RESEARCH ARTICLE OPEN ACCESS

Wi-Fi 7: The Next Generation of Wireless Networking—Technologies, Performance, Applications, Challenges, and Future Directions

Alaa H Ridha[1], Abdullah A Almaatoooq[2]
Computer Department, Higher Institute of Telecommunication and Navigation
Email: hitn2014@gmail.com
Computer Department, Higher Institute of Telecommunication and Navigation
Email: imatooq79@gmail.com

ABSTRACT
Wi-Fi 7, specified by IEEE 802.11be as Extremely High Throughput, is usually introduced through its headline data rate. That shorthand hides the change that matters more: several radio links can now be coordinated within a single association. This survey works from the completed amendment rather than from draft-era projections, and reviews the evolution of the standard, its architecture, the principal mechanisms, reported performance, applications, deployment constraints, and the research outlook. We cover multi-link operation, 320-MHz channels, 4096-QAM, multi-resource-unit allocation, EHT OFDMA and MU-MIMO, preamble puncturing, and the actual status of multi-access-point coordination. An ideal eight-stream link at 320 MHz reaches approximately 23.06 Gbit/s. What a deployment delivers depends instead on the client radio, interference, contention, regulatory limits, and scheduling policy. The available studies point to real gains in capacity, delay distribution, and resilience, but not to a universal improvement under every load or topology, so we keep standardized capability separate from theoretical limits, experimental findings, and vendor-specific implementation. Dependable low-latency WLANs will still require cross-layer scheduling, energy-aware link control, secure automation, careful spectrum planning, and closer integration with sensing, edge computing, 5G, and future 6G systems.

***Keywords* -** IEEE 802.11be, multi-link operation, spectrum efficiency, Wi-Fi 7, wireless LAN

Date of Submission: 09-09-2026 Date of acceptance: 18-09-2026

I. INTRODUCTION

Today's WLANs carry a traffic mix that earlier generations were never dimensioned for. Ultra-high-definition video now shares airtime with cloud-rendered extended reality, industrial control, distributed analytics, and hundreds of enterprise clients at once. In that setting a peak-rate figure says very little about what a user experiences; tail latency, jitter, reliability, coexistence, and energy consumption matter just as much. IEEE Std 802.11be-2024, the Extremely High Throughput (EHT) amendment, responds to this wider problem with operation below 7.25 GHz, channel bandwidths up to 320 MHz, higher-order modulation, and coordinated access over several links. The amendment was approved in September 2024 and published in July 2025 [1]. Wi-Fi Alliance certification had already started in January 2024, based on an interoperable subset of the feature set then emerging [2]. A certified product, the certification program, and the full IEEE amendment are three different things, and conflating them is the source of much of the confusion around Wi-Fi 7.

The distinction matters most when reading the early EHT targets. The task group first aimed at 30 Gbit/s or more at the MAC service access point, together with markedly better worst-case delay and jitter [3]–[6]. Draft-era papers and product marketing went further and projected a 46-Gbit/s single-link rate by combining 320 MHz, 4096-QAM, and sixteen spatial streams. The completed amendment does not allow that configuration. It stops at eight spatial streams, which puts the ideal single-link maximum near 23.06 Gbit/s over 320 MHz. An implementation may advertise a larger aggregate by operating independent links concurrently, but that is a different quantity altogether. How much of either ceiling is reachable depends on radio count, simultaneous-link capability, available channels, guard interval, coding rate, and the design of the client.

This survey therefore works from the final standard and grades its claims by the kind of evidence behind them. Standards and regulatory decisions say what is permitted. Analytical, simulation, and experimental studies show how the mechanisms behave. Vendor documents expose the constraints of shipping radios. We consider literature available through August 2026. The same discipline avoids a category error that recurs in the Wi-Fi 7 literature: broad multi-AP coordination was central to early TGbe research but was never completed as a general 802.11be capability, so we treat it as a forward research direction rather than a feature.

Section 2 traces the path from Wi-Fi 4 to Wi-Fi 7. Sections 3 and 4 examine the EHT architecture and its principal technologies. Section 5 compares Wi-Fi 6, Wi-Fi 6E, and Wi-Fi 7 on common terms and explains why measured results fall short of theoretical limits. Sections 6 and 7 turn to use cases and deployment barriers. Section 8 sets out a research agenda, and Section 9 draws the main conclusions.

Fig. 1. Evolution from Wi-Fi 4 to Wi-Fi 7: successive gains in spatial multiplexing, multi-user efficiency, 6-GHz access, and multi-link coordination.

II. EVOLUTION OF WI-FI STANDARDS

2.1 Wi-Fi 4

IEEE 802.11n, marketed as Wi-Fi 4, is where MIMO became a practical foundation of WLAN design. Channel bonding to 40 MHz, up to four spatial streams, 64-QAM, frame aggregation, and block acknowledgement lifted the PHY ceiling to 600 Mbit/s [7]. The number mattered less than the achievement behind it: spatial multiplexing and substantial MAC aggregation were folded into a backward-compatible system. One limitation showed up quickly in crowded networks, since a transmission opportunity still normally served a single station, so contention and per-frame overhead grew with the client population.

2.2 Wi-Fi 5

Wi-Fi 5, built on IEEE 802.11ac, chased very high throughput in the 5-GHz band. Channels widened to 80 or 160 MHz, 256-QAM arrived, and eight spatial streams were allowed, giving a theoretical maximum of about 6.93 Gbit/s [8]. Downlink MU-MIMO also let an access point (AP) serve several clients in one transmission, though sounding overhead, dissimilar client capabilities, and scheduler complexity trimmed that gain in practice. This generation also surfaced a planning tradeoff that Wi-Fi 7 inherits: widening a channel raises the rate of one cell while leaving fewer channels for spatial reuse across many cells.

2.3 Wi-Fi 6

IEEE 802.11ax moved the target from the fastest link to the efficiency of a dense basic service set. Under the commercial name Wi-Fi 6, OFDMA divides a channel into resource units (RUs) so that several stations can transmit or receive within the same scheduled opportunity. Uplink and downlink MU-MIMO, 1024-QAM, BSS coloring for spatial reuse, and target wake time complete the design [9], [10]. The maximum PHY rate is roughly 9.61 Gbit/s with 160 MHz and eight streams, but the operationally important gain lies elsewhere: the AP can combine many heterogeneous short packets without a fresh contention cycle for every client.

2.4 Wi-Fi 6E

Wi-Fi 6E is 802.11ax operating in newly available 6-GHz spectrum, not another IEEE PHY generation, and the Wi-Fi Alliance certification program dates from 2021 [11]. The United States opened 1,200 MHz from 5.925 to 7.125 GHz to unlicensed devices under specific power classes and incumbent-protection rules [12]; the European Union initially harmonized 480 MHz from 5.945 to 6.425 GHz [13]. That spectrum holds more 80- and 160-MHz channels and carries no pre-Wi-Fi 6 clients. The PHY limits, however, remain those of Wi-Fi 6: 160 MHz, 1024-QAM, and eight streams. Whatever advantage Wi-Fi 6E delivers comes from a cleaner and more spacious radio environment.

2.5 Wi-Fi 7

Wi-Fi 7 pairs the EHT PHY extensions with a deeper change in the MAC: multi-link operation (MLO). IEEE 802.11be introduces 320-MHz transmissions in 6 GHz, 4096-QAM, permitted allocations of several RUs to one station, more flexible puncturing, and multi-link devices that coordinate affiliated interfaces [1], [14], [15]. These mechanisms do not deliver one fixed kind of gain. A deployment can spend them on the peak rate of a single link, on pooled capacity across links, or on picking whichever link is expected to clear first. Choosing among those aims while respecting reliability, coexistence, and energy is what makes Wi-Fi 7 a coordination problem as much as a faster PHY.

Table 1. Evolution of principal Wi-Fi capabilities

Generation	IEEE basis	Bands	Max. channel	Max. modulation	Max. spatial streams	Nominal peak PHY
Wi-Fi 4	802.11n	2.4/5 GHz	40 MHz	64-QAM	4	0.60 Gbit/s
Wi-Fi 5	802.11ac	5 GHz	160 MHz	256-QAM	8	6.93 Gbit/s
Wi-Fi 6	802.11ax	2.4/5 GHz	160 MHz	1024-QAM	8	9.61 Gbit/s
Wi-Fi 6E	802.11ax in 6 GHz	6 GHz plus legacy bands	160 MHz	1024-QAM	8	9.61 Gbit/s
Wi-Fi 7	802.11be	2.4/5/6 GHz	320 MHz	4096-QAM	8	23.06 Gbit/s per link

III. WI-FI 7 ARCHITECTURE

The organizing abstraction in 802.11be is the multi-link device (MLD). An AP MLD is made up of two or more affiliated APs; a non-AP MLD is made up of affiliated stations. Higher protocol layers see one logical device and one association. Underneath, each affiliate keeps its own channel, PHY state, contention process, and band-specific constraints [1], [16], while a shared upper-MAC context handles capabilities, traffic identifiers, sequence handling, security state, and link choice. Applications can then draw on 2.4, 5, and 6 GHz without juggling several unrelated connections.

Multi-link discovery and setup are built so that the same exchange need not be repeated on every radio. An enabled link can advertise information about its peers; the devices then negotiate link identifiers and operating capabilities, including simultaneous transmit and receive. An STR MLD runs independent radios concurrently. A non-STR implementation has to suppress link combinations that would create harmful in-device interference. EMLSR solves the problem differently: the device observes more than one link but moves a main radio to whichever link is selected. That saves components and power at the cost of switching delay, and it rules out continuous aggregation of independent payloads [16], [22].

Interoperability rests on keeping logical identity separate from radio implementation. A tri-band AP may offer three links to a two-link client, and an EMLSR handset and an STR workstation can associate with the same AP MLD while receiving very different service. Older stations still connect to an affiliated AP as single-link devices, so backward compatibility holds, although legacy protection and narrower transmissions consume airtime that newer clients would otherwise use.

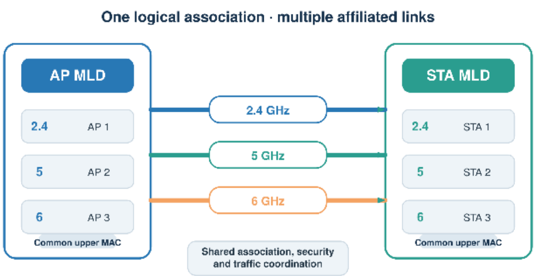

Fig. 2. Wi-Fi 7 multi-link architecture: one logical association coordinates affiliated AP and station interfaces across independent links.

An MLD is only one piece of a deployment. Multi-gigabit Ethernet, Power over Ethernet, edge services, policy control, and Internet or private-cellular gateways can each become the bottleneck, or a useful place to coordinate. A 23-Gbit/s PHY does nothing for a flow that leaves the AP over a 1-Gbit/s uplink, and a short radio-access delay cannot rescue an overloaded rendering or analytics service. Wi-Fi 7 engineering has to look past the access link.

IV. CORE TECHNOLOGIES

4.1 Multi-Link Operation (MLO)

MLO lets a pair of MLDs coordinate several affiliated links under one association [1], [16]. What the scheduler does with them follows from the service objective. Concurrent transmission aggregates throughput when both the radios and the backhaul have headroom. A latency-oriented policy instead sends a packet on the link expected to become available first, buying an extra contention opportunity without necessarily transmitting in parallel. Duplicating selected critical frames across sufficiently independent links lowers the probability of missing a deadline, at the price of airtime. Real deployments will most likely mix these behaviors by traffic class rather than settle on one permanent mode.

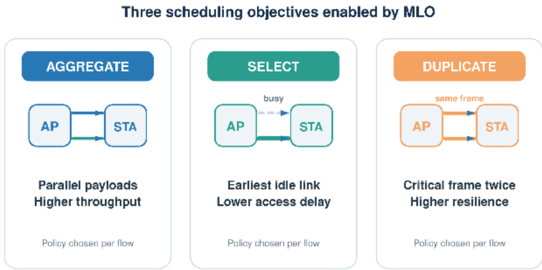

Fig. 3. MLO operating objectives: concurrent aggregation, earliest-link selection, and selective duplication for resilience.

Link count is a poor predictor of MLO gain; the joint occupancy distribution across links tells you far more. Dynamic allocation relieves head-of-line blocking by steering packets away from a long queue [17]. In trace-driven experiments, link selection cut high-percentile access delay by close to an order of magnitude when the candidate links carried similar occupancy, and the same work showed that a poor mapping under asymmetric congestion erases most of that benefit [18]. Standard-compliant evaluations put multi-radio and single-radio modes at different fairness, latency, and throughput operating points [19]. Finite-load analysis agrees: independent contention opportunities compress the delay distribution, and the advantage narrows as every link approaches saturation [20]. Simulations with real-time traffic report similar reductions, without establishing any deterministic guarantee [21].

Radio architecture draws another boundary. STR offers the most room for aggregation but needs several active RF chains and enough isolation between simultaneous transmissions. EMLSR suits mobile equipment, where a primary radio can retune after an initial exchange, yet switching time, unequal link conditions, and the inability to carry two payloads at once all limit it [22]. The open question is no longer whether to use every link. It is when a given link is worth activating, given the packet deadline, queue age, channel state, energy budget, and self-interference.

4.2 320-MHz Channels

Going from 160 to 320 MHz roughly doubles the available data subcarriers and, with modulation and stream count held constant, the attainable PHY rate. EHT defines the mode in 6 GHz with channelization variants suited to different regional allocations [1]. A close-in client with high signal-to-noise ratio exploits the extra bandwidth well. A dense enterprise faces the opposite optimization: the spectrum handed to one 320-MHz cell could instead support several narrower cells that reuse more easily. Network-wide capacity can therefore be higher at 80 or 160 MHz even when each client records a lower peak rate.

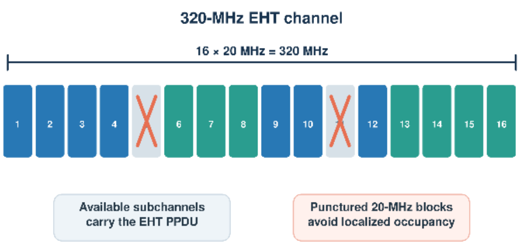

Fig. 4. A 320-MHz EHT channel composed of sixteen 20-MHz subchannels; puncturing preserves usable spectrum when selected subchannels are occupied.

Regulation shapes 320 MHz as much as physics does. The United States allocation can in principle hold three non-overlapping 320-MHz blocks; the initial European lower-6-GHz allocation is wide enough for one, depending on the adopted channel plan and operating class [12], [13]. Power-spectral-density limits, indoor-only device classes, automated frequency coordination for standard-power operation, and incumbent services further constrain coverage and channel availability. The figure "320 MHz" describes a standardized capability, not a uniform worldwide planning rule.

4.3 4096-QAM

Where 1024-QAM conveys 10 bits per modulation symbol, 4096-QAM conveys 12. Holding coding rate and resource allocation constant, that is a raw modulation gain of 20%, and EHT defines the corresponding high modulation and coding schemes for favorable channel conditions [1], [14]. Packing the constellation points closer together raises the required signal-to-noise ratio and tightens error-vector-magnitude limits. Mobility, walls, adjacent-channel leakage, RF nonlinearity, and imperfect channel estimation all shrink the region in which the highest MCS stays stable. 4096-QAM is a short-range peak-rate mechanism, and reading it as a coverage or reliability improvement is a mistake.

4.4 Multi-RU Allocation

Under 802.11ax OFDMA a station normally receives a single RU within a physical protocol data unit. Wi-Fi 7 permits defined combinations of several RUs for one station, so a scheduler can work around interference or around allocations that have already fragmented the channel [1], [14]. This helps when one RU cannot serve a high-rate queue and no sufficiently large contiguous RU is left. It also reduces stranded tones by fitting allocations more closely to heterogeneous demand.

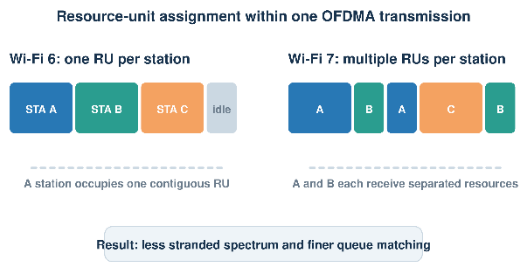

Fig. 5. Multi-RU allocation: Wi-Fi 6 assigns one RU per station, whereas Wi-Fi 7 can combine permitted separated RUs for the same station.

The flexibility comes with a harder control problem. Signaling and scheduler complexity both grow, and separated RUs can see markedly different interference and channel quality. A workable policy has to judge whether reclaiming fragmented spectrum is worth the extra signaling, coding, and receiver power. Timely queue and channel measurements are what make that judgment possible.

4.5 Enhanced OFDMA

EHT extends 802.11ax OFDMA rather than replacing it. Wider channels, additional RU combinations, and revised signaling give the AP more room to serve heterogeneous queues in fragmented spectrum [1], [14]. The underlying benefit is statistical multiplexing: a short control message, a sensor report, a voice frame, and part of a bulk transfer can share one transmission opportunity instead of queueing behind one another. A 320-MHz tone plan multiplies the packing options, and with them the scheduler's state space. Finding low-complexity allocations that still respect deadlines, and coordinating those decisions with MLO, remains an open problem.

4.6 Enhanced MU-MIMO

In Wi-Fi 7, MU-MIMO is combined with EHT signaling, wider bandwidth, and OFDMA/multi-RU scheduling. Early proposals reached for sixteen streams; the completed PHY retains a maximum of eight [1]. The enhancement is not a doubled stream ceiling. Its value lies in how frequency- and spatial-domain multiplexing combine with sounding and the EHT transmission formats. Outcomes still hinge on channel orthogonality, the age of feedback, mobility, antenna placement, and the number of RF chains the AP actually implements. An eight-stream label does not promise eight independent high-rate streams when the clients or the propagation channel cannot produce a well-conditioned MU-MIMO matrix.

4.7 Multi-AP Coordination

Coordination among neighboring APs was prominent in the TGbe research program. Proposals ranged from coordinated scheduling and spatial reuse to joint transmission and distributed MU-MIMO [6], [24], [25], all of them aimed at replacing unmanaged inter-cell contention with deliberate resource sharing. The technical burden climbs steeply as coordination moves from exchanging schedules toward coherent transmission, since synchronization, timely channel-state exchange, backhaul latency, calibration, and protocol overhead all become demanding.

None of these broad multi-AP modes was completed as a normative 802.11be function, a point routinely lost in abbreviated Wi-Fi 7 feature lists. The TGbe work reads better as a technical bridge to IEEE 802.11bn/Wi-Fi 8, where ultra-high reliability and managed multi-AP operation have returned to prominence. Current Wi-Fi 7 controllers do coordinate channel assignment, transmit power, roaming, and load, but network-level management is not a standardized coherent multi-AP PHY.

4.8 Preamble Puncturing

Conventional channel bonding leaves a wide transmission at the mercy of a narrow interferer, because activity in one 20-MHz subchannel can block a much larger bonded channel. Preamble puncturing marks the subchannels to omit and allows transmission on the remainder, and IEEE 802.11be widens the set of puncturing patterns available for the larger EHT channels [1], [23], as illustrated in Fig. 4. The mechanism earns its place when interference is localized and persistent. Waiting is sometimes the better option: heavy fragmentation cuts coding resources and complicates the schedule, and a brief collision may clear before a punctured transmission would have paid off.

V. PERFORMANCE ANALYSIS

Comparing Wi-Fi generations only means something when bandwidth, stream count, coding, and guard interval are stated. The PHY rate is not the rate an application sees. MAC headers, acknowledgements, contention, aggregation efficiency, errors, retransmissions, and transport behavior all take their share of goodput. For the idealized single-link comparison in Table 2, we compute peak spectral efficiency as nominal PHY rate divided by occupied bandwidth. Because spatial multiplexing is included, this is an eight-stream system ceiling rather than per-antenna efficiency.

Table 2. Wi-Fi 6, Wi-Fi 6E, and final Wi-Fi 7 capability comparison

Metric	Wi-Fi 6	Wi-Fi 6E	Wi-Fi 7
IEEE PHY/MAC basis	802.11ax	802.11ax in 6 GHz	802.11be EHT
Maximum channel width	160 MHz	160 MHz	320 MHz (6 GHz)
Highest modulation	1024-QAM	1024-QAM	4096-QAM
Maximum spatial streams	8	8	8
Maximum single-link PHY rate	9.608 Gbit/s	9.608 Gbit/s	23.059 Gbit/s
Representative 2×2 peak	2.402 Gbit/s	2.402 Gbit/s	5.765 Gbit/s
Peak PHY spectral efficiency*	60.0 bit/s/Hz	60.0 bit/s/Hz	72.1 bit/s/Hz
Multi-link operation	No	No	Yes
RU allocation per station	One RU	One RU	Permitted multi-RU combinations

Peak PHY rate divided by maximum channel width, using eight spatial streams and the shortest standard guard interval; this is not measured goodput.

The rise from 9.608 to 23.059 Gbit/s, a factor of about 2.4, comes mostly from doubling the channel width, with another 20% from 4096-QAM; exact tone plans and symbol parameters account for the small remainder. Stream count contributes nothing, since both generations stop at eight. Wi-Fi 6E shows why this arithmetic must not be read as deployed performance: its ceiling matches Wi-Fi 6 exactly, yet access to less congested 6-GHz spectrum often makes wider channels and higher practical throughput far easier to sustain.

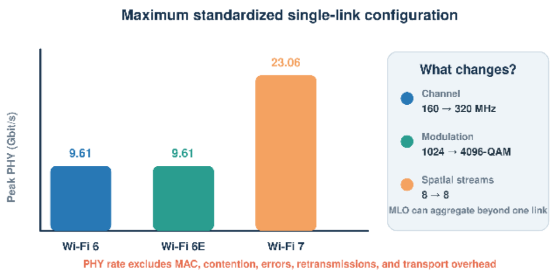

Fig. 6. Theoretical single-link comparison under maximum standardized configurations; bars show peak PHY rate, not user goodput.

Table 3. Critical interpretation of performance dimensions

Dimension	Wi-Fi 6/6E baseline	Wi-Fi 7 mechanism	Evidence-qualified interpretation
Throughput	One 160-MHz link; 1024-QAM	320 MHz, 4096-QAM, MLO	Up to 23.06 Gbit/s on one ideal link; aggregation requires multiple active radios and adequate backhaul.
Latency	OFDMA and scheduled access reduce contention	Earliest-link selection, parallel contention, flexible RUs	Often lower, especially in upper percentiles, but gains collapse under correlated saturation or poor link mapping [18]–[21].
Spectrum efficiency	OFDMA, MU-MIMO, spatial reuse	Multi-RU, puncturing, 4096-QAM	Peak PHY efficiency rises about 20%; network efficiency depends on reuse, scheduler overhead, and interference.

Dimension	Wi-Fi 6/6E baseline	Wi-Fi 7 mechanism	Evidence-qualified interpretation
Reliability	Retransmission and rate adaptation on one link	Link diversity and selective duplication	Potentially higher when failures are weakly correlated; no universal packet-delivery guarantee.
Spatial streams	Up to 8	Up to 8	Final Wi-Fi 7 does not standardize the 16-stream draft-era proposal.

Latency is the metric most often oversold. Several links give several chances to gain access, and the minimum of their access times can have a much shorter upper tail than the access time of any one link. That independence disappears when the links see the same interferer, share a radio, or saturate together. Reordering, retuning, power-save transitions, and uneven queues add delays of their own. MLO gives a scheduler more ways to meet a deadline; it does not convert a contention-based unlicensed channel into deterministic transport.

Reliability is conditional for the same reasons. A scheduler can route around an impaired link, and selective duplication cuts loss when the failure processes are weakly correlated. Neither technique helps against a common backhaul or power failure, correlated interference, or a defect shared by both radios. Demanding services should pair link diversity with admission control, bounded queues, path diversity beyond the AP, and application-level recovery.

Shipping devices occupy a much smaller design space than the theoretical maxima suggest. Cisco notes that 320-MHz operation changes RF planning and wired-uplink requirements together [26]. Intel's client portfolio shows the gap between two-stream configurations at 160 and at 320 MHz [27]. Qualcomm specifies FastConnect 7800 as a 5.8-Gbit/s dual-radio/high-band design rather than an eight-stream client [28], and Broadcom infrastructure chipsets use product-specific 4×4 or dual-2×2 arrangements [29]. These are not shortcomings but engineering compromises among cost, battery life, antenna placement, heat, and form factor.

VI. APPLICATIONS

6.1 Smart Homes

A smart home is an unusually mixed radio environment. Tiny sensor reports and battery-powered actuators share the air with voice assistants, cameras, appliances, and high-rate entertainment, and most of these devices have no use for gigabit service. What Wi-Fi 7 offers here is that MLO and scheduled access can protect responsive control traffic while downloads and video streams keep running. A nearby television or headset may take a wide 6-GHz channel while distant sensors stay on a legacy band. The same split exposes the coverage limitation, since 6 GHz is less forgiving through walls, so placement, fallback, and sometimes a second AP matter more than peak speed.

6.2 Enterprise Networks

An enterprise WLAN has to reconcile dense association, roaming, security policy, heavy uplink demand, and several generations of clients at once. Multi-RU OFDMA packs many short transactions into scheduled opportunities, and MLO can move a capable client off a locally busy link. Channel width stays site-specific: a 320-MHz cell can be defensible in an auditorium or a media workspace and counterproductive on an office floor where reuse determines total capacity. A radio refresh also exposes dependencies on multi-gigabit switching, PoE, controller telemetry, and end-to-end QoS. Gains arrive gradually, because a new AP still spends airtime on Wi-Fi 5 and Wi-Fi 6 stations.

6.3 Industrial Internet of Things

Industrial communication is judged by deadlines and by failure behavior, not by average throughput. EHT MLO supplies frequency diversity and shortens access-delay tails, while scheduled OFDMA and integration with time-sensitive networking improve queue control [30]. All of it still runs in unlicensed spectrum, where correlated interference and contention remain facts of life. A credible industrial design needs admission control, redundant paths, careful RF surveys, synchronized edge processing, and a fail-safe control strategy. Wi-Fi 7 strengthens such a system; it does not substitute for safety engineering.

6.4 Healthcare

Hospitals run telemetry, mobile clinical terminals, imaging, asset tracking, telepresence, and confidential records over one infrastructure. Capacity helps with large studies and video consultation, while link diversity matters more for mobility and resilience. Healthcare studies single out MLO, low-latency scheduling, and support for dense device populations as the promising enablers [31]. Clinical deployment still demands evidence gathered in the actual interference environment, with alarm deadlines, cybersecurity, data governance, and fallback for lost 6-GHz service treated explicitly. A peak-rate mode that has not been validated on site should never become a hidden dependency of critical monitoring.

6.5 Smart Cities

Municipal deployments cover cameras, public access, transport telemetry, environmental sensing, and edge analytics, indoors and outdoors. Wi-Fi 7 fits concentrated demand in transport hubs and public buildings well, while 6-GHz power classes and incumbent protection limit outdoor use. City-scale designs will stay heterogeneous, with fiber and private cellular providing continuity and Wi-Fi adding economical local capacity. Secure onboarding, lifecycle ownership, spectrum planning, and coordination across agencies decide the outcome long before the PHY ceiling does.

6.6 Augmented and Virtual Reality

Cloud- or edge-rendered AR/VR needs sustained throughput, but motion-to-photon delay and its variability usually decide whether the experience works. At short range, 320 MHz carries a high-resolution stream, and MLO can place control traffic on a quieter link or steer the whole flow away from contention. Encoders should track measured tail delay rather than the advertised PHY rate. Retuning and retransmission have to respect the rendering deadline as well, because a detailed frame that arrives after the display deadline is worthless. Once radio access is short, edge placement and wired transport tend to dominate what is left.

6.7 Cloud Gaming

Cloud gaming presents a similar interactive workload across a wider range of displays, clients, and Internet paths. MLO shields small control packets from a busy video link and smooths local contention spikes. It cannot shorten a distant cloud route, unload a saturated server, or remove delay from the display pipeline. The most useful implementation would expose access-delay telemetry to bitrate and frame-pacing logic so the application can react to the part of the path that is actually changing.

6.8 Edge Computing

Wi-Fi 7 and edge computing improve each other only when they are managed together. High-capacity local access makes offloaded inference, rendering, digital-twin updates, and preprocessing practical, while nearby compute removes the delay that sits beyond the radio segment. Per-link MLO telemetry could guide service placement or migration, and an edge controller could prioritize link assignments by application deadline in return. Fig. 7 summarizes the relationship. Optimizing each layer on its own is risky, since the least congested radio link may well lead to the busiest edge path.

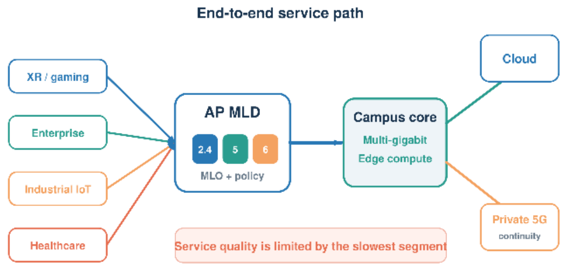

Fig. 7. End-to-end Wi-Fi 7 deployment architecture linking heterogeneous clients, AP MLDs, multi-gigabit transport, edge services, cloud, and private 5G.

Table 4. Application requirements and relevant Wi-Fi 7 mechanisms

Domain	Dominant requirement	Useful Wi-Fi 7 mechanism	Remaining system constraint
Smart home	Mixed traffic and interference	Link selection, OFDMA	Coverage and consumer-device diversity
Enterprise	Density and mobility	Multi-RU, MLO, puncturing	RF reuse, wired uplinks, legacy airtime
Industrial IoT	Tail latency and availability	Scheduled access, link diversity	Admission control and common-mode failure
Healthcare	Mobility, privacy, dependable alarms	MLO and capacity	Validation, security, clinical governance
AR/VR and gaming	High rate with low jitter	320 MHz, earliest-link selection	Rendering, transport, and edge delay
Smart city/edge	Scalable local capacity	EHT access plus edge integration	Regulation, outdoor coverage, orchestration

VII. CHALLENGES

7.1 Cost

The price of a Wi-Fi 7 AP covers only part of an upgrade. Extra RF chains, filters, antennas, faster baseband processing, memory, and multi-gigabit interfaces all raise the bill of materials, and STR MLO is the most demanding of them. A site may also need new switch ports, cabling, PoE capacity, licenses, and controller resources. Procurement should therefore weigh total cost against a defined service objective, such as supported density or tail delay, instead of against the AP price or the advertised PHY rate.

7.2 Deployment

Six-gigahertz operation changes survey results, channel plans, and capacity assumptions. The planner has to decide whether a 320-MHz cell is worth the channels it removes from reuse, and has to decide that for clients with different link sets, radio modes, and power classes. Tools that report only an aggregate association are not adequate for this. Deployment and assurance platforms need per-link load, cross-band coverage, puncturing state, backhaul utilization, and RF-chain constraints, together with a record of which link served each traffic class.

7.3 Compatibility

Backward compatibility keeps older clients connected without turning them into EHT devices. They consume airtime, may invoke protection procedures, and cannot participate in MLO. Wi-Fi 7 products are not equivalent to one another either: certification profiles, drivers, operating systems, radio architectures, and optional capabilities all vary. A serious interoperability program has to exercise STR, non-STR, and EMLSR pairings, along with roaming, power save, mixed security modes, and recovery when one link disappears.

7.4 6-GHz Regulation

The 6-GHz band is not harmonized worldwide. In the United States, low-power indoor, very-low-power, and AFC-controlled standard-power devices operate under different conditions [12], [33]. European Union and United Kingdom frameworks opened the lower portion of the band first, each with its own power limits and outdoor rules [13], [34]. Other administrations have chosen partial allocations, phased ones, or none at all. Devices must obey local operating classes and coordination databases, so a planner cannot assume that the same 320-MHz channel, or even the same pair of MLO links, exists in every country.

7.5 Security

Wi-Fi 7 inherits WPA3 and protected management mechanisms, but MLO enlarges both the protocol state and the attack surface. Association, key handling, link addition, address mapping, roaming, and per-link power transitions all have to stay mutually consistent. The Dragonblood analysis of WPA3's Dragonfly handshake is a useful warning that a modern protocol can still fall to side channels, downgrade paths, or implementation decisions [35]. Multi-link testing should cover replay across links, inconsistent authorization, denial of service aimed at a control link, privacy leakage from persistent link relationships, and manipulation of the telemetry fed to automated controllers.

7.6 Energy Consumption

Energy use climbs when several radios, a 320-MHz analog front end, high-rate conversion, repeated sounding, and continuous link monitoring are all active. EMLSR cuts concurrent hardware and pays for it in retuning and wake-up. Earlier work on Wi-Fi power saving shows that traffic timing and MAC coordination can matter as much as instantaneous PHY power [36]. A sensible MLO controller keeps only the radios needed for the current deadline and throughput target awake, uses target wake time where it fits, and resists leaving a wide idle link on simply because it is there.

VIII. FUTURE RESEARCH DIRECTIONS

8.1 AI-Driven Wi-Fi

Controlling a Wi-Fi 7 network is a partially observed, non-stationary problem. Link occupancy, interference, mobility, and packet deadlines shift far faster than a conventional configuration cycle. Machine learning has already been applied to channel access, rate control, association, scheduling, and anomaly detection [37], and MLO link selection, RU packing, puncturing, and interference prediction are the obvious next targets. Progress should be measured by sample efficiency, safe exploration, and transfer across buildings and device vendors, not by a gain obtained in one simulator.

8.2 Self-Optimizing Networks

A self-optimizing WLAN cannot treat channel width, primary channel, transmit power, the active MLO link set, traffic steering, and edge placement as unrelated settings. Local optima conflict: adding power or widening one cell often degrades several neighbors. A hierarchical architecture looks more promising, with fast local decisions taken inside constraints set by a slower network-wide controller. Because the optimizer acts on production connectivity, stability, explainability, intent checking, and rapid rollback are engineering requirements, not optional governance features.

8.3 Wi-Fi Sensing

IEEE 802.11bf-2025 standardizes enhancements for WLAN sensing [38]. Channel-state and timing measurements can reveal presence, motion, location, health-related activity, and properties of the surrounding environment [39], [40], and Wi-Fi 7 adds bandwidth and link diversity that may sharpen resolution or robustness. The harder questions sit outside raw accuracy: consent, privacy, adversarial manipulation, calibration, transfer between rooms, and competition between sensing exchanges and data traffic. Future APs may have to schedule communication and sensing as two services sharing one medium instead of treating sensing as a free by-product.

8.4 Integration with 5G

Wi-Fi and 5G solve different parts of the access problem. Licensed spectrum and wide-area mobility favor 5G; inexpensive local capacity and LAN integration favor Wi-Fi. Convergence research should keep those strengths intact while coordinating credentials, traffic steering, QoS mapping, time synchronization, and edge-service continuity. At the radio level, 5G NR-U adds a coexistence problem of its own, since two independently scheduled systems may occupy the same 6-GHz unlicensed channel [32]. Work on time-sensitive networking makes the broader point: bounded queues and end-to-end translation are necessary, and a fast air interface by itself is not sufficient [41].

8.5 Integration with 6G

The ITU-R IMT-2030 framework highlights immersive and massive communication, hyper-reliable low-latency service, ubiquitous coverage, AI-and-communication, and integrated sensing [42]. Those themes overlap with the Wi-Fi research agenda without making Wi-Fi a 6G radio interface. The plausible form of convergence is shared policy, identity, edge infrastructure, and sensing across complementary access networks. AI-native 6G visions push further toward semantic service objectives and distributed intelligence [43], [44]. The task for research is to find where such coordination improves end-to-end utility without giving up the low cost, local control, and ownership model that make WLANs attractive in the first place.

8.6 AI Network Management

AI-based management reaches past radio tuning into assurance and day-to-day operations. By correlating per-link telemetry with application delay, client capability, spectrum events, and wired queues, a controller could separate an RF problem from an overloaded service path. Digital twins can test a proposed change before it touches the network, and federated learning can share patterns without exporting raw site data. Automation brings failure modes of its own: biased telemetry, concept drift, poisoning, unsafe actions, and unclear responsibility. Anything operationally credible will need confidence estimates, evidence an operator can inspect, constrained policies, and a secure rollback path.

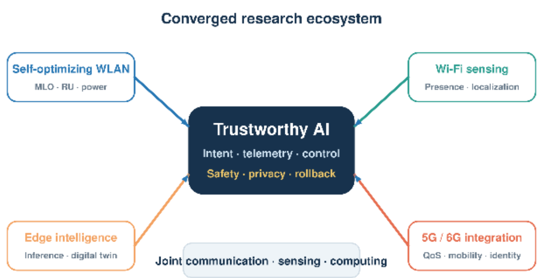

Fig. 8. Future Wi-Fi ecosystem: trustworthy AI coordinates communication, sensing, edge computing, and complementary 5G/6G access under security and energy constraints.

IX. CONCLUSION

Reducing Wi-Fi 7 to "twice the channel width" misses the architectural change at its center. The MLD gathers several links under one association and lets them serve aggregation, earlier access, or selective resilience. Together with 320-MHz channels, 4096-QAM, multi-RU allocation, EHT OFDMA/MU-MIMO, and flexible puncturing, that design lifts the ideal single-link maximum to roughly 23.06 Gbit/s and widens the trade space among capacity, delay, reliability, and energy. It also settles two draft-era claims: the final standard keeps eight spatial streams, and broad multi-AP coordination was never completed as a general 802.11be feature.

What users actually experience depends on the client radio, channel reuse, regional 6-GHz rules, correlated interference, the wired path, security controls, and the scheduler. MLO shortens delay tails and makes individual radio failures easier to absorb, yet it cannot guarantee deterministic service in unlicensed spectrum. The questions worth pursuing therefore concern measurable service objectives, cross-layer and energy-aware control, trustworthy automation, integrated sensing, and coordination with edge, 5G, and 6G systems. Engineered as part of an end-to-end system, Wi-Fi 7 is a capable platform for high-rate, low-latency local access. Judged by its peak PHY number alone, its most important contribution stays invisible.

REFERENCES

[1] IEEE, "IEEE 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 2: Enhancements for extremely high throughput (EHT)," IEEE Std 802.11be-2024, Jul. 2025. [Online]. Available: https://standards.ieee.org/ieee/802.11be/7516/

[2] Wi-Fi Alliance, "Wi-Fi Alliance introduces Wi-Fi CERTIFIED 7," Jan. 8, 2024. [Online]. Available: https://www.globenewswire.com/news-release/2024/01/08/2805409/0/en/wi-fi-alliance-introduces-wi-fi-certified-7.html

[3] D. López-Pérez, A. Garcia-Rodriguez, L. Galati-Giordano, M. Kasslin, and K. Doppler, "IEEE 802.11be extremely high throughput: The next generation of Wi-Fi technology beyond 802.11ax," IEEE Commun. Mag., vol. 57, no. 9, pp. 113–119, Sep. 2019, doi: 10.1109/MCOM.001.1900338.

[4] E. Khorov, I. Levitsky, and I. F. Akyildiz, "Current status and directions of IEEE 802.11be, the future Wi-Fi 7," IEEE Access, vol. 8, pp. 88664–88688, 2020, doi: 10.1109/ACCESS.2020.2993448.

[5] C. Deng, X. Fang, X. Han, et al., "IEEE 802.11be Wi-Fi 7: New challenges and opportunities," IEEE Commun. Surveys Tuts., vol. 22, no. 4, pp. 2136–2166, 2020, doi: 10.1109/COMST.2020.3012715.

[6] A. Garcia-Rodriguez, D. Lopez-Perez, L. Galati-Giordano, and G. Geraci, "IEEE 802.11be: Wi-Fi 7 strikes back," IEEE Commun. Mag., vol. 59, no. 4, pp. 102–108, Apr. 2021, doi: 10.1109/MCOM.001.2000711.

[7] IEEE, "Enhancements for higher throughput," IEEE Std 802.11n-2009, Oct. 2009. [Online]. Available: https://standards.ieee.org/ieee/802.11n/3952/

[8] IEEE, "Enhancements for very high throughput for operation in bands below 6 GHz," IEEE Std 802.11ac-2013, Dec. 2013. [Online]. Available: https://standards.ieee.org/standard/802_11ac-2013.html

[9] B. Bellalta, "IEEE 802.11ax: High-efficiency WLANs," IEEE Wireless Commun., vol. 23, no. 1, pp. 38–46, Feb. 2016, doi: 10.1109/MWC.2016.7422404.

[10] IEEE, "Enhancements for high-efficiency WLAN," IEEE Std 802.11ax-2021, May 2021. [Online]. Available: https://standards.ieee.org/ieee/802.11ax/7180/

[11] Wi-Fi Alliance, "Wi-Fi Alliance delivers Wi-Fi 6E certification program," Jan. 7, 2021. [Online]. Available: https://www.globenewswire.com/news-release/2021/01/07/2154890/0/en/wi-fi-alliance-delivers-wi-fi-6e-certification-program.html

[12] Federal Communications Commission, "Unlicensed use of the 6 GHz band," Report and Order and Further Notice of Proposed Rulemaking, FCC 20-51, Apr. 23, 2020. [Online]. Available: https://docs.fcc.gov/public/attachments/FCC-20-51A3.pdf

[13] European Commission, "Commission Implementing Decision (EU) 2021/1067 on the harmonised use of radio spectrum in the 5945–6425 MHz frequency band," Jun. 17, 2021. [Online]. Available: https://eur-lex.europa.eu/eli/dec_impl/2021/1067/oj/eng

[14] C. Chen, X. Chen, D. Das, D. Akhmetov, and C. Cordeiro, "Overview and performance evaluation of

Wi-Fi 7," IEEE Commun. Standards Mag., vol. 6, no. 2, pp. 12–18, Jun. 2022, doi: 10.1109/MCOMSTD.0001.2100082.

[15] M. Yang and B. Li, "Survey and perspective on extremely high throughput (EHT) WLAN—IEEE 802.11be," Mobile Netw. Appl., vol. 25, no. 5, pp. 1765–1780, Oct. 2020, doi: 10.1007/s11036-020-01567-7.

[16] Á. López-Raventós and B. Bellalta, "Multi-link operation in IEEE 802.11be WLANs," IEEE Wireless Commun., vol. 29, no. 4, pp. 94–100, Aug. 2022, doi: 10.1109/MWC.006.2100404.

[17] Á. López-Raventós and B. Bellalta, "Dynamic traffic allocation in IEEE 802.11be multi-link WLANs," IEEE Wireless Commun. Lett., vol. 11, no. 7, pp. 1404–1408, Jul. 2022, doi: 10.1109/LWC.2022.3171442.

[18] M. Carrascosa-Zamacois, G. Geraci, E. Knightly, and B. Bellalta, "Wi-Fi multi-link operation: An experimental study of latency and throughput," IEEE/ACM Trans. Netw., vol. 32, no. 1, pp. 308–322, Feb. 2024, doi: 10.1109/TNET.2023.3283154.

[19] M. Carrascosa-Zamacois, L. Galati-Giordano, A. Jonsson, G. Geraci, and B. Bellalta, "Performance and coexistence evaluation of IEEE 802.11be multi-link operation," in Proc. IEEE WCNC, 2023, pp. 1–6, doi: 10.1109/WCNC55385.2023.10118829.

[20] B. Bellalta, M. Carrascosa, L. Galati-Giordano, and G. Geraci, "Delay analysis of IEEE 802.11be multi-link operation under finite load," IEEE Wireless Commun. Lett., vol. 12, no. 4, pp. 595–599, Apr. 2023, doi: 10.1109/LWC.2023.3235001.

[21] G. Naik, D. Ogbe, and J.-M. J. Park, "Can Wi-Fi 7 support real-time applications? On the impact of multi link aggregation on latency," in Proc. IEEE ICC, 2021, pp. 1–6, doi: 10.1109/ICC42927.2021.9500256.

[22] X. Lan, X. Zu, and J. Yang, "Enhanced multilink single-radio operation for the next-generation IEEE 802.11be Wi-Fi systems," Security Commun. Netw., vol. 2022, Art. no. 7018360, 11 pp., 2022, doi: 10.1155/2022/7018360.

[23] S. Kim and J.-H. Yun, "Wider-bandwidth operation of IEEE 802.11 for extremely high throughput: Challenges and solutions for flexible puncturing," IEEE Access, vol. 8, pp. 213840–213853, 2020, doi: 10.1109/ACCESS.2020.3040429.

[24] P. Imputato, S. Avallone, and D. Magrin, "Multi-AP coordination in Wi-Fi 7 exploiting time resources sharing," in Proc. IEEE MeditCom, 2022, pp. 166–171, doi: 10.1109/MeditCom55741.2022.9928663.

[25] S. Verma, T. K. Rodrigues, Y. Kawamoto, M. M. Fouda, and N. Kato, "A survey on multi-AP coordination approaches over emerging WLANs: Future directions and open challenges," IEEE Commun. Surveys Tuts., vol. 26, no. 2, pp. 858–889, 2024, doi: 10.1109/COMST.2023.3344167.

[26] Cisco, "Wi-Fi 7 and the growing future of wireless—Design guide," Apr. 8, 2025. [Online]. Available: https://www.cisco.com/c/en/us/products/collateral/networking/wireless/wifi7-future-of-wireless-dg.html

[27] Intel, "Intel Wi-Fi 7 series," 2026. [Online]. Available: https://www.intel.com/content/www/us/en/products/details/wireless/wi-fi-7-series.html

[28] Qualcomm Technologies, "Qualcomm FastConnect 7800," 2022. [Online]. Available: https://www.qualcomm.com/wi-fi/products/fastconnect/fastconnect-7800

[29] Broadcom, "BCM67263/BCM6726: 4×4 802.11be Wi-Fi 7 access point radio," Apr. 2022. [Online]. Available: https://www.broadcom.com/products/wireless/wireless-lan-infrastructure/bcm67263

[30] T. Adame, M. Carrascosa-Zamacois, and B. Bellalta, "Time-sensitive networking in IEEE 802.11be: On the way to low-latency WiFi 7," Sensors, vol. 21, no. 15, Art. no. 4954, Jul. 2021, doi: 10.3390/s21154954.

[31] Y. A. Qadri, Zulqarnain, A. Nauman, A. Musaddiq, E. Garcia-Villegas, and S. W. Kim, "Preparing Wi-Fi 7 for healthcare Internet-of-Things," Sensors, vol. 22, no. 16, Art. no. 6209, Aug. 2022, doi: 10.3390/s22166209.

[32] G. Naik, J.-M. Park, J. Ashdown, and W. Lehr, "Next generation Wi-Fi and 5G NR-U in the 6 GHz bands: Opportunities and challenges," IEEE Access, vol. 8, pp. 153027–153056, 2020, doi: 10.1109/ACCESS.2020.3016036.

[33] Federal Communications Commission, "Unlicensed use of the 6 GHz band; expanding flexible use in mid-band spectrum between 3.7 and 24 GHz," Memorandum Opinion and Order, FCC 21-100, Sep. 2021. [Online]. Available: https://docs.fcc.gov/public/attachments/FCC-21-100A1_Rcd.pdf

[34] Ofcom, "Improving spectrum access for Wi-Fi: Spectrum use in the 5 and 6 GHz bands," Jul. 24, 2020. [Online]. Available: https://www.ofcom.org.uk/spectrum/frequencies/improving-spectrum-access-for-wi-fi

[35] M. Vanhoef and E. Ronen, "Dragonblood: Analyzing the Dragonfly handshake of WPA3 and EAP-pwd," in Proc. IEEE Symp. Security Privacy, 2020, pp. 517–533, doi: 10.1109/SP40000.2020.00031.

[36] E. Guérin, T. Begin, and I. Guérin Lassous, "An overview of MAC energy-saving mechanisms in Wi-

Fi," Comput. Commun., vol. 203, pp. 129–145, 2023, doi: 10.1016/j.comcom.2023.03.003.

[37] S. Szott, K. Kosek-Szott, P. Gawłowicz, et al., "Wi-Fi meets ML: A survey on improving IEEE 802.11 performance with machine learning," IEEE Commun. Surveys Tuts., vol. 24, no. 3, pp. 1843–1893, 2022, doi: 10.1109/COMST.2022.3179242.

[38] IEEE, "Enhancements for wireless local area network (WLAN) sensing," IEEE Std 802.11bf-2025, 2025. [Online]. Available: https://standards.ieee.org/ieee/802.11bf/11574/

[39] R. Du, H. Hua, H. Xie, et al., "An overview on IEEE 802.11bf: WLAN sensing," IEEE Commun. Surveys Tuts., vol. 27, no. 1, pp. 184–217, 2025, doi: 10.1109/COMST.2024.3408899.

[40] Y. Ma, G. Zhou, and S. Wang, "WiFi sensing with channel state information: A survey," ACM Comput. Surv., vol. 52, no. 3, Art. no. 46, pp. 1–36, Jun. 2019, doi: 10.1145/3310194.

[41] M. K. Atiq, R. Muzaffar, Ó. Seijo, I. Val, and H.-P. Bernhard, "When IEEE 802.11 and 5G meet time-sensitive networking," IEEE Open J. Ind. Electron. Soc., vol. 3, pp. 14–36, 2022, doi: 10.1109/OJIES.2021.3135524.

[42] ITU-R, "Framework and overall objectives of the future development of IMT for 2030 and beyond," Recommendation ITU-R M.2160-0, Nov. 2023. [Online]. Available: https://www.itu.int/rec/R-REC-M.2160-0-202311-I/en

[43] K. B. Letaief, W. Chen, Y. Shi, J. Zhang, and Y.-J. A. Zhang, "The roadmap to 6G: AI empowered wireless networks," IEEE Commun. Mag., vol. 57, no. 8, pp. 84–90, Aug. 2019, doi: 10.1109/MCOM.2019.1900271.

[44] W. Saad, M. Bennis, and M. Chen, "A vision of 6G wireless systems: Applications, trends, technologies, and open research problems," IEEE Netw., vol. 34, no. 3, pp. 134–142, May/Jun. 2020, doi: 10.1109/MNET.001.1900287.