is the number of spatial streams that should be directed to STA i. This matrix is used to design the precoding for subcarrier k. For example, using zero-forcing the beamforming steering matrix is $\mathbf{Q}_k = \tilde{\mathbf{V}}_{\mathbf{k}} \left(\tilde{\mathbf{V}}_{\mathbf{k}}^{\dagger} \tilde{\mathbf{V}}_{\mathbf{k}.} \right)^{-1}$. In the ideal case, this makes the streams directed to one STA orthogonal to the channel between the AP and each of the other STAs, thus reducing interference.

Note that as the feedback is transmitted back unencrypted, security issues may arise. Indeed, a malicious user in the network may capture the legitimate feedback of other stations and properly craft a malicious feedback that leads to a wrong precoding which drastically degrades the network performance [77], [104], [105], or enable eavesdropping on a victim's traffic [106]–[109].

3) Uplink MU-MIMO and MU-MIMO-OFDMA: In 802.11ax, MU-MIMO was extended to support both downlink and uplink transmissions. Uplink MU-MIMO, allows multiple STAs to transmit data simultaneously to the AP. The process is orchestrated by the AP, which sends a Basic Trigger frame specifying the RU and timing for each STA. The STAs then transmit data in their assigned RU, and the AP decodes the individual data streams. As discussed for uplink OFDMA, uplink MU-MIMO reduces contention, as multiple STAs can transmit concurrently, and is particularly beneficial in environments with a high density of connected STAs [110]. To manage the added complexity introduced by uplink MU-MIMO transmissions, 802.11ax introduced synchronization mechanisms for precise alignment of time, frequency, and power across STAs to avoid interference and ensure efficient decoding at the AP [15]. The alignment in time is achieved by requiring the STAs to transmit their UL data a SIFS after the Basic Trigger frame transmitted by the AP. The AP also indicates in the Basic Trigger frame the power that each STA should use for transmission to ensure that the different UL transmissions are received at almost the same power at the AP. A pre-compensation of the CFO is also performed using the training fields in the Basic Trigger Frame [111].

VI. Wi-Fi Energy Savings

Power consumption in wireless networks has become a growing concern, particularly as energy efficiency becomes a priority for both STAs and APs. An 802.11 card can account for anywhere between 3% and 10% of a smartphone's total power consumption [112], which underscores the need for power-saving mechanisms. This section explores the energy-saving solutions adopted (power save mode (PSM) and TWT) and those under consideration for 802.11bn (DPS).

A. Power Save Mode (PSM)

PSM was introduced in the initial release of 802.11 to alleviate the power consumption of battery-powered 802.11 devices. It introduced two power states for STAs, *awake* (active) and *doze* (power save mode), depending on their expected capability to transmit or receive data. The main procedures held within PSM are as follows.

- When an STA enters doze mode, it notifies the AP by setting the Power Management bit to 1 in the Frame Control field of any frame. This allows the AP to buffer

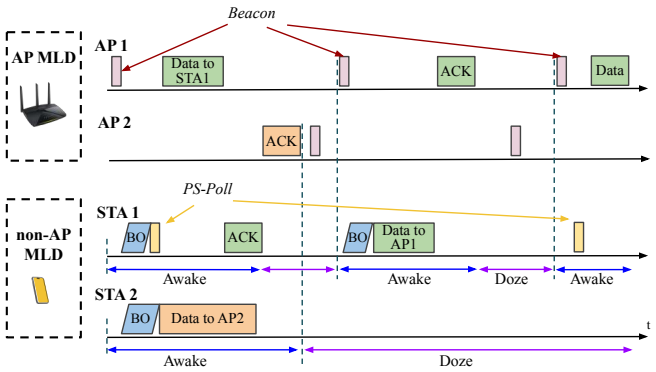

Fig. 22: Example of the multi-link power-save operation. The non-AP MLD activates only one link (STA 1) to listen for Beacon frames from the AP MLD. This mechanism allows the non-AP MLD to activate the other links (STA 2) only when required for data transmission, thereby saving energy.

the frames intended for such an STA and keep them until it wakes up again.

- Dozed STAs wake up to receive the Beacon frames from their AP, in order to check if there are any bufferable units (BUs) available for them. The existence of buffered traffic for a particular STA is informed by the corresponding AID in the traffic indication map (TIM) bitmap (1 indicates the existence of unicast data).
- When an STA wakes up, it notifies the AP with a PS-Poll or unscheduled automatic power save delivery (U-APSD) trigger frame. This action makes the AP aware that the STA is ready to receive the buffered data.
- Finally, the AP starts sending the buffered frames, which can be done in multiple transmissions if the More Data bit in the Frame Control is set to 1.

While effective at the time it was introduced, PSM forced STAs to wake up frequently—often several times per second—to check the Beacon and send a Trigger frame for data retrieval, even if they did not have any data to transmit and receive. This approach resulted in significant power consumption for low-duty-cycle STAs. Later, with the introduction of voice-over-Wi-Fi in 802.11e, a more advanced mechanism, U-APSD, was introduced. U-APSD allowed STAs to sleep during beacon intervals and wake up only when they needed to transmit or receive data, reducing unnecessary wake-ups. In addition, 802.11v introduced a new sleep mode, the wireless network management (WNM) sleep mode, whereby STAs do not need to listen to every delivery traffic indication map (DTIM) Beacon. However, PSM was still limited to relatively short sleep intervals and did not provide the long-term scheduling flexibility required by many modern IoT devices [113].

Today, PSM has turned into multi-link power-save operation in 802.11be, which extends the feature to multiple links, thus unleashing a high flexibility to perform more efficient energy saving operations thanks to dynamic link switching. As shown in Fig. 22, one non-AP multi-link device (MLD) (described later in Section VII) can wake up to listen to Beacons from an AP MLD on one link only, thus allowing the other link(s) to remain in doze mode when there is no need to use them.

B. Target Wake Time (TWT)

To address the limitations of earlier power saving mechanisms like PSM, TWT was introduced in 802.11ah (commercially known as Wi-Fi HaLow), a specialized amendment designed for low-power, long-range IoT networks in the sub-1 GHz license-exempt bands, but its adoption was limited due to the niche application of 802.11ah devices. While TWT was explicitly designed for power efficiency and extensive range requirements (at the cost of high speed), its wider adoption was initially limited due to the niche nature of 802.11ah devices.

To meet both power efficiency and high speed performance, the 802.11ax amendment reintroduced and expanded TWT, making it a widely applicable power-saving feature for modern devices. TWT is especially valuable for battery-operated STAs that communicate infrequently, such as sensors and other IoT endpoints. TWT overcomes some of the limitations of PSM and, instead of relying on frequent wake-ups, allows AP and STAs to negotiate a wake-up schedule tailored to the specific needs of the devices. Therefore, TWT moves away from statically defined sleeping periods and introduces a much more flexible approach, where the AP and the STAs define sleep schedules jointly. The core of TWT lies in the negotiation and agreement of SPs between the AP and the STAs. The SPs define the periods when an STA is active for communication. In what follows, we explain how TWT agreements are established and how SPs are implemented [114].

1) TWT Agreement: To initiate the negotiation of a TWT agreement, an STA sends a TWT Request frame to the AP. A request can include parameters such as Target Wake Time (offset time to start the first SP), TWT Wake Duration (the awake time within a SP), TWT Wake Interval (the interval between consecutive SPs), Minimum TWT Wake Duration (the minimum period a device stays awake to exchange frames), TWT Channel (the primary channel on which communication will occur during the TWT session), TWT Protection (indication to protect TWT sessions through mechanisms like RTS/CTS), and Negotiation Type (indicating soft or strict requirements). Then, the AP responds with a TWT Reply, which can simply accept the request or renegotiate it.

TWT agreements can be *individual* or *broadcast* (collective). In individual TWT, an STA negotiates a wake-up schedule that matches its traffic requirements, minimizing unnecessary wake-ups (Fig. 23a). In broadcast TWT, the AP establishes a shared wake-up schedule for a group of STAs, optimizing spectrum use and reducing contention (Fig. 23b). Broadcast TWT agreements are announced in Beacon frames, and STAs wake up only for relevant Beacons and follow the broadcast schedule. STAs can request membership in an existing broadcast agreement or propose new ones to the AP. Each STA can maintain up to eight TWT agreements with an AP, identified by a 3-bit value. This capability allows devices to adapt to traffic patterns and changing network conditions, improving both performance and power efficiency.

2) SP Operation: During the agreed SP, the STA wakes up from its doze state and becomes active to transmit and receive data. Depending on the direction of the data, DL or UL, the STA can use different operation modes. First, uplink data can be transmitted in two ways: *trigger-enabled* mode,

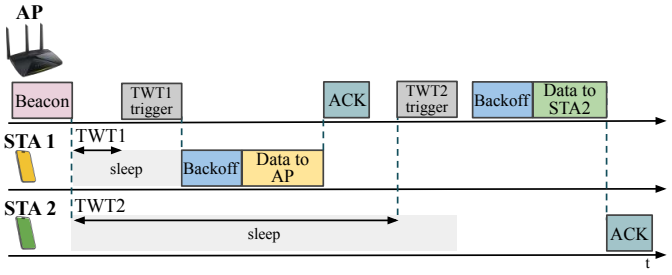

(a) Stations using trigger-enabled TWT under an individual agreement.

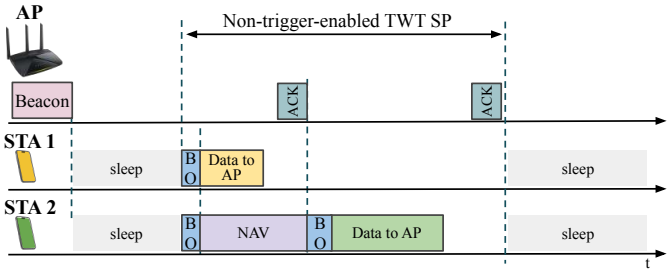

(b) Stations using non-trigger-enabled TWT under a broadcast agreement.

Fig. 23: Example of a two-STA network operating using TWT. (a) Trigger-enabled TWT: the AP sends a trigger frame to announce the transmission or reception of data for specific STAs during the wake intervals. (b) Non-trigger-enabled TWT: the different STAs allocated to a given SP contend for the medium during the wake intervals.

whereby the AP uses trigger frames to explicitly schedule the transmission and reception of data during the current TWT wake interval (Fig. 23a), and *non-trigger-enabled* mode, where the STA autonomously decides when to transmit or receive data during its wake interval (Fig. 23b). Regarding the downlink, two more modes are supported: *announced TWT*, where the STA must send a request (e.g., a PS-Poll) to the AP to retrieve buffered data, and *unannounced TWT*, where the AP can proactively deliver buffered data without waiting for a request, leveraging the fact that the STA is awake at the scheduled time.

TWT enables devices to remain in low-power sleep mode for extended intervals, often minutes, hours, or even days, depending on the application. When the agreed wake time arrives, the STA wakes up, exchanges data with the AP, and returns to sleep. This process ensures highly efficient communication for STAs with predictable or periodic data requirements. Finally, with the enhancements provided by MLO (see Section VII), non-AP MLDs can negotiate multiple TWT SPs, one per active link, by sending multiple requests through one of the links.

C. Dynamic Power Save

STA DPS: Unlike previous power-saving methods, 802.11bn defines a DPS mechanism that allows nodes to dynamically switch between high capability (HC) and low capability (LC) modes in real time, based on the exchange of specific control frames. DPS applies to both STAs and APs, though it is expected to be more commonly used by STAs, which typically face stricter power constraints.

The HC mode functions similarly to the traditional awake state, whereas LC is a partially active mode, not a full doze. In LC mode, the STA can receive a limited subset of legacy frames, constrained by parameters such as bandwidth, spatial streams, and maximum MCS. For example, LC mode may restrict operation to 20 MHz, one spatial stream, and a low MCS [115]. This enables the STA to wake instantly upon receiving an Initial Control frame (ICF), transmitted by its associated device.

Wake-up introduces latency due to the need to initialize PLLs and activate RF and digital circuits. During a preliminary configuration phase, the STA specifies its required delay, which the associated device accommodates by inserting padding into the ICF. An intermediate Frame Check Sequence (i-FCS) is included before the padding to allow decoding in advance, ensuring the STA is ready by the end of the ICF transmission. Upon meeting certain conditions (e.g., traffic reduction), an STA in HC mode may return to LC mode.

AP DPS: Historically, Wi-Fi power-saving mechanisms have focused on STAs, given their reliance on batteries, while APs are typically mains-powered. However, APs still consume significant energy while idle, due to continuous medium sensing. With increasing environmental concerns, energy costs, and AP complexity, reducing their power consumption is a growing priority [115], [116].

To address this, 802.11bn introduces a DPS mode for mobile APs, with possible extension to non-mobile APs under consideration [117]. Like STA DPS, the mechanism defines LC and HC states that alternate based on communication needs. AP LC mode may include strategies such as maintaining one active link, operating on 20 MHz, reducing spatial streams, or capping MCS.

A key challenge for AP power saving, unlike for STAs, is maintaining compatibility with legacy (pre-802.11bn) STAs. Any power-saving scheme must preserve uninterrupted communication with the associated STAs. At the time of writing, the final specification for AP power save remains open. Several proposals exist: one suggests that the AP remains in LC mode for listening, switching to HC only when transmitting or receiving [112]. Another proposes that a mobile AP stays in LC mode by default unless explicitly triggered by an STA, e.g., via an ICF [118]. Additionally, periodic AP power-saving schedules may be coordinated using broadcast TWT [119].

VII. Wi-Fi Multi-link Operation

MLO is widely regarded as the standout innovation introduced in 802.11be [97], enabling devices to operate dynamically on multiple channels. MLO opens up great opportunities, including increased throughput via link aggregation, reduced latency by improving channel access likelihood, enhanced reliability through redundancy across links, and the ability to separate traffic types—such as control and data planes—across different frequency bands [120], [121]. Existing Wi-Fi devices already use multiple links in the 2.4 GHz, 5 GHz, and 6 GHz bands, but these links operate through independent associations. In contrast, MLO integrates and manages multiple links through a single association, leveraging them in a coordinated and highly dynamic manner for maximum efficiency and performance. MLO is also regarded as a future-proof framework for integrating and coordinating additional

links, such as those operating at higher frequency carriers. This presents an opportunity to reconsider existing 802.11 amendments for mmWave technology (see Appendix A-B).

A. Multi-link Architecture and Key Procedures

MLO-capable devices adopt a new architecture to support multi-link operations. In turn, MLO comes with new namings for AP and STA devices, which are now referred to as AP MLD and STA (non-AP) MLD. Unlike previous definitions of APs and STAs, MLDs have multiple PHY interfaces, each capable of operating on different frequency bands or channels. Therefore, an AP/STA MLD comprises two or more legacy APs/STAs that support multi-link operations. Links are established between pairs of legacy APs and STAs within the MLDs. The architecture divides responsibilities across two sublayers: the upper MAC sublayer handles MLD-level shared functions, the lower MAC sublayer manages per-link operations.

In the following, we delve into the key aspects of MLO, beginning with its discovery and setup procedures and extending to the practical implementations and control mechanisms that enable multi-link functionality.

1) Multi-link Discovery: This phase involves either passive or active scanning procedures. In passive scanning, an MLD listens on each link and waits to detect Beacon or unsolicited Probe Response frames. In contrast, active scanning involves the MLD sending Probe Request frames on each of its available links, soliciting faster responses. The discovery process works as follows: an STA identifies an AP and requests information to determine whether to associate with it. Multi-link discovery extends this process to MLDs, where an STA MLD gathers information about multiple APs across multiple links by allowing a MLD to advertise information about all its available links through a single link. As a result, an STA MLD does not have to perform separate scanning on each band/link, reducing the need for redundant scanning and accelerates link setup. During multi-link discovery, the MLD can share either basic or complete information:

- *Basic information* includes essential details about a reported AP on a specific link, such as operating channel, basic service set identifier (BSSID), short service set identifier (SSID), and BSS parameters. To facilitate this, each legacy AP in an AP MLD includes a Reduced Neighbor Report (RNR) element in its Beacon or Probe Response frames. This RNR element provides basic information about all other APs within the same AP MLD.
- *Complete information* comprises all parameters related to a reported AP on a specific link, equivalent to the data included in a Beacon or Probe Response for that link. For this purpose, a new frame type, the Multi-Link Probe Request, allows an STA MLD to request complete information for all APs in an AP MLD in a single transmission on one link. Upon receiving the request, the AP MLD responds with a Multi-Link Probe Response frame, containing complete information for all affiliated APs.

2) Multi-link Setup and Traffic-to-link Mapping: Multi-link setup is a key procedure introduced to streamline the association and configuration of multiple links within an MLD. The association request and response can be performed on a single link, carrying the necessary information to set up all the links in the MLD and eliminating the need for separate associations for each link.

Traditionally, the association process allowed only one link to be set up per association, meaning three links would require three separate association executions. Multi-link setup addresses this limitation by enabling the exchange of capabilities and setup procedures for all links in a single execution. This is achieved by reusing existing (Re)Association Request and (Re)Association Response frames, enhanced with a new Multi-link Element. The Multi-link Element includes a Common Info field that carries MLD-level information shared across all STAs, plus one or more STA Profile sub-elements that provide complete information for each STA operating on a specific link. While legacy BA agreements are link-specific, MLD BA agreements apply across all links. Additionally, the receive status of QoS Data frames may be signaled on other links, ensuring seamless multi-link operation.

A critical procedure in multi-link operation is mapping traffic across different links. This includes modifying traffic indication maps to reflect the entire MLD and associating different TID with individual links to optimize traffic separation and prioritization. The TIM element indicates buffered data for the entire MLD, providing a unified view across links. In Default Mode, all TIDs are mapped to all links. In Optional Mode, specific subsets of TIDs can be mapped to particular links, enabling traffic separation and prioritization. A link is considered enabled if at least one TID is mapped to it. By combining the multi-link setup and efficient traffic-to-link mapping, 802.11 networks can achieve enhanced performance, better traffic management, and improved resource utilization across multiple links.

B. Multi-link Implementations

MLO can be implemented in various forms, each offering distinct capabilities and trade-offs. These implementations include enhanced multi-link single radio (EMLSR) and enhanced multi-link multi-radio (EMLMR), illustrated in Fig. 24 and summarized below.

Enhanced Multi-link Single-radio (EMLSR): An EMLSR device listens to two or more links simultaneously, but only for CCA and limited control frame reception. This approach supports opportunistic spectrum access at a reduced cost, as it requires only one fully functional 802.11be radio, supplemented by low-capability radios that can decode control frame preambles. Upon receiving an initial control frame on one of the links, the device switches to that link and operates using all antennas.

Enhanced Multi-link Multi-radio (EMLMR): In EMLMR, all radios are 802.11be-compliant, enabling concurrent operation across multiple links. EMLMR is further classified into:

- *Simultaneous transmit and receive (STR):* Allows simultaneous uplink and downlink transmissions over a pair of links, maximizing flexibility and throughput.

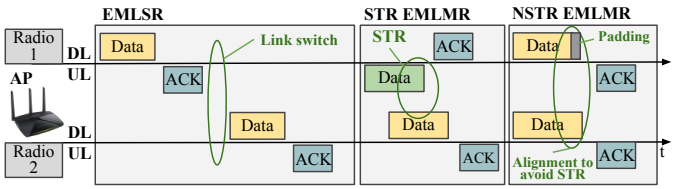

Fig. 24: Illustration of different MLO implementations. Left: EMLSR operation, where traffic is switched between two links (radios) over time. Middle: STR EMLMR enables simultaneous transmissions on both links with overlapping data and ACK frames. Right: NSTR EMLMR uses padding and temporal alignment so that transmissions on the two links do not overlap, thereby avoiding STR operation.

- *Non-simultaneous transmit and receive (NSTR):* Prevents simultaneous transmissions to avoid self-interference. For example, data and ACK transmissions over multiple links must start and end at the same time, as illustrated in Fig. 24.

The flexibility of EMLMR lies in its ability to allocate variable RF chains across links. For example, with four RF chains, three could be assigned to a 5 GHz link and one to a 6 GHz link, allowing the device to optimize the number of spatial streams on each link.

Considerations on MLO implementations: Studies indicate that under low contention, STR EMLMR—the most flexible implementation—can support significantly higher traffic loads and throughput compared to single-link configurations, given the same delay requirements [120]–[125]. However, in high-load and high-contention scenarios, STR EMLMR devices frequently access multiple links, potentially blocking contending BSSs. This can occasionally result in larger worst-case delays compared to legacy single-link systems [126], [127].

C. Seamless Roaming

Seamless roaming in 802.11bn [11] builds upon the MLO framework and is designed to enhance the mobility of STA MLDs by allowing them to transition smoothly between different AP MLDs without losing connectivity. This capability is crucial for maintaining continuous service in environments with multiple overlapping wireless networks, such as large enterprises or public areas. The seamless roaming process consists of two main procedures: preparation and execution, as detailed next.

1) Roaming Preparation Procedure: This involves the transfer or renegotiation of the context related to the STA MLD from the current AP MLD to the target AP MLD. The context includes essential information required for maintaining the connection and ensuring a seamless transition. The preparation also includes setting up the necessary link(s) with the target AP MLD. This step ensures that the STA MLD is ready to switch to the new AP MLD when the execution phase begins, while avoiding any interruption of the data stream thanks to the possibility of maintaining connectivity with the current AP MLD. The specifics of what context can be transferred or renegotiated are yet to be determined.

2) Roaming Execution Procedure: The STA MLD initiates the execution phase by sending a Request frame to the current AP MLD. This frame signals the intention to transition to a new AP MLD. Upon receiving the request, the current AP MLD may continue to send individually addressed downlink data frames to the STA MLD for a predetermined period. This ensures that the STA MLD continues to receive data during the transition. The current AP MLD is responsible for transferring the necessary context to enable operations with the target AP MLD. This context transfer is crucial for service continuity. After the context transfer or renegotiation is completed, the current AP MLD sends a response frame to the STA MLD, confirming the completion of the transition process.

Overall, seamless roaming in 802.11bn is focused on minimizing the time during which connectivity is lost, ensuring that applications requiring real-time data transmission, such as voice over IP (VoIP) and video streaming, are unaffected by the transitions between AP MLDs.

VIII. Wi-Fi Multi-AP Coordination

Due to their distributed nature, one of the primary limitations of 802.11 networks arises from inter-BSS interactions, which can lead to collisions and unpredictable worst-case delays. This issue is especially pronounced in dense and heavily loaded environments, where contention and interference effects are magnified. Although recent 802.11 amendments have introduced mechanisms to improve channel access within a BSS, such as BSS coloring (see Section IV-F1), the absence of explicit coordination among APs belonging to different BSSs remains a major obstacle to achieving high reliability in Wi-Fi. In the past, standardization bodies like WFA or IETF have developed protocols and functionalities that provide inter-APs backhauls (e.g., Wi-Fi EasyMesh, CAPWAP), enabling the development of wireless controllers and unlocking functionalities such as roaming, mesh networking, or optimized radio resource management. However, these solutions often lack interoperability and are incapable of providing high-grained performance optimization (because they operate at layers above the PHY and MAC) and are instead better suited for longer-term network management (e.g., channel allocation).

The 802.11bn amendment aims to overcome the limitations of spectrum sharing in unlicensed bands and advance the reliability of future Wi-Fi. MAPC is expected to play a central role in this by enabling multiple APs to coordinate over-the-air and perform more efficient operations, thereby reducing channel contention and interference, and improving resource utilization. This section presents the framework and coordination modes currently under consideration in 802.11bn [128]: coordinated TDMA (Co-TDMA), coordinated R-TWT (Co-RTWT), Co-SR, and Co-BF.

A. MAPC Framework

Each MAPC mechanism targets specific challenges in resource allocation, interference mitigation, and performance enhancement, offering the flexibility needed to adapt to diverse deployment scenarios. As a result, the implementation complexity and signaling requirements of each MAPC approach may vary, depending on the volume, type, and periodicity of data exchanged between APs.

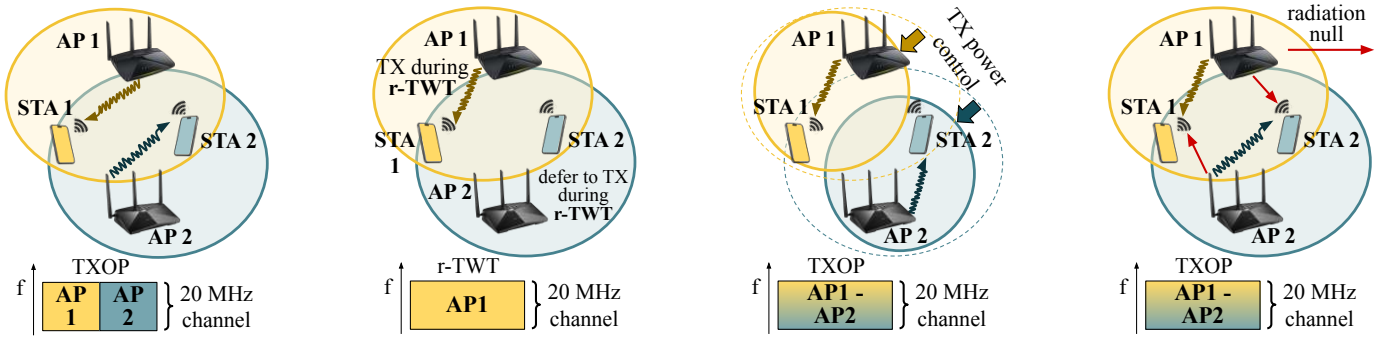

(a) Co-TDMA. (b) Co-R-TWT. (c) Co-SR. (d) Co-BF.

Fig. 25: Illustration of four different MAPC mechanisms for 802.11bn in a scenario with two overlapping APs and two STAs. In (a), Co-TDMA enables time-division sharing of the same 20 MHz channel between AP1 and AP2 within a TXOP. In (b), Co-R-TWT reserves an rTWT service period for AP1 while AP2 defers its transmissions. In (c), Co-SR uses coordinated transmit-power control so both APs can transmit simultaneously with limited inter-BSS interference. In (d), Co-BF employs joint beamforming to create radiation nulls toward non-intended receivers, enabling fully overlapping concurrent transmissions on the same channel.

802.11bn defines a common framework for MAPC, designed to support multiple coordination schemes. This framework includes the following key operations:

- *MAPC discovery:* To initiate coordination, APs must advertise their coordination capabilities, e.g., support for responding in a TB PPDU. These capabilities are typically broadcast through management frames, such as Beacons or Probe Request/Response frames, allowing neighboring APs to discover and evaluate potential coordination partners.
- *MAPC negotiation:* Once neighboring APs discover one another and agree to coordinate, they establish a formal agreement defining the conditions for coordinated transmissions. This includes negotiating roles (e.g., coordinator vs. responder), maximum TXOP durations for simultaneous transmissions, and transmit power limits. Negotiation occurs via management frames, such as Public Action frames or new individually-addressed Action frames. While 802.11bn limits each MAPC agreement to two APs, it permits multiple concurrent agreements.
- *MAPC session:* With an agreement in place, APs can initiate coordinated transmissions. Although each MAPC scheme has distinct requirements, a unified mechanism is defined to trigger coordination. Specifically, the *sharing AP*, i.e., the one holding the TXOP, sends a Trigger frame to notify the *shared AP* of the start of a coordinated transmission.
- *MAPC teardown:* A coordinated session can be terminated using new teardown messages, which may be initiated by any participating AP.

All the aforementioned procedures require a significant architectural transformation from 802.11be, which currently lacks the necessary messages and mechanisms to support inter-AP communication. As noted earlier, new fields, formats, and even message types, such as MAPC Discovery, MAPC Negotiation Request, and MAPC Negotiation Response, will need to be defined. The 802.11bn amendment considers over-the-air communication among APs, which, unlike alternative wired backhauls, allows a sub-millisecond synchronization that is necessary to perform joint transmission operations in real time.

To ensure the MAPC framework is generalizable across coordination schemes, it is essential to introduce flexible signaling structures, such as a dynamic Universal Signal (U-SIG) field, where MAPC-specific information can be adapted as needed. Additionally, new MAC layer management entity (MLME) primitives will be required to support core operations, e.g., to trigger the transmission of a MAPC Negotiation Response upon receipt of a corresponding request.

From a management perspective, identifiers such as AP IDs and session IDs must be maintained to track eligible and coordinated APs, as well as active and past coordination sessions. Finally, security mechanisms must be extended to enable the secure exchange and management of cryptographic keys between coordinating APs, potentially leveraging mechanisms like pre-association security negotiation (PASN).

B. MAPC Mechanisms

This section provides an overview of the four candidate MAPC mechanisms proposed for 802.11bn, as illustrated in Fig. 25.

1) Coordinated TDMA (Co-TDMA): Co-TDMA is one of the foundational approaches to MAPC, leveraging coordination in the time domain [129]. In Co-TDMA, the sharing AP divides its TXOP into multiple time slots, distributing each slot among itself and the shared APs. This method aims to reduce channel contention and improve resource allocation efficiency. As illustrated in Fig. 25a, two APs (AP 1 and AP 2) can split the duration of a TXOP to transmit to their respective STAs, showcasing the fundamental principle of time-based coordination in Co-TDMA.

Co-TDMA operates in three phases [129]: the sharing AP first polls shared APs with an ICF; then allocates TXOP slots based on responses and access category priorities; and finally, allows early TXOP return if a shared AP finishes before its slot ends. The detailed scheduling logic is implementation-specific and may vary across vendors and system objectives. To prevent interference during coordination, Co-TDMA uses

enhanced NAV protection mechanisms, including short NAV settings during polling, dynamic updates during TXOP return, and backup strategies like EDCA timers or RTS gating [129].

A key limitation of Co-TDMA is its inability to achieve spatial reuse. The time slots allocated to shared APs are carved directly from the sharing AP's own TXOP, resulting in no net increase in overall airtime utilization. Furthermore, performance may degrade in the presence of mismatched traffic loads or access categories. For instance, in unbalanced OBSS scenarios [129], an AP serving mostly low-latency traffic may find limited benefit in sharing with an AP focused on throughput-heavy traffic, leading to poor slot utilization, wasted return time, or coordination inefficiencies.

Despite these limitations, Co-TDMA provides a simple and deterministic framework for multi-AP coordination. By avoiding contention and introducing structured access, it improves latency predictability, particularly in dense or delay-sensitive deployments.

2) Coordinated R-TWT: Co-RTWT is a mechanism that enhances reliability by allowing APs to share and coordinate their R-TWT schedules (see Section IV-D) [130]. The ability to secure unlicensed spectrum resources across multiple BSSs represents a major breakthrough for Wi-Fi. Fig. 25b illustrates the Co-RTWT principle, where AP 1 successfully protects its transmission to STA 1. This is achieved through a coordination agreement with another BSS (AP 2 and STA 2), which remain silent during AP 1's scheduled SP. Such lightweight coordination makes Co-RTWT particularly suitable for reliability-sensitive applications in OBSS scenarios.

To enable Co-RTWT, the two APs must negotiate an R-TWT schedule by exchanging frames, either existing or newly defined management frames (to this date, the specific negotiation and advertisement of R-TWT schedules between APs has not been defined). These frames convey key information such as the TWT Interval, the Broadcast TWT ID, or the Nominal Minimum TWT Wake Duration (notice that other parameters might also be included later, such as the parameters needed to time synchronize the APs). Multiple R-TWT agreements can be negotiated in parallel, each of which is included as a Co-RTWT Parameter Set field. Once these agreements are in place, each AP must ensure that neither its own transmissions nor those from its associated STAs overlap with the scheduled SPs of the other AP. For that, each involved AP must advertise, e.g., by extending Beacon frames, the OBSS R-TWT schedules to its STAs.

Similarly to Co-TDMA, Co-RTWT focuses on a better utilization of time resources, so no net gain is expected. However, Co-RTWT offers good potential for achieving reliability in OBSS scenarios.

3) Coordinated Spatial Reuse (Co-SR): Co-SR enables APs to cooperatively manage their transmit power, facilitating concurrent transmissions and improving spectral efficiency in dense environments [131]. This approach enhances the moderately adopted spatial reuse operation (see Section IV-F), where one AP transmits at full power, while others reduce their power, often resulting in suboptimal SINR for some STAs. Fig. 25c illustrates the principle of Co-SR, where AP 1 and AP 2 adjust their transmit power to enable concurrent transmissions within a shared TXOP, maximizing spectral reuse without compromising link quality.

To enable Co-SR, two APs exchange the necessary information to derive transmit power levels that ensure an acceptable SINR at all receiving STAs during simultaneous transmissions. A practical method involves measuring and sharing the received signal strength indicator (RSSI) of interfering links, e.g., between the sharing AP and the shared AP's STA. Existing mechanisms support frequent RSSI measurements and centralized reporting at the AP, making Co-SR viable with minimal coordination overhead.

Once transmit power values are negotiated, the coordinated transmission is triggered by the shared AP through the transmission of a Trigger frame (which one remains to be decided), indicating the shared TXOP duration. From the shared AP side, the transmit power employed cannot exceed the one previously negotiated. 802.11bn defines two Co-SR transmission modes, differentiated by the indication of different information in the U-SIG: *Mode 1*, whereby the different combinations of UHR and EHT STAs are considered, and *Mode 2*, which defines transmissions involving UHR devices only. In either case, and for the sake of ensuring reliable decoding, both the sharing and shared AP must begin and complete their PPDU transmissions simultaneously. In the UHR-only case (Mode 2), the coordinated transmission starts a SIFS after the trigger.

Co-SR is a feature that only accounts for SU DL transmissions and which limits the maximum number of spatial streams to four. However, its potential for improving the delay is substantial thanks to the added capability of enabling concurrent transmissions from multiple BSSs, provided those BSSs have data ready to send at the same TXOP.

4) Coordinated Beamforming (Co-BF): Co-BF is a spatial-domain coordination mechanism in which multi-antenna APs suppress interference to and from neighboring STAs belonging to other BSSs [132]. By leveraging spatial degrees of freedom, each AP can serve its own associated STAs, while simultaneously placing radiation nulls toward non-associated STAs [133], [134]. This suppression of spatial interference reduces contention and enables concurrent transmissions, improving both spectral reuse and worst-case latency in dense deployments. Fig. 25d illustrates a concurrent transmission scenario enabled by Co-BF, where AP 1 and AP 2 steer radiation nulls towards STA 2 and STA 1, respectively, to suppress mutual interference.

In Co-BF, the sounding phase is critical, and therefore, it is being carefully designed. In particular, two sounding procedures are envisioned in 802.11bn: (i) *Sequential sounding*, whereby NDPA and NDP exchanges and CFR feedback reporting occur in one BSS at a time (the same EHT Compressed Beamforming/CQI report and procedures detailed in Section III-D are adopted), and (ii) *Joint sounding*, where a new protocol allows for multiple APs to simultaneously perform sounding with a given STA. The type and conditions associated with sounding, e.g., the number of OBSS sounding reports that can be stored at a time, are agreed upon during the negotiation phase of Co-BF. Furthermore, before starting each sounding phase, the APs involved in Co-BF are expected to exchange certain frames (which are to be decided) to indicate

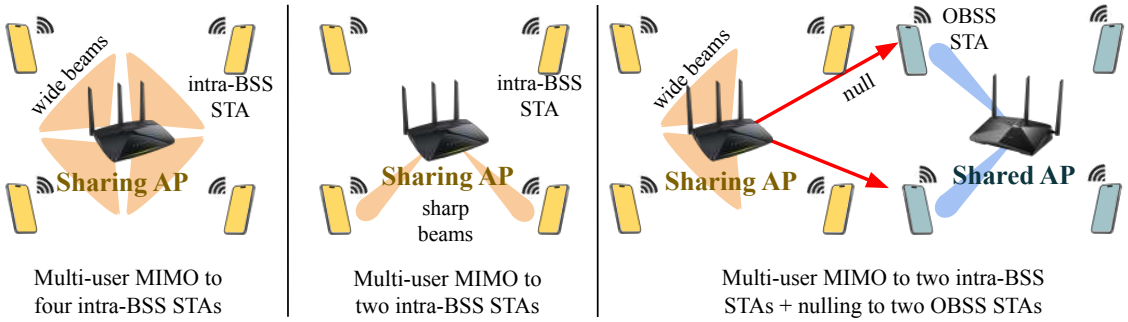

Fig. 26: Three examples of employing four spatial degrees of freedom (DoF) in Co-BF. Left: all four DoF are used to spatially multiplex four intra-BSS STAs, resulting in limited beamforming gain. Middle: two DoF are used to spatially multiplex two intra-BSS STAs, while the remaining two DoF provide additional beamforming gain. Right: two DoF are used to spatially multiplex two intra-BSS STAs, and the remaining two DoF place nulls at two OBSS STAs, enabling spatial reuse.

availability for participating in the process and inform about the capabilities of the STAs to be sounded. Finally, to address CFO correction, it is worth noting that the sharing and shared APs take the role of *sync-reference* and *sync-follower* APs, respectively.

Once beamforming information is available at both AP sides, the coordinated transmission can take place. Before initiating the transmission of the Co-BF PPDUs, the sharing AP first sends a newly defined Co-BF Invite frame, which is used to indicate parameters such as the minimum and maximum number of supported OFDM symbols, the selected bandwidth, information regarding the punctured channel, the total number of spatial streams allowed for the sharing AP, or the number of Co-BF STAs and corresponding spatial streams in the sharing BSS. Such a message is replied to by the shared AP with a Co-BF Response frame, which indicates the acceptance or not of undergoing a Co-BF transmission, plus the suggested number of OFDM symbols (not smaller than the initially proposed) and information about the number of STAs and their assigned MCS and spatial streams in the shared BSS. The setup is concluded by a Co-BF Sync frame, which the sharing AP sends to confirm the final set of parameters.

Unlike joint transmission schemes where all APs cooperate to jointly serve all STAs [135], Co-BF does not require joint data processing or tight backhaul integration. Each STA is served by a single AP, and only the channel information is required to coordinate null placement. This design makes Co-BF more practical for decentralized deployments, albeit still dependent on reliable CFR acquisition.

However, the performance of Co-BF hinges on several critical factors. First, spatial DoF—limited by each AP's antenna array—must be split among competing goals: (i) spatial multiplexing of multiple intra-BSS STAs, (ii) enhancing SINR via beamforming gain, and (iii) placing nulls to suppress interference to OBSS STAs. Fig. 26 shows three representative configurations using four DoF, highlighting the trade-offs among these objectives.

Second, the accuracy of available CFR is constrained by the channel coherence time, which depends on STA mobility and environmental dynamics. As the channel ages, outdated CFR may result in imperfect nulls, degrading SINR, and reducing throughput. While allocating more resources to channel estimation can reduce these effects, it introduces overhead that must be carefully managed.

Lastly, each AP must intelligently select which intra-BSS STAs to serve and which OBSS STAs to null. These decisions are influenced by the quality of CFR, the SNR of each candidate STA, and the spatial correlation between them. Together, these challenges underscore the importance of efficient channel information acquisition, robust scheduling, and dynamic user selection to achieve the full potential of Co-BF.

IX. Conclusions

With Wi-Fi over 25 years old, it is perhaps not surprising that there are decades of Wi-Fi tutorials published already, including some by the authors of this paper. Notwithstanding, this paper is the first tutorial that spans Wi-Fi's full evolution from 802.11b to the emerging 802.11bn (Wi-Fi 8). We describe how use cases have evolved from basic web browsing to high-resolution video requiring high throughput, to real-time critical applications requiring low latency and high reliability. To keep up with, and indeed drive such application demand, the underlying technologies of Wi-Fi have radically transformed over these generations. Major early milestones include introduction of OFDM and MIMO. The next generations brought features such as multi-user technologies including uplink and downlink OFDMA and multi-user MIMO. Energy saving methods such as Power Saving Mode and Target Wake Time are increasingly important as Wi-Fi overwhelmingly dominates data usage in mobile phones. Recently, Multi-Link Operation and Multi-AP Coordination are yielding new capabilities for providing not only higher throughput, but also radically improving reliability and reducing latency. This article aims to be a comprehensive resource for researchers and practitioners to describe Wi-Fi's technical journey, current status and the road ahead.

Building on the advancements envisioned for 802.11bn, the Appendix explores parallel efforts shaping the future of Wi-Fi. Key areas of focus include mmWave operations, sensing, security and privacy, and the integration of AI/ML techniques into Wi-Fi.

Appendix A
Task Group 802.11bq: Integrated mmWave Operations

The primary fuel for innovation in wireless technology has always been the availability of spectrum. The regulatory

experiment in the 2.4 GHz band with approximately 80 MHz allocated for unlicensed operation gave rise to technologies such as Wi-Fi and Bluetooth in late 1990s. Realizing the success of these technologies, regulators around the world made additional unlicensed spectrum available in the 5 GHz band in early 2000s and, more recently, in the 6 GHz band. However, wireless spectrum is a finite resource with a demand much greater than the available supply. Therefore, there is a growing fear that finding a large contiguous swath of unlicensed spectrum below 10 GHz will be extremely difficult, if not impossible. Among other things, this means that increasing the channel bandwidths of present Wi-Fi above 320 MHz or adding new channels will be challenging, which then limits Wi-Fi's future ability to meet the stringent demand of emerging applications in highly dense environments, such as use of VR/AR in classrooms and wireless docking stations in enterprises.

To address this challenge, proposals have emerged to leverage the large swaths of unlicensed spectrum available in the 60 GHz mmWave band (e.g., 57 GHz to 71 GHz in the USA and EU) to enable an extension to Wi-Fi that can deliver, among other things, channels wider than 320 MHz for high throughput operation, ultra-low latency wireless connectivity, and networks that can operate in highly dense environments given the large number of channels that can be created. Needless to say, operation in mmWave bands is not without its challenges. Besides a higher energy consumption, operating in these bands entails facing rapid signal attenuation and susceptibility to blockages. One approach to mitigate these losses is through beamforming gain, where most spatial degrees of freedom are allocated for beamforming rather than for transmitting multiple spatial streams. This enables improved link reliability and performance, even under challenging propagation conditions.

In early 2025, the 802.11bq Task Group, also known as Integrated mmWave (IMMW), was formed to explore these opportunities further. The goal is to develop a new amendment for dynamically operating mmWave links using the PHY/MAC functionalities derived from 802.11ac, 802.11be, and 802.11bn. Future high-end devices are expected to operate not only in the traditional sub-7 GHz bands, but also in the 60 GHz band.

A. Background: 802.11ad and 802.11ay (WiGig)

801.11bq is not the first project in 802.11 dealing with operation in mmWave. Between 2010 and 2020, two other amendments were developed that defined operation in mmWave: 802.11ad and 802.11ay, collectively referred to as WiGig. 802.11ad operates on a single channel with a bandwidth of 2.16 GHz, while 802.11ay supports channel bonding (up to four 2.16 GHz channels) and MU-MIMO (up to 4 spatial streams), significantly enhancing throughput and capacity. Despite these improvements, the market adoption of 802.11ad and 802.11ay has been limited and mainly in specific segments such as fixed wireless access. The main reasons for this limited adoption are the costs and complexity associated with defining a completely new mmWave modem that does not reuse any component of the existing 2.4/5/6 GHz modem. These have a ripple impact on the design, development, silicon area, power consumption, and validation of such solutions.

B. Overview of 802.11bq

802.11bq is the current project that promises to overcome the shortcomings faced with 802.11ad and 802.11ay. 802.11bq has a stated goal of building on the existing 802.11ac/be/bn PHY and MAC features and channelization schemes. This means that the basic modulation scheme of 802.11bq will be based on OFDM, upclocking the modem used in the 2.4/5/6 GHz bands. As for channelization, 802.11bq could adopt channel sizes that are multiple of 160 MHz or 320 MHz as defined in the lower bands.

In the MAC, a key innovation will be the reuse of the MLO framework introduced in 802.11be (see Section VII), which could be extended to manage operations across the lower and mmWave bands, dynamically activating links as needed to adapt to varying conditions. With MLO, the AP would be able to dynamically (in the order of microseconds) switch between lower and mmWave bands to exchange packets with STAs. Moreover, given that links in the lower bands will be more robust than those in mmWave, using MLO will make it possible to move all control and management traffic to the lower bands while dedicating the use of the high capacity and ultra-low-latency links in mmWave bands for data exchanges only. For example, the lower bands could be used for beacon frame transmissions, authentication and key establishment, association, and similar management exchanges, which would also encompass operation in mmWave. Then when data transmission is needed with an STA that finds itself in mmWave coverage, the AP can choose whether to use the lower, mmWave, or even both bands to exchange data with the STA.

APPENDIX B
TASK GROUP 802.11BF: WI-FI SENSING

Sensing is a transformative application of Wi-Fi, which will allow devices to "see" the surrounding environment, providing new services to the users, and improving the network management procedures [136]. The increasing interest in this feature led to the instantiation of a new TG in 2020 to design an amendment, identified as 802.11bf, that will support sensing in the sub-7 GHz and mmWave portions of the radio spectrum. The key idea is that wireless devices are inherently able to sense the environment thanks to the channel estimation procedure that is continuously performed for data decoding and precoding as thoroughly detailed in Section V. The CFR provides information about how transmitted signals are modified on their way toward the receiver due to the multipath propagation phenomenon. As the multiple paths are generated by objects and people in the environment, the CFR can be effectively used as a proxy of the physical environment to passively sense it. Importantly, the sensing targets do not need to carry a Wi-Fi device: their presence is detected as they act as reflectors, diffractors, or scatterers for the signal propagation. By analyzing this information through over-the-top signal

processing algorithms and machine learning tools, researchers have been able to design and implement several applications such as human activity and gesture recognition [137], people tracking [138], vital sign monitoring [139], and many others [140], [141].

A. Motivation

A question that always arises when talking about Wi-Fi sensing is about the motivations behind and the benefits of performing sensing through radio waves, and in particular through a wireless network. Indeed, several other technologies are already available in the market to perform sensing, such as cameras and lidars. However, these technologies face problems in bad illumination conditions and when obstacles are present in the environment. Moreover, the use of camera systems comes with privacy issues when they operate in the presence of people. Radio waves allow overcoming all these issues as they can provide sensing information also in the dark and in the presence of smoke, and they do not capture the images of people around. For these reasons, radar devices have been extensively used to perform sensing. The use of Wi-Fi signals for this comes with additional benefits linked with the widespread adoption of this technology that, in turn, enables reusing already deployed infrastructures, extending their utility beyond traditional connectivity. However, this requires designing new procedures to integrate the two functionalities, which have quite different requirements [43]. While communication waveforms are generated only when there is data to transmit, sensing requires a continuous transmission of signals to trigger the channel estimation at the monitoring devices. Moreover, communication devices are usually characterized by a smaller bandwidth with respect to radars, which impacts the sensing resolution. Indeed, the resolution is inversely proportional to the bandwidth B as $c/(2\times B)$ or c/B for mono-static and bi-static sensing, respectively, where c is the speed of light. Operating in the 2.4 and 5 GHz bands, where bandwidth is limited, provides meter-level resolution, while larger bandwidths available in the mmWave portion of the spectrum enable centimeter-level resolution. However, sensing accuracy improves with multiple receivers and a larger number of Wi-Fi transmitters, making Wi-Fi sensing particularly promising given the proliferation of Wi-Fi [142].

Current sensing algorithms mostly rely on custom firmware modifications which are crafted for some specific devices, such as Nexmon CSI [143], AX CSI [144], and PicoScenes [145]. Another approach is to rely on the beamforming feedback matrices introduced in Section III-D as described in [146]. This last strategy is hardware-agnostic as all Wi-Fi devices operating in SU-MIMO or MU-MIMO mode have to feed back such information. However, all these approaches require triggering the transmission of data for generating the CFR or the beamforming feedback matrices. Instead, the goal of the 802.11bf TG has been to define modifications to the 802.11 PHY and MAC to enhance sensing capabilities in 802.11-compliant devices. Most sensing operations are based on existing beamforming protocols and are time-separated from data communications. Note that while the measurement

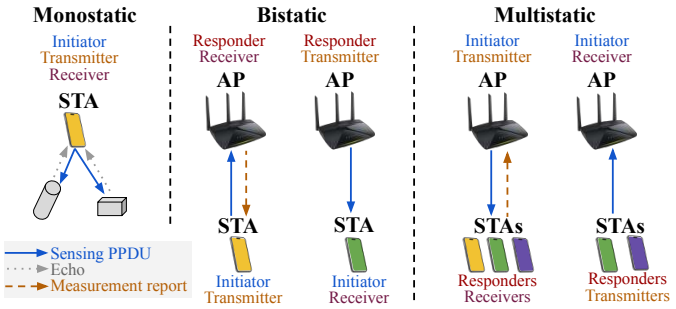

Fig. 27: Examples of monostatic, bistatic and multistatic sensing configurations. In the monostatic setting, the sensing is entirely executed by a single device that initiates the procedure by transmitting the sensing PPDU and collects the echoes from the environment. In bistatic sensing, the transmitter and receiver are distinct devices, e.g., an AP and a STA, and each of them can act as initiator. In multistatic, more sensing transmitter or receiver are involved.

report formats have been standardized through the 802.11bf amendment, specific sensing algorithms and application layers remain open for vendor-specific implementation [43].

B. Sensing Procedure in 802.11bf

802.11bf defines a unified procedure for obtaining channel measurements between two or more devices (bistatic and multistatic sensing), or between the transmitter and receiver antennas of a single device (monostatic sensing) [42]. The devices involved can be either APs or STAs, and their roles can be identified as follows:

- *Sensing transmitter*: Transmits PPDUs for sensing measurements.
- *Sensing receiver*: Receives PPDUs from the sensing transmitter and obtains sensing measurements. In monostatic sensing, this device is the same as the sensing transmitter, while in bistatic sensing, they are distinct entities.
- *Sensing initiator*: Initiates the sensing procedure and is where the sensing application is hosted. It may participate in sensing as a sensing transmitter, or sensing receiver, or both a transmitter and a receiver, or neither a transmitter nor a receiver.
- *Sensing responder*: Participates in the sensing procedure. It may participate in sensing as a sensing transmitter, or sensing receiver, or both a transmitter and a receiver.

In addition to monostatic and bistatic sensing, multistatic sensing can be performed by involving multiple transmitters and/or receivers. Some examples of sensing configurations in the three different scenarios are depicted in Fig. 27. The IEEE 802.11bf amendment also introduces a "sensing by proxy" mode to enable an STA to enhance sensing accuracy through diversity, with assistance from an AP. While an STA can only perform sensing measurements on the link between itself and the AP, the latter has access to and can perform measurements of the channel of all other STA [42].

A typical sensing procedure consists of four phases:

- *Sensing capabilities exchange*: Devices capable of sensing discover each other and establish security. For associated STAs, this is handled automatically through regular association.

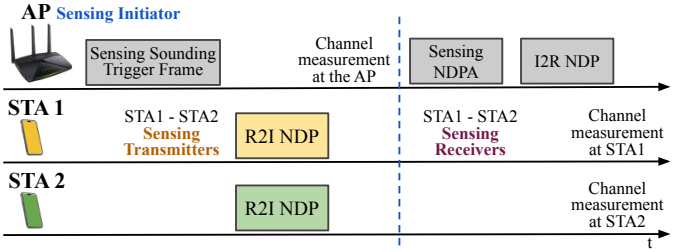

Fig. 28: Illustration of two AP-initiated Sub-7 GHz sensing approaches. On the left, the STAs transmit NDP for channel estimation in response to the AP transmitted trigger frame; the channel is then measured at the AP. On the right, the AP transmits a NDP for channel measurement at the STAs, preceded by a NDPA announcement frame. 'R2I' indicates a frame from the responder to the initiator, while 'I2R' indicates a frame from the initiator to the responder.

- *Sensing measurement session*: The sensing initiator and responder(s) negotiate operational parameters specific to the application, such as role assignments, PHY parameters, measurement report types, and scheduling details.
- *Sensing measurement exchange*: The actual sensing measurements are performed as detailed in Section B-C and Section B-D.
- *Sensing measurement session termination*: The procedure concludes after the channel information collection is completed.

Note that for monostatic sensing, the *sensing session setup* and *sensing measurement setup* phases are not required.

C. Sensing in sub-7 GHz Bands

Only bistatic and multistatic sensing are supported at sub-7 GHz, while monostatic sensing is not allowed. The process may be initiated by the AP or an STA. AP-initiated sensing follows a TB mechanism, as depicted in Fig. 28. The AP can transmit a Sensing Sounding Trigger Frame to stimulate STAs to transmit an NDP for channel measurement at the AP (left part of Fig. 28). The AP can also transmit a Sensing NDPA frame followed by a Sensing NDP to collect measurements performed by the STA through the latter frame (downlink sensing, right part of Fig. 28). The STA-initiated sensing is a non-TB mechanism where the STA initiates the sensing by transmitting a Sensing NDPA frame to the AP (single responder) followed by a Sensing NDP frame for uplink sensing. The AP subsequently transmits a Sensing NDP frame for downlink sensing.

D. Sensing in the 60 GHz Band

Sensing in the 60 GHz band is built on the 802.11ad and 802.11ay amendments [42]. In monostatic sensing, the same STA transmits an (E)DMG PPDU, receives the echo, performs channel measurements leveraging the training fields, and reports the results to the initiator. Bistatic sensing is instead based on the beam refinement protocol (BRP): the sensing transmitter sends a BRP frame with appended training fields to the sensing receiver for measurements. The receiver then responds with a BRP frame containing the channel measurements. Multistatic sensing in the 60 GHz band requires more coordination among devices: the initiator sends a DMG Sensing Request frame to each responder, each of which replies with a DMG Sensing Response and aligns its receiver antennas toward the initiator. To accommodate varying antenna configurations, sensing PPDUs are appended with unique synchronization fields for each STA.

Appendix C
Task Groups 802.11bh and 802.11bi: Security and Privacy Extensions

Security and privacy have been a major focus of every new generation of Wi-Fi [147]. In March 2021, two new task groups were formed to improve the security and privacy of 802.11 networks: 802.11bh and 802.11bi.

A. The 802.11bh Standard Amendment

802.11bh examined the specific implications of random and changing MAC addresses. MAC address randomization can improve user privacy, but at the same time its use can have a significant implication in how existing Wi-Fi networks and services on these networks operate. For example, many existing networks and applications running on such networks rely on MAC addresses to recognize devices, e.g., captive portals, DHCP, and troubleshooting. If these addresses change often, this can disrupt the seamless onboarding of devices into the network, device operation within a network and the delivery of network service to devices. As a result, 802.11bh defines two opt-in methods to identify a device to a network. Both methods are based on sharing information securely over an initial connection that can then be used when the devices (STA and any AP in the network) next encounter each other:

- Identifiable random MAC address (IRM): an STA tells the AP what MAC address it will use the next time it engages with the AP.
- Device identifier (ID): an AP gives the STA a secret identifier to be provided to the network the next time the two engage.

An STA participates in this identification only when it opts-in. The mechanisms are designed to prevent third-party tracking of the STA by its identification. The intention is to only allow the AP to track the STA (i.e., recognize it on future interactions) when there is a shared security context, which prevents network spoofing. 802.11bh also introduces mechanisms that provide the ability to support a device identification pre-association through PASN extension and a Beacon Measurement extension.

B. The 802.11bi Standard Amendment

802.11bi was formed with a focus on improving personal and device privacy, since MAC address randomization (as defined in 802.11bh) has been shown to be necessary but insufficient to ensure user privacy [148]. Users and regulatory agencies are concerned about protecting personal information such as location, movement, contacts, and activities of STAs and their users. 802.11 devices are ubiquitous and protecting users from tracking and profiling attacks has become a major

market need. While Wi-Fi is already being capable of encrypting all frame exchanges between an AP and a station after the association process, prior to association frame exchanges are still transmitted in the open. This can lead to privacy attacks and fingerprinting. To address these challanges, 802.11bi defines key derivation in Authentication frames and subsequent encryption of (Re)Association Request and Response frames. By encrypting (Re)Association Request and Response frames, all configuration-specific information that today is carried in the open within these frames and that could otherwise be used for fingerprinting becomes protected, thereby ensuring a very high level of privacy for an STA and its user.

APPENDIX D
AI/ML TOPIC INTEREST GROUP / STANDING COMMITTEE

The integration of AI into 802.11 is another promising direction for the future of Wi-Fi [47]. The adoption of ML in Wi-Fi, however, requires a major transformation of the 802.11 standard, which lacks the necessary elements and procedures to support specialized ML operations such as training/inference data collection, model training, or model output application. In this regard, the IEEE 802.11 WG formed the AIML TIG in 2022 (turned into SC in 2024), which was entrusted to investigate the potential and feasibility of integrating ML into future 802.11 amendments. The group's primary focus consisted in identifying requirements for specific ML-based use cases for Wi-Fi, which were collected in a technical report [149]. The included use cases are as follows:

- *AIML-based CSI feedback compression:* The sounding procedures in Wi-Fi lead to large overheads (see Section V-B), especially when multiple users, spatial streams, and even multiple BSSs are involved. Henceforth, a promising and much needed use case for AI is to reduce CFR overheads, which can be achieved in multiple ways, e.g., leveraging similar STA channels using unsupervised learning [150] or quantizing CFR matrices efficiently using auto-encoders [79].
- *Deep-learning based distributed channel access:* While valid for many years, Wi-Fi's channel access procedures based on CSMA/CA are inefficient and unreliable. For that reason, ML is identified as a potential game changer in this use case. In particular, it is proposed to perform centralized CW optimization using the DL [151] or even relying on deep reinforcement learning (DRL) to drive the access to the channel [152].
- *AIML enhanced roaming:* Denser deployments make roaming very challenging since a large number of candidate APs need to be handled, thus high roaming times can be experienced. In this use case, AIML is envisioned to provide meaningful information about the different candidate APs for roaming, indicating the most likely ones on a per-STA basis, e.g., based on learned STA patterns [153]. AIML enhanced roaming is proposed be integrated into the 802.11k neighbor reporting mechanism.
- *AIML based multi-AP transmission:* The added complexity introduced by MAPC (see Section VIII) introduces new challenges, such as deciding the set of APs that should collaborate. It has been suggested that AIML could help in identifying devices and networks that can benefit from specific coordination features, e.g., Co-BF.
- *Efficient AIML model sharing:* This use case, which was the only one categorized as a use case that *enables* AIML in WLANs, aims to address the fact that ML applications require large information exchanges, including data, metadata, and models. Regardless of whether ML applications that exchange data over Wi-Fi are used to improve Wi-Fi or not, it is clear that enhancing ML data distribution would be beneficial to preserve QoS. In this regard, potential changes include the possibility of exchanging models at the MAC layer rather than at the application layer. With this, applications based on paradigms such as federated learning (FL) could benefit from fast model exchanges.

So far, it remains unclear what the role of AI in 802.11 is going to be and what ML features could be enabled. The fact that 802.11 focuses on the MAC and PHY layers limits the range of AI functionalities and operations that could be defined by the standard, especially when compared to frameworks like 3rd Generation Partnership Project (3GPP), where both core and radio access network (RAN) are defined. Furthermore, the inherent nature of Wi-Fi and its principles of decentralized channel access and backward compatibility make the adoption of AIML a unique challenge for this technology.

REFERENCES

[1] Wi-Fi Alliance, "Wi-Fi Alliance - The worldwide network of companies that brings you Wi-Fi®," 2025. [Online]. Available: https://www.wi-fi.org/

[2] M. Kasslin, "Confessions of a Wi-Fi pioneer," https://www.nokia.com/blog/confessions-of-a-wi-fi-pioneer/, Dec. 2024.

[3] L. Berlinsky-Schine. Gen Z: 18 statistics about today's newest workers. Fairygodboss. [Online]. Available: https://fairygodboss.com/career-topics/gen-z-statistics

[4] Wi-Fi Alliance, "Global economic value of Wi-Fi® to reach $5 trillion in 2025," 2024. [Online]. Available: https://www.wi-fi.org/system/files/Economic$_{\mathrm{V}}$alue$_{\mathrm{o}}$f$_{\mathrm{W}}$i-Fi$_{\mathrm{H}}$ighlights$_2$02407.pdf

[5] Wi-Fi Alliance, "COVID-19 AND THE ECONOMIC VALUE OF Wi-Fi," 2020. [Online]. Available: https://www.wi-fi.org/system/files/COVID-19$_{\mathrm{E}}$conomic$_{\mathrm{V}}$alue$_{\mathrm{W}}$i-Fi$_2$02012.pdf

[6] K. Pahlavan and P. Krishnamurthy, "Evolution and impact of wi-fi technology and applications: A historical perspective," *International Journal of Wireless Information Networks*, vol. 28, pp. 3–19, 2021.

[7] IEEE Standards Association, "The Evolution of Wi-Fi Technology and Standards," May 2023, accessed: April 24, 2025. [Online]. Available: https://standards.ieee.org/beyond-standards/the-evolution-of-wi-fi-technology-and-standards/

[8] E. Reshef and C. Cordeiro, "Future directions for Wi-Fi 8 and beyond," *IEEE Communications Magazine*, vol. 60, no. 10, pp. 50–55, 2022.

[9] L. Galati Giordano *et al.*, "What will Wi-Fi 8 be? A Primer on IEEE 802.11bn Ultra High Reliability," *IEEE Communications Magazine*, vol. 62, no. 8, pp. 126–132, 2024.

[10] X. Liu *et al.*, "Wi-Fi 8: Embracing the Millimeter-Wave Era," *IEEE Communications Magazine*, vol. 63, no. 3, pp. 69–75, 2024.

[11] "IEEE standard Draft 1.1 for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. Amendment 6: Enhancements for ultra high reliablity (UHR)," *IEEE P802.11bn/D1.3*, 2025.

[12] R. Karmakar *et al.*, "Impact of IEEE 802.11 n/ac PHY/MAC high throughput enhancements on transport and application protocols—A survey," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2050–2091, 2017.

[13] B. Chen *et al.*, "Coexistence of LTE-LAA and Wi-Fi on 5 GHz With Corresponding Deployment Scenarios: A Survey," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 1, pp. 7–32, 2017.

[14] B. Bellalta *et al.*, "Next generation IEEE 802.11 Wireless Local Area Networks: Current status, future directions and open challenges," *Computer Communications*, vol. 75, pp. 1–25, 2016.

[15] E. Khorov *et al.*, "A tutorial on IEEE 802.11 ax high efficiency WLANs," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 1, pp. 197–216, 2019.

[16] E. Mozaffariahrar *et al.*, "A survey of Wi-Fi 6: Technologies, advances, and challenges," *Future Internet*, vol. 14, no. 10, p. 293, 2022.

[17] E. Khorov *et al.*, "Current status and directions of IEEE 802.11 be, the future Wi-Fi 7," *IEEE access*, vol. 8, pp. 88 664–88 688, 2020.

[18] C. Deng *et al.*, "IEEE 802.11 be Wi-Fi 7: New challenges and opportunities," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 4, pp. 2136–2166, 2020.

[19] S. Verma *et al.*, "A Survey on Multi-AP Coordination Approaches Over Emerging WLANs," *IEEE Communications Surveys & Tutorials*, vol. 26, no. 2, pp. 858–889, 2024.

[20] S. Szott *et al.*, "Wi-Fi meets ML: A survey on improving IEEE 802.11 performance with machine learning," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 3, pp. 1843–1893, 2022.

[21] E. J. Oughton *et al.*, "The Future of Wireless Broadband in the Peak Smartphone Era: 6G, Wi-Fi 7, and Wi-Fi 8," *IEEE Wireless Communications*, 2025.

[22] D. Lopez-Perez *et al.*, "IEEE 802.11 be extremely high throughput: The next generation of Wi-Fi technology beyond 802.11 ax," *IEEE Communications Magazine*, vol. 57, no. 9, pp. 113–119, 2019.

[23] M. Yang and B. Li, "Survey and perspective on extremely high throughput (EHT) WLAN—IEEE 802.11 be," *Mobile Networks and Applications*, vol. 25, no. 5, pp. 1765–1780, 2020.

[24] A. Garcia-Rodriguez *et al.*, "IEEE 802.11 be: Wi-Fi 7 strikes back," *IEEE Communications Magazine*, vol. 59, no. 4, pp. 102–108, 2021.

[25] X. Liu *et al.*, "IEEE 802.11 be Wi-Fi 7: Feature summary and performance evaluation," *arXiv preprint arXiv:2309.15951*, 2023.

[26] R. Liu and N. Choi, "A first look at Wi-Fi 6 in action: Throughput, latency, energy efficiency, and security," *Proc. of the ACM on Measurement and Analysis of Computing Systems*, vol. 7, no. 1, pp. 1–25, 2023.

[27] S. S. Murad *et al.*, "Introduction to Wi-Fi 7: A Review of History, Applications, Challenges, Economical Impact and Research Development," *Mesopotamian Journal of Computer Science*, vol. 2024, pp. 110–121, 2024.

[28] M. S. Afaqui *et al.*, "IEEE 802.11 ax: Challenges and requirements for future high efficiency WiFi," *IEEE wireless communications*, vol. 24, no. 3, pp. 130–137, 2016.

[29] Q. Qu *et al.*, "Survey and performance evaluation of the upcoming next generation WLANs standard-IEEE 802.11 ax," *Mobile Networks and Applications*, vol. 24, pp. 1461–1474, 2019.

[30] M. Yang *et al.*, "MAC Technology of IEEE 802.11 ax: Progress and Tutorial," *Mobile Networks and Applications*, vol. 26, pp. 1122–1136, 2021.

[31] X. Wang *et al.*, "A survey of LTE Wi-Fi coexistence in unlicensed bands," *GetMobile: Mobile Computing and Communications*, vol. 20, no. 3, pp. 17–23, 2017.

[32] S. Zinno *et al.*, "On a fair coexistence of LTE and Wi-Fi in the unlicensed spectrum: A Survey," *computer communications*, vol. 115, pp. 35–50, 2018.

[33] G. Naik *et al.*, "Next generation Wi-Fi and 5G NR-U in the 6 GHz bands: Opportunities and challenges," *IEEE Access*, vol. 8, pp. 153 027–153 056, 2020.

[34] E. Khorov *et al.*, "A survey on IEEE 802.11 ah: An enabling networking technology for smart cities," *Computer communications*, vol. 58, pp. 53–69, 2015.

[35] M. Meera and S. N. Rao, "A survey of the state of the art of 802.11 ah," in *Proc. of the IEEE International Conference on Computational Intelligence and Computing Research (ICCIC)*. IEEE, 2017.

[36] R. C. Carrano *et al.*, "IEEE 802.11 s multihop MAC: A tutorial," *IEEE Communications Surveys & Tutorials*, vol. 13, no. 1, pp. 52–67, 2010.

[37] P. Zhou *et al.*, "IEEE 802.11 ay-based mmWave WLANs: Design challenges and solutions," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 3, pp. 1654–1681, 2018.

[38] P. Tarafder and W. Choi, "MAC protocols for mmWave communication: A comparative survey," *Sensors*, vol. 22, no. 10, p. 3853, 2022.

[39] Y. Koda *et al.*, "Survey, taxonomy, and unification of standard mmWave channel models for WPAN, WLAN, and cellular systems in 6G," *IEEE Communications Standards Magazine*, vol. 8, no. 3, pp. 44–52, 2024.

[40] A. Khalili *et al.*, "Wi-Fi sensing: applications and challenges," *The Journal of Engineering*, vol. 2020, no. 3, pp. 87–97, 2020.

[41] N. Singh *et al.*, "Machine learning based indoor localization using Wi-Fi RSSI fingerprints: An overview," *IEEE access*, vol. 9, pp. 127 150–127 174, 2021.

[42] C. Chen *et al.*, "Wi-Fi sensing based on IEEE 802.11 bf," *IEEE Communications Magazine*, vol. 61, no. 1, pp. 121–127, 2022.

[43] F. Meneghello *et al.*, "Toward integrated sensing and communications in IEEE 802.11 bf Wi-Fi networks," *IEEE Communications Magazine*, vol. 61, no. 7, pp. 128–133, 2023.

[44] J. Dai *et al.*, "A survey of latest Wi-Fi assisted indoor positioning on different principles," *Sensors*, vol. 23, no. 18, p. 7961, 2023.

[45] R. Du *et al.*, "An overview on IEEE 802.11 bf: WLAN sensing," *IEEE Communications Surveys & Tutorials*, 2024.

[46] F. Wang *et al.*, "A Survey on Wi-Fi Sensing Generalizability: Taxonomy, Techniques, Datasets, and Future Research Prospects," *arXiv preprint arXiv:2503.08008*, 2025.

[47] F. Wilhelmi *et al.*, "Machine Learning and Wi-Fi: Unveiling the Path Toward AI/ML-Native IEEE 802.11 Networks," *IEEE Communications Magazine*, 2024.

[48] S. Barrachina-Muñoz *et al.*, "Wi-Fi channel bonding: An all-channel system and experimental study from urban hotspots to a sold-out stadium," *IEEE/ACM Transactions on Networking*, vol. 29, no. 5, pp. 2101–2114, 2021.

[49] Wi-Fi Alliance, "Regulations Enabling 6 GHz Wi-Fi," 2025. [Online]. Available: https://www.wi-fi.org/regulations-enabling-6-ghz-wi-fi

[50] S. Dogan-Tusha *et al.*, "Evaluation of Indoor/Outdoor Sharing in the Unlicensed 6 GHz Band," *arXiv preprint arXiv:2505.18359*, 2025.

[51] S. Dogan-Tusha *et al.*, "Evaluating the interference potential in 6 GHz: An extensive measurement campaign of a dense indoor Wi-Fi 6E network," in *Proc. of the 17th ACM workshop on wireless network testbeds, experimental evaluation & characterization (WiNTECH)*, 2023.

[52] S. Doğan-Tusha *et al.*, "Spectrum Sharing Characterization Using Smartphones: Exploring 6 GHz Sharing Through Large-Scale Wi-Fi 6E Measurements," *IEEE Communications Magazine*, vol. 63, no. 2, pp. 70–76, 2025.

[53] Wi-Fi Alliance, "Regulations Enabling 6 GHz Standard Power Wi-Fi Devices Under Control of Automated Frequency Coordination (AFC) System," 2025. [Online]. Available: https://www.wi-fi.org/regulations-enabling-6-ghz-standard-power-wi-fi

[54] Wi-Fi Alliance, "Wi-Fi CERTIFIED 7™," https://www.wi-fi.org/wi-fi-certified-7, accessed May 2025.

[55] Wi-Fi Alliance, "Wi-Fi CERTIFIED WPA3™," https://www.wi-fi.org/discover-wi-fi/security, accessed May 2025.

[56] Wi-Fi Alliance, "Wi-Fi Direct™," https://www.wi-fi.org/discover-wi-fi/wi-fi-direct, accessed May 2025.

[57] Wi-Fi Alliance, "Wi-Fi CERTIFIED n." [Online]. Available: https://www.wi-fi.org/discover-wi-fi/wi-fi-certified-n

[58] M. Gast, *802.11n: a survival guide*. " O'Reilly Media, Inc.", 2012.

[59] Wi-Fi Alliance, "Wi-Fi CERTIFIED ac." [Online]. Available: https://www.wi-fi.org/discover-wi-fi/wi-fi-certified-ac

[60] M. Gast, *802.11ac: A survival guide*. " O'Reilly Media, Inc.", 2013.

[61] Wi-Fi Alliance, "Wi-Fi CERTIFIED 6." [Online]. Available: https://www.wi-fi.org/discover-wi-fi/wi-fi-certified-6

[62] S. G. Sankaran and S. R. Gulasekaran, *Wi-Fi 6: Protocol and Network*. Artech House, 2021.

[63] Wi-Fi Alliance, "Wi-Fi Alliance® delivers Wi-Fi 6E certification program." [Online]. Available: https://www.wi-fi.org/news-events/newsroom/wi-fi-alliance-delivers-wi-fi-6e-certification-program

[64] *Broadband Radio Access Networks (BRAN); 5 GHz high performance RLAN; Harmonised Standard for access to radio spectrum*, European Telecommunications Standards Institute (ETSI) Std. EN 301 893, current year, e.g., 2024.

[65] N. Benvenuto *et al.*, *Algorithms for communications systems and their applications*. John Wiley & Sons, 2021.

[66] J. Henry *et al.*, *Wi-Fi 7 in Depth: Your Guide to Mastering Wi-Fi 7, the 802. 11be Protocol, and Their Deployment*. Pearson Education, 2024.

[67] A. Goldsmith, *Wireless communications*. Cambridge Univ. press, 2005.

[68] R. W. Heath Jr and A. Lozano, *Foundations of MIMO communication*. Cambridge University Press, 2018.

[69] E. Björnson *et al.*, "Twenty-five years of signal processing advances for multiantenna communications: From theory to mainstream technology," *IEEE Signal Processing Magazine*, vol. 40, no. 4, pp. 107–117, 2023.

[70] M. S. Gast, *802.11 ac: a survival guide: Wi-Fi at gigabit and beyond*. "O'Reilly Media, Inc.", 2013.

[71] F. Meneghello *et al.*, "DeepCSI: Rethinking Wi-Fi radio fingerprinting through MU-MIMO CSI feedback deep learning," in *Proc. of the IEEE 42nd International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2022.

[72] J. Zhang *et al.*, "Implementation and evaluation of IEEE 802.11 ax channel sounding frame exchange in ns-3," in *Proc. of the ACM Workshop on ns-3 (WNS3)*, 2023.

[73] K. F. Haque *et al.*, "Wi-BFI: Extracting the IEEE 802.11 Beamforming Feedback Information from Commercial Wi-Fi Devices," in *Proc. of the 17th ACM Workshop on Wireless Network Testbeds, Experimental evaluation & Characterization (WiNTECH)*, 2023.

[74] F. Meneghello *et al.*, "Evaluating the Impact of Channel Feedback Quantization and Grouping in IEEE 802.11 MIMO Wi-Fi Networks," *IEEE Wireless Communications Letters*, pp. 1–1, 2024.

[75] C. B. Peel *et al.*, "A vector-perturbation technique for near-capacity multiantenna multiuser communication-part I: channel inversion and regularization," *IEEE Transactions on Communications*, vol. 53, no. 1, pp. 195–202, 2005.

[76] E. Perahia and R. Stacey, *Next Generation Wireless LANs: Throughput, Robustness, and Reliability in 802.11n*. Cambridge Univ. Press, 2008.

[77] F. Meneghello *et al.*, "WHACK: Adversarial Beamforming in MU-MIMO Through Compressed Feedback Poisoning," *IEEE Transactions on Wireless Communications*, vol. 23, no. 11, pp. 17 252–17 265, 2024.

[78] K. Rumman *et al.*, "SHRINK: Reducing MIMO Feedback Overhead in Wi-Fi with Dynamic Data-Driven Channel Sounding," in *Proceedings of the Twenty-sixth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing*, 2025, pp. 11–20.

[79] P. K. Sangdeh *et al.*, "LB-SciFi: Online learning-based channel feedback for MU-MIMO in wireless LANs," in *Proc. of the IEEE 28th International Conference on Network Protocols (ICNP)*. IEEE, 2020.

[80] Z. Hu *et al.*, "A learnable optimization and regularization approach to massive MIMO CSI feedback," *IEEE Transactions on Wireless Communications*, vol. 23, no. 1, pp. 104–116, 2023.

[81] N. Bahadori *et al.*, "SplitBeam: Effective and Efficient Beamforming in Wi-Fi Networks Through Split Computing," in *Proc. of IEEE ICDCS*, Hong Kong, China, 2023.

[82] H. Wu *et al.*, "MIMO Channel as a Neural Function: Implicit Neural Representations for Extreme CSI Compression," in *Proc. of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025.

[83] O. Bejarano *et al.*, "MUTE: Sounding Inhibition for MU-MIMO WLANs," in *Proc. of IEEE International Conference on Sensing, Communication, and Networking (SECON)*. IEEE, 2014, pp. 135–143.

[84] P. K. Sangdeh and H. Zeng, "DeepMux: Deep-learning-based Channel Sounding and Resource Allocation for IEEE 802.11 ax," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 8, pp. 2333–2346, 2021.

[85] J. Mei *et al.*, "Learning Aided Closed-loop Feedback: A Concurrent Dual Channel Information Feedback Mechanism for Wi-Fi," *IEEE Transactions on Wireless Communications*, vol. 13, no. 12, pp. 19 275–19 288, 2024.

[86] (2020) Discussion on EHT PPDU Formats. [Online]. Available: https://mentor.ieee.org/802.11/dcn/20/11-20-0031-02-00be-considerations-on-eht-ppdu-formats.pptx

[87] (2024) Enhanced Long Range (ELR) PPDU Design. [Online]. Available: https://mentor.ieee.org/802.11/dcn/24/11-24-1478-00-00bn-elr-ppdu-design.pptx

[88] (2024) Design Targets and Considerations for Enhanced Long Range. [Online]. Available: https://mentor.ieee.org/802.11/dcn/24/11-24-0873-00-00bn-design-targets-and-considerations-for-enhanced-long-range.pptx

[89] (2024) Specification Framework for TGbn. [Online]. Available: https://mentor.ieee.org/802.11/dcn/24/11-24-0209-12-00bn-specification-framework-for-tgbn.docx

[90] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE Journal on selected areas in communications*, vol. 18, no. 3, pp. 535–547, 2000.

[91] L. Sanabria-Russo and B. Bellalta, "Traffic differentiation in dense collision-free WLANs using CSMA/ECA," *Ad Hoc Networks*, vol. 75, pp. 33–51, 2018.

[92] (2024) Low latency channel access follow up. [Online]. Available: https://mentor.ieee.org/802.11/dcn/24/11-24-1144-01-00bn-hip-edca-proposal-follow.pptx

[93] D. Skordoulis *et al.*, "IEEE 802.11n MAC frame aggregation mechanisms for next-generation high-throughput WLANs," *IEEE Wireless Communications*, vol. 15, no. 1, pp. 40–47, 2008.

[94] B. Bellalta and K. Kosek-Szott, "AP-initiated multi-user transmissions in IEEE 802.11ax WLANs," *Ad Hoc Networks*, vol. 85, pp. 145–159, 2019.

[95] S. Barrachina-Munoz *et al.*, "Dynamic channel bonding in spatially distributed high-density WLANs," *IEEE Transactions on Mobile Computing*, vol. 19, no. 4, pp. 821–835, 2019.

[96] C. Michaelides *et al.*, "Lessons Learned From a Large-Scale Virtual Reality Experience Over Wi-Fi," *IEEE Wirless Communications*, 2025.

[97] "IEEE Approved Draft Standard for Information technology–Telecommunications and information exchange between systems Local and metropolitan area networks–Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment: Enhancements for Extremely High Throughput (EHT)," *IEEE P802.11be/D7.0, August 2024*, pp. 1–1089, 2024.

[98] B. Bellalta *et al.*, "Performance Analysis of IEEE 802.11bn Non-Primary Channel Access," *arXiv preprint arXiv:2504.15774*, 2025.

[99] F. Wilhelmi *et al.*, "Spatial reuse in IEEE 802.11ax WLANs," *Computer Communications*, vol. 170, pp. 65–83, 2021.

[100] E. de Carvalho Rodrigues *et al.*, "On the latency of IEEE 802.11ax WLANs with parameterized spatial reuse," in *Proc. of the IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2020.

[101] H. Yin and S. Alamouti, "OFDMA: A broadband wireless access technology," in *Proc. of the IEEE sarnoff symposium*. IEEE, 2006.

[102] S. Avallone *et al.*, "Will OFDMA improve the performance of 802.11 WiFi networks?" *IEEE Wireless Communications*, vol. 28, no. 3, pp. 100–107, 2021.

[103] D. Magrin *et al.*, "Performance evaluation of 802.11 ax ofdma through theoretical analysis and simulations," *IEEE Transactions on Wireless Communications*, vol. 22, no. 8, pp. 5070–5083, 2023.

[104] T. Hou *et al.*, "MUSTER: Subverting user selection in MU-MIMO networks," in *Proc. of the IEEE International Conference on Computer Communications (INFOCOM)*, 2022.

[105] F. Meneghello *et al.*, "How to BREAK MU-MIMO Precoding in IEEE 802.11 Wi-Fi Networks," in *Proc. of the IEEE International Conference on Computer Communications (INFOCOM)*, 2025.

[106] Y.-C. Tung *et al.*, "Vulnerability and Protection of Channel State Information in Multiuser MIMO Networks," in *Proc. of ACM SIGSAC Conference on Computer and Communications Security*, New York, NY, USA, 2014.

[107] X. Wang *et al.*, "On eavesdropping attacks and countermeasures for MU-MIMO systems," in *Proc. of the IEEE Military Communications Conference (MILCOM)*, 2017.

[108] S. Wang *et al.*, "On user selective eavesdropping attacks in MU-MIMO: CSI forgery and countermeasure," in *Proc. of the IEEE International Conference on Computer Communications (INFOCOM)*, 2019.

[109] Y. Yang *et al.*, "Securing channel state information in multiuser MIMO with limited feedback," *IEEE Transactions on Wireless Communications*, vol. 19, no. 5, pp. 3091–3103, 2020.

[110] G. Naik *et al.*, "Performance analysis of uplink multi-user OFDMA in IEEE 802.11 ax," in *Proc. of the IEEE international conference on communications (ICC)*. IEEE, 2018.

[111] D.-J. Deng *et al.*, "IEEE 802.11 ax: Highly efficient WLANs for intelligent information infrastructure," *IEEE Communications Magazine*, vol. 55, no. 12, pp. 52–59, 2017.

[112] (2022) Follow up on the low power listening mode. [Online]. Available: https://mentor.ieee.org/802.11/dcn/22/11-22-1841-00-0uhr-follow-up-on-the-low-power-listening.pptx

[113] E. Mozaffariahrar *et al.*, "R-TWT in Wi-Fi 7 and Beyond: Enabling Bounded Latency, Energy Efficiency, and Reliability," in *ETFA 2025*. IEEE, 2025.

[114] M. Nurchis and B. Bellalta, "Target wake time: Scheduled access in IEEE 802.11 ax WLANs," *IEEE Wireless Communications*, vol. 26, no. 2, pp. 142–150, 2019.

[115] R. Sanchez-Vital *et al.*, "A primer on AP power save in Wi-Fi 8: overview, analysis, and open challenges," *IEEE Wireless Communications*, 2025.

[116] E. Guérin *et al.*, "An overview of MAC energy-saving mechanisms in Wi-Fi," *Computer Communications*, vol. 203, pp. 129–145, 2023.

[117] (2025) Follow-up on AP Power Save. [Online]. Available: https://mentor.ieee.org/802.11/dcn/25/11-25-0749-00-00bn-follow-up-on-ap-power-save.pptx

[118] (2024) A Proposal For UHR Mobile-AP Power Save. [Online]. Available: https://mentor.ieee.org/802.11/dcn/24/11-24-0450-03-00bn-a-proposal-for-uhr-soft-ap-power-save.pptx

[119] (2025) TWT-Based AP Power Save. [Online]. Available: https://mentor.ieee.org/802.11/dcn/25/11-25-0578-00-00bn-twt-based-ap-power-save.pptx

[120] M. Carrascosa-Zamacois *et al.*, "Understanding Multi-link Operation in Wi-Fi 7: Performance, Anomalies, and Solutions," in *Proc. of the IEEE 34th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2023.

[121] M. Carrascosa-Zamacois *et al.*, "Performance and Coexistence Evaluation of IEEE 802.11be Multi-link Operation," in *Proc. of the IEEE Wireless Communications and Networking Conference (WCNC)*, 2023.

[122] G. Lacalle *et al.*, "Analysis of Latency and Reliability Improvement with Multi-Link Operation over 802.11," in *Proc. of the IEEE 19th International Conference on Industrial Informatics (INDIN)*, 2021.

[123] B. Bellalta *et al.*, "Delay Analysis of IEEE 802.11be Multi-Link Operation Under Finite Load," *IEEE Wireless Communications Letters*, vol. 12, no. 4, pp. 595–599, 2023.

[124] M. Carrascosa-Zamacois *et al.*, "Performance Evaluation of MLO for XR Streaming: Can Wi-Fi 7 Meet the Expectations?" in *Proc. of the IEEE 29th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)*, 2024.

[125] N. Korolev *et al.*, "Analytical Model of Multi-Link Operation in Saturated Heterogeneous Wi-Fi 7 Networks," *IEEE Wireless Communications Letters*, vol. 11, no. 12, pp. 2546–2549, 2022.

[126] M. Carrascosa-Zamacois *et al.*, "Wi-Fi multi-link operation: An experimental study of latency and throughput," *IEEE/ACM Transactions on Networking*, vol. 32, no. 1, pp. 308–322, 2024.

[127] B. Bellalta *et al.*, "Delay analysis of IEEE 802.11be multi-link operation under finite load," *IEEE Wireless Communications Letters*, vol. 12, no. 4, pp. 595–599, 2023.

[128] "IEEE P802.11bn Project Authorization Request (PAR)," https://development.standards.ieee.org/myproject-web/public/view.html#pardetail/10639, 2023, accessed on February 6, 2024.

[129] I. Val *et al.*, "Wi-Fi 8 Unveiled: Key Features, Multi-AP Coordination, and the Role of C-TDMA," 2025.

[130] J. Haxhibeqiri *et al.*, "Coordinated SR and restricted TWT for time sensitive applications in WiFi 7 networks," *IEEE Communications Magazine*, vol. 62, no. 8, pp. 118–124, 2024.

[131] F. Wilhelmi *et al.*, "Throughput analysis of IEEE 802.11 bn coordinated spatial reuse," in *Proc. of the IEEE Conference on Standards for Communications and Networking (CSCN)*. IEEE, 2023.

[132] L.-H. Shen *et al.*, "Coordinated multiple access point multiuser beamforming training protocol for millimeter wave WLANs," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 11, pp. 13 875–13 889, 2020.

[133] G. Geraci *et al.*, "Operating massive MIMO in unlicensed bands for enhanced coexistence and spatial reuse," *IEEE Journal on Selected Areas in Communications*, vol. 35, no. 6, pp. 1282–1293, 2017.

[134] A. Garcia-Rodriguez *et al.*, "Massive MIMO unlicensed: A new approach to dynamic spectrum access," *IEEE Communications Magazine*, vol. 56, no. 6, pp. 186–192, 2018.

[135] G. Geraci *et al.*, "Indoor massive MIMO deployments for uniformly high wireless capacity," in *Proc. IEEE WCNC*, 2018, pp. 1–6.

[136] F. Adib and D. Katabi, "See through walls with WiFi!" in *Proc. of the ACM SIGCOMM conference on SIGCOMM*, 2013.

[137] F. Meneghello *et al.*, "SHARP: Environment and Person Independent Activity Recognition with Commodity IEEE 802.11 Access Points," *IEEE Transactions on Mobile Computing*, vol. 22, no. 10, pp. 6160–6175, 2022.

[138] J. Pegoraro *et al.*, "RAPID: Retrofitting IEEE 802.11 ay access points for indoor human detection and sensing," *IEEE Transactions on Mobile Computing*, vol. 23, no. 5, pp. 4501–4519, 2023.

[139] X. Wang *et al.*, "On CSI-based vital sign monitoring using commodity WiFi," *ACM Transactions on Computing for Healthcare*, vol. 1, no. 3, pp. 1–27, 2020.

[140] Y. Ma *et al.*, "WiFi Sensing with Channel State Information: A Survey," *ACM Computing Surveys*, vol. 52, no. 3, 2019.

[141] F. Liu *et al.*, "Integrated Sensing and Communications: Toward Dual-Functional Wireless Networks for 6G and Beyond," *IEEE Journal on Selected Areas in Communications*, vol. 40, no. 6, 2022.

[142] K. F. Haque *et al.*, "Beamsense: Rethinking Wireless Sensing with MU-MIMO Wi-Fi Beamforming Feedback," *Computer Networks*, vol. 258, p. 111020, 2025.

[143] F. Gringoli *et al.*, "Free Your CSI: A Channel State Information Extraction Platform For Modern Wi-Fi Chipsets," in *Proc. of the 13th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (WiNTECH)*. Association for Computing Machinery, 2019.

[144] F. Gringoli *et al.*, "AX-CSI: Enabling CSI Extraction on Commercial 802.11ax Wi-Fi Platforms," in *Proc. of the 15th ACM Workshop on Wireless Network Testbeds, Experimental Evaluation & CHaracterization (WiNTECH)*. Association for Computing Machinery, 2022.

[145] Z. Jiang *et al.*, "Eliminating the Barriers: Demystifying Wi-Fi Baseband Design and Introducing the PicoScenes Wi-Fi Sensing Platform," *IEEE Internet of Things Journal*, vol. 9, no. 6, pp. 4476–4496, 2022.

[146] K. F. Haque *et al.*, "Wi-BFI: Extracting the IEEE 802.11 Beamforming Feedback Information from Commercial Wi-Fi Devices," in *Proc. of the 17th ACM Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (WiNTECH)*. New York, NY, USA: Association for Computing Machinery, 2023.

[147] D. Ficara *et al.*, "A tutorial on privacy, RCM and its implications in WLAN," *IEEE Communications Surveys & Tutorials*, vol. 26, no. 2, pp. 1003–1040, 2024.

[148] M. Vanhoef and et al., "Why MAC address randomization is not enough: an analysis of Wi-Fi network discovery mechanisms," *ASIA CCS*, 2016.

[149] X. Wang *et al.*, "IEEE 802.11-22/0987r25, AIML TIG Technical Report Draft," 2023, accessed on May 28, 2025. [Online]. Available: https://mentor.ieee.org/802.11/dcn/22/11-22-0987-25-aiml-aiml-tig-technical-report-draft.doc

[150] M. Deshmukh *et al.*, "Intelligent feedback overhead reduction (iFOR) in Wi-Fi 7 and beyond," in *Proc. of the IEEE 95th Vehicular Technology Conference:(VTC2022-Spring)*. IEEE, 2022.

[151] W. Wydmański and S. Szott, "Contention window optimization in IEEE 802.11 ax networks with deep reinforcement learning," in *Proc. of the IEEE wireless communications and networking conference (WCNC)*. IEEE, 2021.

[152] Z. Guo *et al.*, "Multi-agent reinforcement learning-based distributed channel access for next generation wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 40, no. 5, pp. 1587–1599, 2022.

[153] F. Wilhelmi *et al.*, "A flexible machine-learning-aware architecture for future WLANs," *IEEE Communications Magazine*, vol. 58, no. 3, pp. 25–31, 2020.