telecom

MDPI

Article

Multi-Scenario Performance Evaluation of 5G NR Networks: From Physical-Layer MIMO to Deployment Density, Traffic Heterogeneity, and Self-Optimization

Alina Yanko [1,*], Yuriy Ponochovnyi [2], Anatolii Shostia [3], Olena Ponochovna [4] and Yevhen Zhyvylo [1]

[1] Department of Computer and Information Technologies and Systems, National University "Yuri Kondratyuk Poltava Polytechnic", 36011 Poltava, Ukraine; ziviloevgenij651@gmail.com
[2] Department of Information Systems and Technologies, Poltava State Agrarian University, 36003 Poltava, Ukraine; yuriy.ponochovnyy@pdau.edu.ua
[3] Department of Animal Husbandry Production Technology, Poltava State Agrarian University, 36003 Poltava, Ukraine; anatoliy.shostya@pdau.edu.ua
[4] Department of Economics and Public Administration, Poltava State Agrarian University, 36003 Poltava, Ukraine; olena.ponochovna@pdau.edu.ua
* Correspondence: al9_yanko@ukr.net; Tel.: +380-95-423-8943

Abstract

The performance of fifth-generation New Radio (5G NR) networks is typically evaluated layer by layer, using disconnected methodologies that make it difficult to compare the returns of design decisions across layers or to trace how deficiencies at one layer propagate to another. This article proposes a unified metric taxonomy that enables direct comparison of results spanning four network layers (physical, deployment, traffic, and autonomous management), obtained from five independently developed simulation studies that were not co-simulated. We quantify: (i) the diversity–multiplexing trade-off and massive multiple-input multiple-output (MIMO) scaling, showing a 3.5× spectral-efficiency gain from 4 to 64 antennas with diminishing returns; (ii) the density–capacity trade-off between dense indoor and sparse rural enhanced mobile broadband (eMBB) deployment, showing that the two regimes are limited by different mechanisms (link-level signal-to-noise ratio indoors, path loss and coverage radius in rural areas) and therefore respond to different scaling levers (MIMO order versus frequency-differentiated multi-band layering); (iii) massive machine-type communications (mMTC) and ultra-reliable low-latency communications (URLLC) channel-provisioning sensitivity, where increasing channels from 1 to 4 resolves most contention-driven loss while aggregate metrics can mask class-specific URLLC reliability erosion; and (iv) self-organizing network (SON) self-optimization, showing a disproportionate improvement in worst-case link quality (+9.9%) relative to already-adequate aggregate quality (+1.65%), with two key performance indicator (KPI) targets unreached. The cross-scenario comparison reveals a consistent threshold–saturation pattern across all four layers, attributable to structurally analogous mechanisms, from which we derive cross-scenario network-planning guidelines.

Keywords: 5G New Radio; Massive MIMO; network deployment planning; mMTC; URLLC; self-organizing networks; reinforcement learning; cross-layer performance evaluation

Check for updates

Academic Editor: Michael Hempel

Received: 24 August 2026
Revised: 12 October 2026
Accepted: 15 October 2026
Published: 23 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Fifth-generation New Radio (5G NR) networks are designed around three structurally distinct usage scenarios defined by 3GPP—enhanced Mobile Broadband (eMBB), massive Machine-Type Communications (mMTC), and Ultra-Reliable Low-Latency Communications (URLLC)—each with different, sometimes conflicting, radio-resource requirements [1]. In practice, a single physical network must simultaneously support all three: a dense office deployment carrying high-throughput video, a rural site serving sparse sensor traffic, and a mixed cell handling both massive IoT reporting and latency-critical control signals—all built on the same underlying physical-layer antenna technology and increasingly managed by autonomous, self-organizing control loops rather than static, manually tuned configuration [2].

This fragmentation has a concrete practical cost. A network planner deciding how to allocate a fixed capital budget across antenna arrays (physical layer), site density and frequency planning (deployment layer), channel provisioning (traffic layer), and autonomous-optimization tooling (management layer) has no single, evidence-based framework for comparing the relative return on investment across these four categories of decision—only separate, non-comparable bodies of literature for each. Without such a framework, planning decisions risk either over-investing in one layer past the point of diminishing returns while under-provisioning another, or relying on a later layer (e.g., self-organizing network (SON)-based optimization) to compensate for deficiencies that originate earlier in the design process (e.g., insufficient deployment density) and that later layers cannot, in fact, correct.

This article addresses the gap identified above by proposing a unified metric taxonomy that enables direct comparison of results across four network layers: physical, deployment, traffic, and autonomous management. The comparison draws on five independently developed and validated 5G NR simulation studies—physical-layer multiple-input multiple-output (MIMO) characterization, dense indoor eMBB deployment, sparse rural eMBB deployment, mixed mMTC/URLLC traffic, and SON-based autonomous optimization—spanning these four layers, with the deployment layer represented by two independent studies (dense indoor and sparse rural/private-network deployment). The studies were run and validated in isolation, without dynamic co-simulation between layers. The proposed taxonomy does not compare simulation tools; it compares network layers: for each layer, it expresses the response of a layer-specific key performance indicator (KPI) to a layer-specific resource variable (Sections 2 and 4.1), so that the shape of this response and the dependencies between layers can be assessed on a common basis. The scenarios (dense indoor, sparse rural, and mixed mMTC/URLLC traffic) are the contexts in which the response of each layer is measured; they are not themselves the objects of comparison. Rather than presenting the five studies as unrelated case studies, this article makes the following contributions:

1. A unified methodology and metric taxonomy (Section 2) that maps the five studies, spanning four network layers (Table 1), onto one consistent set of KPIs, enabling direct cross-layer comparison for the first time across this combination of scenarios.
2. Quantitative characterization of four network layers (Section 3)—physical-layer diversity–multiplexing trade-offs and Massive MIMO scaling; the density–capacity trade-off between dense indoor and sparse rural eMBB deployment; mMTC/URLLC channel-provisioning sensitivity and class-specific reliability erosion; and SON-based self-configuration and self-optimization outcomes, including an explicit, transparent distinction between the tabular method actually implemented and the deep-learning extension described but not executed in the source study.

https://doi.org/10.3390/telecom7060137

3. A cross-scenario synthesis (Section 4) identifying a recurring threshold–saturation response pattern across all four layers, tracing it to four distinct but structurally analogous mechanisms (coding-threshold behavior, array-gain saturation, queueing-blocking decay, and exploration-schedule decay), and formalizing three evidenced cross-layer dependencies into a set of cross-scenario design guidelines—the article's principal contribution beyond what any single constituent study established.

Table 1. Mapping of component studies to the unified framework.

Component	Scenario/Layer	Environment	Core Method
Physical-layer MIMO model	PHY foundation (all scenarios)	Python (NumPy)	Analytical closed-form BER (SISO, diversity) and Monte Carlo BER/SE simulation (SIMO-MRC, MIMO-ZF), 3GPP TR 38.901 UMa NLOS path loss
Indoor 5G NR model	eMBB—dense deployment	MATLAB 5G Toolbox	3GPP CDL channel + full PHY chain (LDPC, HARQ)
Private 5G rural model	eMBB—sparse deployment	Python	Empirical propagation models (Hata-Okumura, COST-231, 3GPP RMa/UMa)
Multichannel mMTC/URLLC model	mMTC/URLLC traffic heterogeneity	Python (SimPy)	Discrete-event simulation, stochastic traffic models
SON model	Autonomous network optimization	Python (NumPy)	Multi-agent tabular Q-learning

The limitations of this approach—including the absence of full model co-simulation and the parameter-regime-specific nature of the identified thresholds—are stated explicitly in Section 4.5 and Section Future Work, together with a concrete future-work agenda.

Throughout this article, scalars are denoted by italic letters (e.g., γ, N_t), vectors by boldface lowercase letters, and matrices by boldface uppercase letters. All metrics in Table 2 are real-valued scalars, whereas the channel matrix $\mathbf{H} \in \mathbb{C}^{N_r \times N_t}$ and the zero-forcing matrix $\mathbf{W} \in \mathbb{C}^{N_t \times N_r}$ of the physical-layer model (Section 2) are complex-valued; $(\cdot)^H$ denotes the Hermitian (conjugate) transpose.

Table 2. Unified notation and metric definitions.

Symbol	Metric	Unit	Domain of Applicability
γ	Signal-to-noise ratio (SNR)	dB	PHY, eMBB
PL	Path loss	dB	PHY, eMBB
σ_{SF}	Shadow-fading standard deviation	dB	eMBB (indoor)
τ_{DS}	RMS delay spread	ns	eMBB (indoor)
BER	Bit error rate	–	PHY
BLER	Block error rate	–	eMBB (indoor)
η (SE)	Spectral efficiency	bit/s/Hz	PHY, SON
R	Throughput	Mbit/s	eMBB
D	Latency (mean/P95)	ms	eMBB, mMTC/URLLC
PDR/SR	Packet delivery/success rate	%	mMTC/URLLC
R_{cov}	Coverage radius	km	eMBB (rural)
NQI	Aggregated network-quality index	0–1	SON

The remainder of this article is organized as follows. Section 1.1 reviews related work across the four layers examined and formalizes the cross-layer synthesis gap this study addresses. Section 2 describes the unified simulation methodology and metric taxonomy (Materials and Methods). Section 3 presents Results across the four layers: physical-layer MIMO performance (Section 3.1), dense indoor and sparse rural eMBB deployment (Section 3.2), mMTC/URLLC traffic coexistence (Section 3.3), and SON-based

self-configuration and self-optimization (Section 3.4). Section 4 synthesizes findings across all four layers into a unified threshold–saturation pattern and derives cross-scenario design guidelines. Section 5 concludes and outlines directions for future work. Figure 1 depicts the resulting four-layer structure and the cross-layer dependencies established in Section 4.

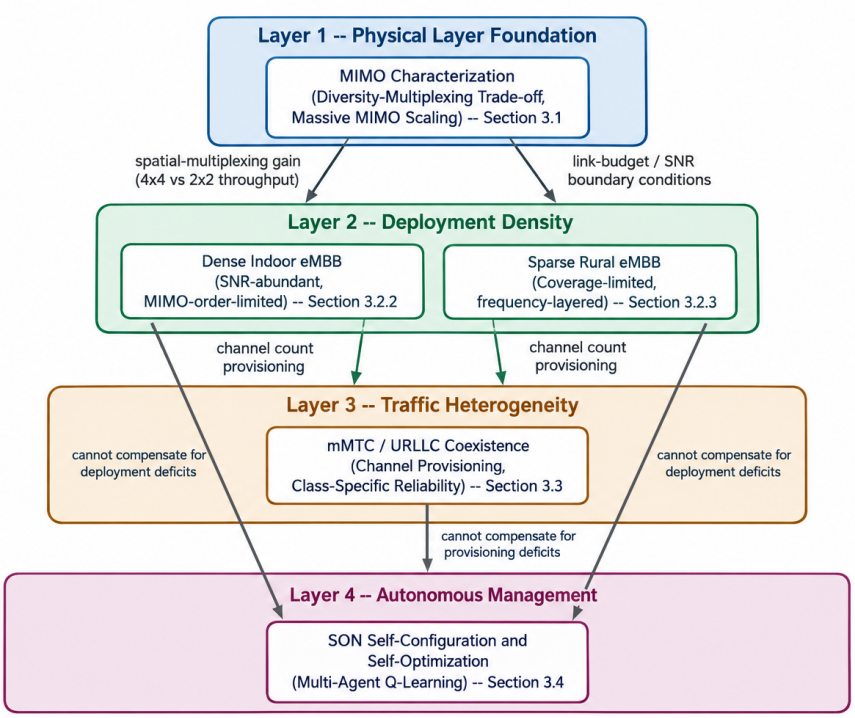


Figure 1. Structural framework of the unified multi-scenario methodology, showing the four layers examined (physical, deployment, traffic, autonomous management) and the cross-layer dependencies evidenced in Section 4. Abbreviations: MIMO, multiple-input multiple-output; eMBB, enhanced mobile broadband; mMTC, massive machine-type communications; URLLC, ultra-reliable low-latency communications; SON, self-organizing network.

The remainder of this section situates this framework within the existing literature on each of its four constituent layers before Section 2 details the unified methodology underlying it.

1.1. Related Work

1.1.1. Physical-Layer MIMO and Massive MIMO Evolution

Massive MIMO remains the subject of intense survey activity, reflecting its role as a foundational 5G technology now being extended toward 6G. Recent comprehensive reviews trace the evolution from conventional MIMO to massive MIMO (mMIMO), extremely large MIMO (XL-MIMO), and cell-free massive MIMO (CF-mMIMO), documenting spectral-efficiency gains of up to ∼20 bit/s/Hz through advanced beamforming and simultaneous multi-user service [3,4]. Antenna-hardware-focused surveys specifically address the design challenges of sub-6 GHz massive MIMO arrays for 5G, including pilot contamination, element isolation, and envelope correlation coefficient (ECC)—practical constraints closely related to the spatial-correlation analysis conducted in this study (Section 3.1.5) [5]. A parallel research direction integrates massive MIMO with intelligent reflecting surfaces (IRS) and cell-free architectures as candidate 6G enablers, extending the antenna-count scaling studied here (Section 3.1.4) toward distributed, coordinated deployments that remove the single-cell boundary entirely [6,7]. Industry-oriented documentation from equipment vendors and standards bodies further confirms that Massive MIMO evolution toward 3GPP Release 18 remains an active area of both academic and standardization work [8,9].

Nonetheless, this literature is predominantly link- or array-level: it characterizes MIMO performance as an end in itself, rarely connecting antenna-configuration trade-offs quantitatively to downstream deployment-density or traffic-layer outcomes—the connection established in Sections 3.1.6 and 4.4 of this article.

1.1.2. eMBB Deployment: Dense Indoor and Sparse Rural/Private Networks

Indoor 5G NR performance has been examined empirically in industrial and enterprise settings, with recent work reporting field measurements of 5G NR in indoor industrial environments that corroborate the general link-quality behavior (SNR-dependent throughput saturation) modeled in Section 3.2.2 of this study [10]. Broader surveys of indoor coverage-improvement techniques catalog the range of engineering responses—distributed antenna systems, repeaters, small cells, beamforming—available once the SNR/BLER threshold identified in Section 3.2.2 is not met by macro-cell signal alone [11]. On the opposite end of the density spectrum, private and rural 5G has attracted substantial applied and market-oriented attention: agriculture is repeatedly identified as a leading vertical for private 5G deployment, encompassing crop sensing, variable-rate irrigation, and autonomous machinery, with recent industry reports projecting rapid market growth (∼29% CAGR, 2025–2034) for 5G-in-agriculture specifically [12,13]. Field-deployment case studies confirm the practical relevance of the coverage-radius/frequency trade-offs modeled in Section 3.2.3, reporting comparable engineering responses—tower-mounted sector antennas, multi-kilometer coverage radii per site—to those obtained independently in this study's simulation [14,15]. Regulatory and policy literature documents that the coverage gap motivating rural/private deployment is a persistent, only partially resolved global challenge as of 2025–2026, reinforcing that the sparse-deployment scenario modeled in Section 3.2.3 addresses an active rather than a largely solved problem [16,17]. What is largely absent from this literature is a direct, quantitative comparison of the indoor-dense and rural-sparse regimes within a single consistent metric framework—most sources address one regime or the other, but not, as attempted in Section 3.2.4, the structural trade-off between them.

1.1.3. mMTC/URLLC Coexistence and Traffic Heterogeneity

The problem of serving eMBB, mMTC, and URLLC traffic classes on shared radio resources has an established communication-theoretic foundation, most influentially formalized by Popovski et al.'s treatment of 5G network slicing as an information-theoretic coexistence problem [18], which remains a standard reference point for subsequent work. Building on this foundation, recent research has shifted toward learning-based resource management: deep-reinforcement-learning approaches for intelligent eMBB/URLLC resource slicing [19], and hybrid NOMA-based network-slicing schemes explicitly evaluated for simultaneous eMBB/URLLC/mMTC coexistence, reporting spectrum-efficiency and QoS-satisfaction improvements over orthogonal multiple access baselines [20]. Access-layer studies specifically targeting mMTC contention—e.g., LSTM-based access-class-barring schemes for mixed IoT traffic—address the same random-access contention problem modeled via the ALOHA-based approach in Section 2.6, but via adaptive/learned barring rather than static channel provisioning [21]. Comparative studies of decoding architectures for eMBB-URLLC-mMTC network slicing confirm that channel-count/resource-provisioning sensitivity—the central finding of Section 3.3.2—remains an actively studied lever, though this literature typically optimizes the access scheme itself rather than establishing the more basic channel-count sufficiency threshold this study identifies [22]. The threshold-driven diminishing-returns behavior found in Section 3.3.2 (success rate saturating after ≈4 channels) is broadly consistent with the queueing-theoretic expectations of this literature but has not, to the

authors' knowledge, been reported as an explicit, quantified "knee" applicable across both channel-count and device-population scaling dimensions simultaneously, as done in Sections 3.3.2 and 3.3.3.

1.1.4. Self-Organizing Networks and Reinforcement-Learning-Based Optimization

SON remains an active research and standardization area. Klaine et al. provide a comprehensive survey of machine learning techniques applied to self-organizing cellular networks, covering self-configuration, self-optimization, and self-healing approaches [23], complemented by Moysen and Giupponi's treatment of self-organized network management meeting machine learning across the 4G-to-5G transition [24]. The self-healing dimension of SON depends fundamentally on accurate fault diagnosis: misclassified or delayed diagnostics can themselves become a safety-relevant failure mode, as shown by indicator- and requirement-based analyses of the influence of diagnostics errors on system safety [25], a consideration directly relevant to the PCI-conflict and fault-detection procedures examined in Section 3.4. A recent survey specifically addresses handover-parameter self-optimization for 6G mobile networks, noting persistent gaps between academic SON proposals and field-deployable solutions [26]. Industry analyses confirm SON's growing commercial significance and its anticipated centrality to 6G-era ultra-dense network management [27], a trajectory reinforced by the O-RAN Alliance's own use-case and deployment guidance for RAN Intelligent Controller-based autonomous optimization [28]. Notably, this literature—including the source study underlying Section 3.4—consistently identifies deep reinforcement learning, following the Deep Q-Network (DQN) formulation of Mnih et al. [29] and beyond, as the performance frontier, with tabular Q-learning treated as a transitional or pedagogical baseline; the limitation acknowledged in Section 3.4.2 (implemented method is tabular, not DQN) is therefore consistent with, rather than contradicting, the trajectory of this literature.

1.1.5. Synthesis: The Gap Addressed by This Study

Across all four sub-fields, a consistent structural gap emerges: each layer is a mature, independently active research area, but cross-layer synthesis is rare. MIMO surveys do not connect antenna-configuration trade-offs to deployment-density outcomes; indoor and rural deployment studies are not typically compared within one consistent framework; traffic-heterogeneity research generally assumes deployment-layer parameters as fixed inputs rather than examining their joint sensitivity with deployment decisions; and SON research, while increasingly sophisticated methodologically, is rarely evaluated for its dependency on—and inability to compensate for—deficiencies at the deployment and traffic layers below it. No source identified above reports the recurring threshold–saturation pattern documented independently at all four layers in this study (Section 4.1) as a structural property linking them, nor derives the cross-layer design guidelines presented in Section 4.4. This study's contribution is therefore not a new result at any single layer—each subsection of Section 3 individually reproduces results consistent with, and corroborated by, the literature reviewed above—but the systematic cross-layer synthesis connecting them, positioned as filling the gap this review identifies.

2. Materials and Methods

2.1. Overview and Rationale for a Unified Framework

The performance evaluation of 5G New Radio (NR) systems spans several logically distinct layers—from the physical layer (channel propagation and antenna processing) through network-level deployment (coverage and capacity planning) and heterogeneous traffic engineering (mMTC/URLLC coexistence) up to autonomous network management

(self-organizing optimization). Existing studies typically address a single layer in isolation, using tool-specific parameterizations that are difficult to compare directly. To enable a coherent cross-layer analysis, this study compares five independently developed simulation models—covering, respectively, MIMO physical-layer characterization, dense indoor deployment, sparse rural deployment, mixed mMTC/URLLC traffic, and self-organizing network (SON) optimization—on the basis of one unified metric taxonomy, described in this section.

The framework is organized along the taxonomy defined by 3GPP usage scenarios—eMBB, mMTC, and URLLC [1]—extended with an autonomous-management dimension corresponding to SON functionality as defined in 3GPP TS 32.500-series specifications [2]. Table 1 summarizes the correspondence between the five original simulation studies, the scenario/layer they represent, and the software environment used.

The indoor eMBB model was implemented in MATLAB R2025b (version 25.2.0.2998904; The MathWorks, Inc., Natick, MA, USA) using the 5G Toolbox and the Communications Toolbox (both version 25.2, R2025b; The MathWorks, Inc.). The remaining four models were implemented in Python 3.14.6 with NumPy 2.5.1 (physical-layer, traffic, and SON models), SciPy 1.18.1 (special functions in the physical-layer model), pandas 3.0.5 and Matplotlib 3.11.1 (traffic-model post-processing and figure generation), and the SimPy 4.1.2 library (discrete-event traffic model); the rural coverage planner uses only the Python standard library. All Python packages listed are open-source software. Diagram source code was rendered with PlantUML 1.2026.9 and Graphviz (via Viz.js 3.17.0).

Comparing quantitative results across these five heterogeneous environments requires a common set of metrics and symbols, defined next.

2.2. Unified System of Metrics and Notation

To allow direct comparison of results obtained with heterogeneous tools, all five component models are mapped onto a common set of key performance indicators (KPIs), summarized in Table 2.

These symbols are used consistently throughout Sections 3.1–3.4. The following five subsections describe, in turn, the simulation methodology underlying each component model.

2.3. Physical-Layer Simulation Methodology (MIMO Foundation)

The physical-layer component was implemented in Python (NumPy) to characterize the fundamental trade-offs of Multiple-Input Multiple-Output (MIMO) transmission that underlie all subsequent deployment- and traffic-level analyses. Path loss for the reference urban macro (UMa) non-line-of-sight scenario was computed according to the 3GPP TR 38.901 model [30]:

$$PL_{\text{UMa-NLOS}} = 13.54 + 39.08 \log_{10}(d_{3D}) + 20 \log_{10}(f_c) - 0.6(h_{UT} - 1.5), \quad (1)$$

where d_{3D} is the three-dimensional distance between base station and user equipment, f_c the carrier frequency (3.5 GHz), and h_{UT} the UE antenna height (1.5 m). For the reference distance of 150 m, the model yields $PL = 109.46$ dB, consistent with expected values for dense urban NLOS conditions.

Figure 2 summarizes this Monte Carlo BER-estimation procedure. After verifying the rank condition $N_t \leq N_r$ and fixing the random seed for reproducibility, the algorithm sweeps the SNR range; at each SNR point g, the noise variance is set to $\sigma^2 = N_t/g$ and the realization loop begins. In each realization, an $N_r \times N_t$ i.i.d. complex Gaussian channel matrix $\mathbf{H}$ is drawn and $\mathbf{H}^H\mathbf{H}$ is tested for singularity; singular realizations are caught and skipped, while well-conditioned ones proceed to solve $\mathbf{H}^H\mathbf{H}\mathbf{W} = \mathbf{H}^H$ for the ZF matrix

W. A random BPSK-modulated bit vector is transmitted through **H** with additive complex Gaussian noise, detected via **W**, and demodulated by sign; bit errors are accumulated over all realizations to yield the empirical BER at that SNR point.

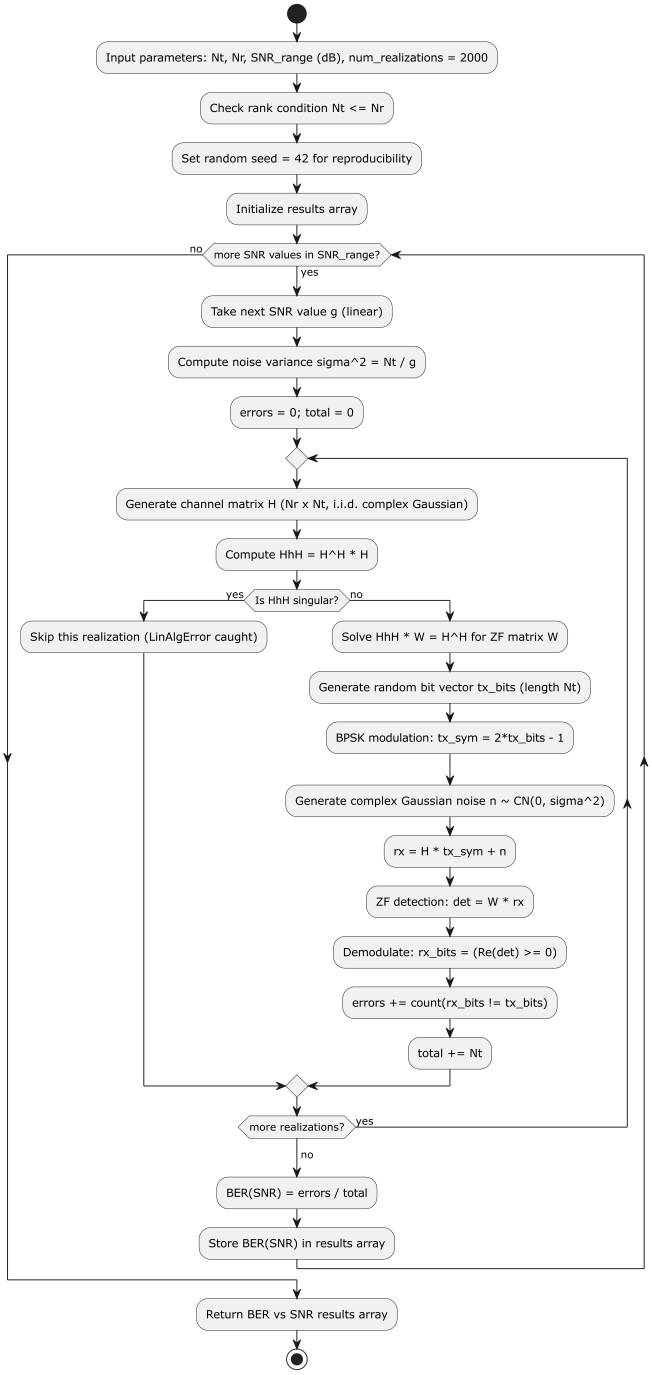


Figure 2. Monte Carlo algorithm for MIMO Zero-Forcing BER estimation.

Two distinct methods were used to obtain the BER curves reported in Section 3.1, matching the actual implementation rather than a uniform Monte Carlo procedure. SISO BER curves (AWGN and Rayleigh channels, five modulation orders) and the receive-diversity curves ($d = 1, 2, 4, 8$) were computed from closed-form analytical expressions (standard Q-function-based formulas for AWGN and Rayleigh fading, and the corresponding MRC diversity-combining expression), not simulated. Only the SIMO-MRC and MIMO Zero-Forcing configurations were obtained via Monte Carlo simulation with 2000 channel realizations per SNR point (the corresponding numbers of transmitted bits are given in

Section 3.1.3). The Zero-Forcing precoding/detection matrix was obtained not by explicit matrix inversion but by solving the linear system $(\mathbf{H}^H\mathbf{H})\mathbf{W} = \mathbf{H}^H$ via numerically stable linear-solve routines, avoiding explicit inversion of a potentially ill-conditioned matrix. Reproducibility was ensured by fixing the pseudo-random generator seed prior to each simulation run. Spatial correlation between antenna elements was modeled via a Toeplitz correlation matrix based on the zero-order Bessel function of inter-element spacing (uniform linear array, isotropic scattering assumption)—the classical narrowband fading-correlation model of Jakes [31]—with correlation square-roots computed via eigenvalue decomposition for numerical stability at high correlation coefficients ($\rho \to 1$), consistent with the capacity-degradation mechanism analyzed by Shiu et al. for correlated multi-element antenna systems [32].

The next four subsections describe the deployment, traffic, and management-layer simulation methodologies, beginning with the dense indoor eMBB scenario.

2.4. Dense-Deployment (Indoor eMBB) Simulation Methodology

The dense-deployment scenario was implemented in MATLAB using the 5G Toolbox and Communications Toolbox, which provide a standards-compliant implementation of the 5G NR physical-layer processing chain—transport-block generation, LDPC channel coding, PDSCH/PUSCH mapping, MIMO precoding, and demodulation—together with the 3GPP-defined Clustered Delay Line (CDL) channel models and standard modulation-and-coding-scheme (MCS) tables for physical-layer data procedures [33,34]. Two indoor environments were modeled according to the 3GPP InH (Indoor Hotspot) specification: InH-OfficeMixed (partitioned office) and InH-OfficeOpen (open-plan office), extended in the load-scaling analysis to InH-ShoppingMall and InH-Factory.

Figure 3 depicts this simulated chain end to end: on the transmitter side, bits with CRC-24A attachment are LDPC-encoded, rate-matched and 16-QAM-modulated, mapped onto the resource grid together with DM-RS, and OFDM-modulated; the resulting waveform passes through the channel (AWGN, static, Rayleigh, or Rician with $K = 10$) with noise added at the target SNR; the receiver performs OFDM demodulation, estimates the channel from DM-RS, equalizes the received grid, computes PDSCH log-likelihood ratios, and recovers the transport block through rate recovery, LDPC decoding, and de-segmentation with a CRC check; block/bit errors and throughput are accumulated over successive transport blocks to obtain the BLER, BER, and throughput statistics reported in Section 3.2.2.

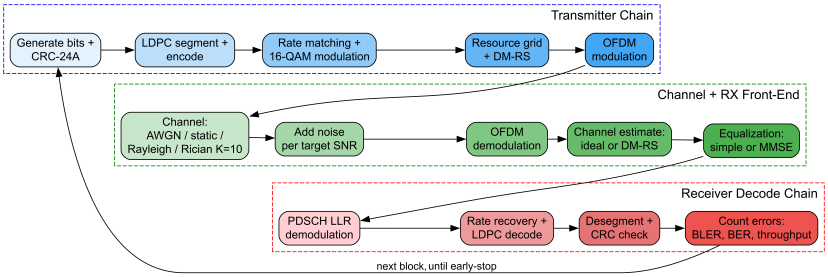


Figure 3. Simplified 5G NR physical-layer processing chain used for the indoor BLER/throughput simulation.

The sparse-deployment methodology, addressing the opposite end of the density spectrum, is described next.

2.5. Sparse-Deployment (Rural Private-Network) Simulation Methodology

The sparse-deployment scenario was implemented in Python as a coverage-planning tool supporting a hybrid propagation-modeling approach, selected by carrier frequency:

the modified Okumura-Hata model for frequencies below 1500 MHz, the extended COST-231 Hata model for 1500–2000 MHz, and a simplified 3GPP TR 38.901 RMa/UMa model above 2000 MHz [30,35]. The governing path-loss expression for the Hata-family models used in this study is:

$$L = a + b\log_{10}(f) + c\log_{10}(h_{bs}) - 13.82\log_{10}(h_{rx}) + (44.9 - 6.55\log_{10}(h_{bs}))\log_{10}(d), \quad (2)$$

where L is path loss (dB), a, b, c are terrain-dependent empirical coefficients, f is carrier frequency (MHz), h_{bs} and h_{rx} are base-station and receiver antenna heights (m), and d is link distance (km).

2.6. Heterogeneous Traffic (mMTC/URLLC) Simulation Methodology

The traffic-heterogeneity component was implemented in Python using the SimPy discrete-event simulation framework [36]. Device traffic was modeled using three stochastic arrival processes: a Poisson process for periodic mMTC reporting, a Beta distribution for bursty/event-driven traffic, and an On/Off model for intermittent sensor activity. mMTC and URLLC device populations shared a common pool of N logical channels under a priority-aware access and load-balancing scheme, with URLLC packets granted preferential channel access consistent with the 3GPP latency target of $\leq$10 ms and reliability target of 99.999% [37].

Figure 4 shows the class structure of the simulator: an abstract `IoTDevice` base class is specialized into `mMTCDevice` and `uRLLCDevice` subclasses implementing the respective traffic models (Poisson, Beta, On/Off, and Periodic for mMTC; a sparse high-priority process for uRLLC); each generated packet is represented as a `PacketEvent` record; a `BaseStation` class implements priority-aware channel selection over the configured access protocol; and a top-level `IoT5GSimulator` class orchestrates the device population, channel pool, and statistics collection, with all run parameters captured in a single `SimulationConfig` object.

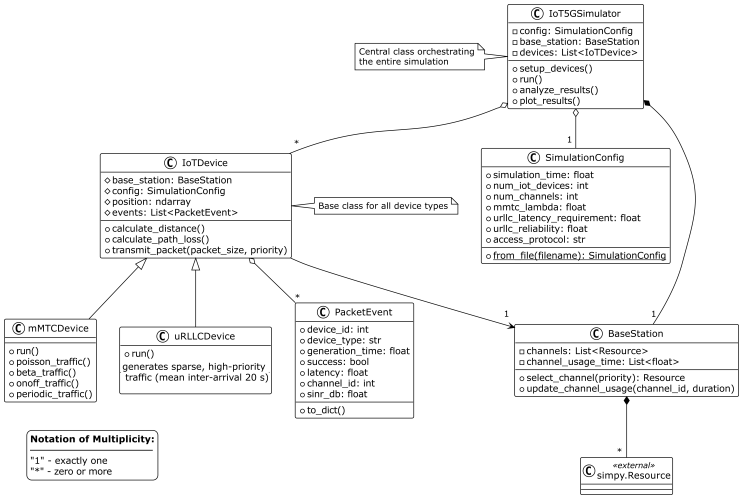


Figure 4. UML class diagram of the mMTC/URLLC multichannel simulator.

The final component model—autonomous network optimization—is described next.

2.7. Self-Organizing Network (SON) Simulation Methodology

The autonomous-optimization component was implemented in Python (NumPy), modeling a heterogeneous network of 27 base stations (7 macro cells, 20 small cells) within a 5 $\times$ 5 km simulation area. The self-configuration stage implements the 3GPP TS 32.501-defined plug-and-play sequence: initial radio-parameter calculation, Automatic Neighbour

Relation (ANR) table construction, Physical Cell Identity (PCI) conflict detection and resolution, and OSS registration [38].

The architecture is centralized in its environment computation but decentralized in its learning update: at each simulation step, the network environment centrally applies all agents' actions, recomputes the network-wide SINR map and KPI set, and returns a single shared scalar reward together with a global state vector. Each agent then extracts only its own local slice of that state (four normalized components—transmit power, antenna downtilt, own-cell load, and average served-UE SINR—discretized to five levels per dimension, $5^4 = 625$ states per agent) and updates its own independent Q-table using the shared reward—there is no explicit inter-agent messaging or X2/Xn-like signaling of the kind standardized for inter-eNB/gNB coordination [39] implemented; coordination emerges solely through the shared reward signal and the shared environment dynamics, consistent with a centralized-training/decentralized-execution-style approach to multi-agent reinforcement learning [40], and avoiding the combinatorial state-space explosion that a single fully joint-state agent would face at network scale. The action space contains five discrete actions per agent (± 2 dBm power step, $\pm 1^\circ$ tilt step, no-change), selected via an ε-greedy policy with Bellman update:

$$Q(s,a) \leftarrow Q(s,a) + \alpha \left[r + \gamma \max_{a'} Q(s',a') - Q(s,a) \right], \tag{3}$$

with learning rate $\alpha = 0.10$ and discount factor $\gamma = 0.95$. The shared reward function combines four weighted components—SINR quality (weight 0.4), coverage (0.3), load-balancing uniformity (0.2), and handover success (0.1)—consistent with established KPI-based SON reward-weighting practice [41]. Training was conducted over 200 episodes of 50 steps, with exploration rate ε decaying multiplicatively from 1.0 to a floor of 0.05 (decay factor 0.975).

Figure 5 shows the resulting software architecture: a `NetworkEnvironment` module composes `BaseStation` and `UserEquipment` objects and centrally computes the SINR map and KPI set each step; a `SelfConfiguration` module implements the plug-and-play sequence described above; a `SelfOptimization` module coordinates one independent `QTable` per agent under the shared-reward update rule; and a top-level driver loads the run configuration, executes the baseline and training phases in sequence, and persists results and diagnostic plots.

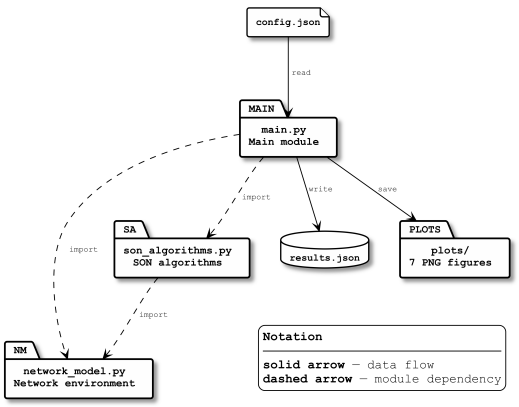


Figure 5. Software architecture of the SON simulation framework.

Having described each of the five component models individually, the following two subsections address, respectively, the cross-model consistency checks applied and the use of generative AI tools in preparing this manuscript.

2.8. Cross-Model Consistency and Reproducibility

All five component models share the following consistency safeguards: (i) fixed pseudo-random generator seeds to ensure reproducibility across repeated runs; (ii) explicit statistical validation of Monte Carlo results against closed-form theoretical expressions where available; (iii) results exported in structured, machine-readable form (JSON/CSV) enabling the metric unification described above without manual transcription error. The name and version of all software used (MATLAB 5G Toolbox, Python, NumPy, SimPy) are given above; further details on data and code availability are given in the Data Availability Statement.

2.9. Use of Generative AI Tools

Generative AI (Claude Sonnet 4.5 and Claude Sonnet 5, Anthropic, San Francisco, CA, USA; and ChatGPT Images 2.5, OpenAI, San Francisco, CA, USA, free tier, used to generate the structural diagrams of the framework and of the cross-layer dependencies (Figure 1 and Section 4.3)) was used to assist in drafting and editing manuscript text, was used to assist in drafting and editing manuscript text, in generating Python/MATLAB plotting and diagram-source code used to render figures from the authors' own simulation outputs, and in cross-referencing quantitative findings across the five underlying studies into the unified framework presented in Section 3. All simulation models, experimental configurations, and raw numerical results reported in this article were designed, implemented, and executed independently by the authors using their own MATLAB and Python code; no simulation data were generated, altered, or interpreted by AI. All AI-assisted text and code were reviewed and verified against source data by the authors, who take full responsibility for the accuracy and integrity of the content.

3. Results

3.1. Physical-Layer Foundation: MIMO Performance

3.1.1. Rationale and Scope

Every scenario examined in this study—dense indoor deployment, sparse rural deployment, and mixed mMTC/URLLC traffic—ultimately operates on a physical layer whose error and capacity behavior is governed by antenna configuration and channel statistics. This section establishes the physical-layer boundaries within which those scenarios operate, using a generic urban macro (UMa) NLOS reference channel at 3.5 GHz. The diversity–multiplexing trade-off is the organizing concept for this section: increasing the number of antennas can be used either to combat fading (diversity, reducing BER at fixed SNR) or to transmit multiple independent data streams (multiplexing, increasing throughput at fixed bandwidth), and these two uses compete for the same spatial degrees of freedom [42].

3.1.2. Diversity Gain: SISO, SIMO-MRC, and the Fading Penalty

The baseline SISO BER behavior in a Rayleigh channel confirms the well-known "fading penalty" relative to an ideal AWGN channel: at SNR = 10 dB, BPSK yields BER = 2.33×10^{-2}, rising to 3.96×10^{-2} (16-QAM), 5.83×10^{-2} (64-QAM), and 8.14×10^{-2} (256-QAM). To reach BER = 10^{-3} with BPSK, the Rayleigh channel requires $\gamma \approx 18$ dB, versus $\gamma \approx 7$ dB in AWGN—an 11 dB fading penalty. Figure 6 plots this BER-vs-SNR behavior across all four modulation orders for both channels, visually confirming that the fading penalty widens further toward higher-order modulation.

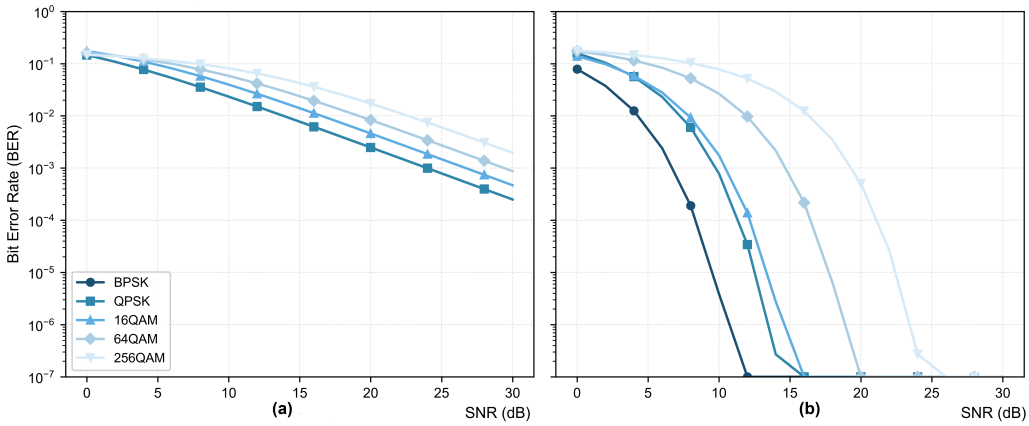


Figure 6. SISO BER performance: (**a**) Rayleigh fading channel; (**b**) AWGN channel.

Spatial diversity closes this gap efficiently. Table 3 summarizes the effect of diversity order d at fixed SNR = 0 dB.

Table 3. BER as a function of diversity order (SNR = 0 dB, BPSK, Rayleigh channel).

Diversity Order d	BER	Gain vs. $d = 1$
1 (SISO)	1.464×10^{-1}	–
2	5.81×10^{-2}	$\times 2.5$
4	1.11×10^{-2}	$\times 13$
8	5.16×10^{-4}	$\times 284$

This diversity gain, however, is obtained by dedicating all spatial degrees of freedom to combating fading; the next subsection examines the opposite use of those same degrees of freedom—spatial multiplexing—and the trade-off between the two.

3.1.3. The Diversity–Multiplexing Trade-Off: MIMO Zero-Forcing Spatial Multiplexing

When the same antenna resources are instead used for spatial multiplexing rather than diversity combining, the BER behavior reverses sharply, as shown in Table 4 and Figure 7.

Table 4. BER for ZF-detected MIMO spatial multiplexing configurations, Rayleigh channel (BPSK, Monte Carlo).

SNR, dB	SIMO 1×4 MRC	MIMO 2×4 ZF	MIMO 3×4 ZF	MIMO 4×4 ZF
0	1.25×10^{-2}	6.85×10^{-2}	1.59×10^{-1}	2.72×10^{-1}
8	$<1.5 \times 10^{-3}$	2.00×10^{-3}	2.37×10^{-2}	1.05×10^{-1}
16	$<1.5 \times 10^{-3}$	$<7.5 \times 10^{-4}$	1.33×10^{-3}	2.09×10^{-2}
20	$<1.5 \times 10^{-3}$	$<7.5 \times 10^{-4}$	$<5.0 \times 10^{-4}$	1.10×10^{-2}

Entries of the form $< x$ in Table 4 denote zero observed bit errors; $x \approx 3/N$ is the corresponding 95% upper confidence bound (rule of three), where N is the number of bits transmitted per SNR point: $N = 2000$ for SIMO 1×4 and $N = 2000\,N_t$ for MIMO $N_t \times 4$ ZF (i.e., 4000, 6000, and 8000 bits for $N_t = 2, 3, 4$).

These single-user diversity/multiplexing trade-offs assume a fixed small number of antennas; the next subsection scales the base-station antenna count into the massive-MIMO regime under a multi-user precoding scheme.

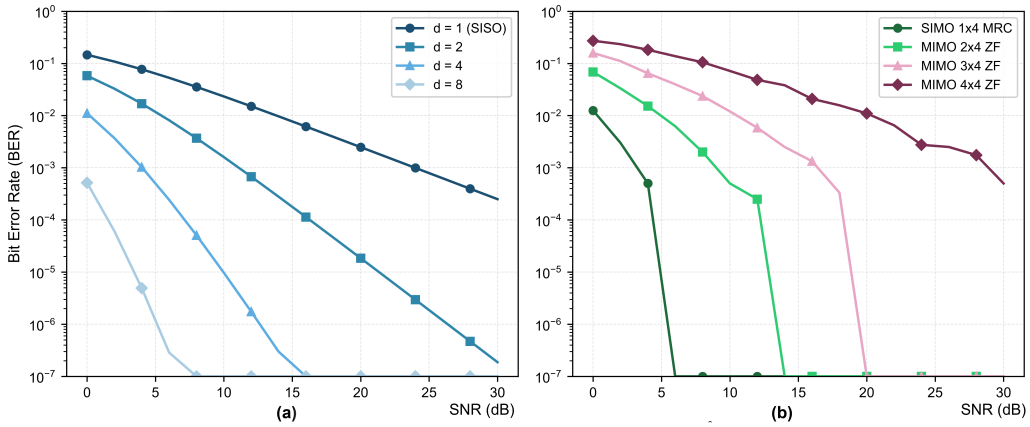


Figure 7. Diversity-multiplexing trade-off: (**a**) receive diversity order (theory, BPSK); (**b**) spatial multiplexing (Monte Carlo, BPSK); points with zero observed errors are drawn at a display floor of 10^{-7} and carry no statistical meaning.

3.1.4. Massive MIMO: Spectral Efficiency and Channel Hardening

For the downlink multi-user MIMO (MU-MIMO) case with $K = 4$ users and a ZF precoder, spectral efficiency (SE) and average per-user SINR both increase monotonically—but with diminishing returns—as the number of base-station antennas N_t grows from 4 to 64 (Table 5, Figure 8).

Table 5. Spectral efficiency and SINR vs. base-station antenna count ($K = 4$ users, $\gamma = 10$ dB).

N_t	SE, bit/s/Hz	Incremental SE Gain	SINR, dB	SINR Gain vs. Input SNR
4	10.47	–	11.08	+1.08 dB
8	21.78	+108%	16.22	+6.22 dB
16	27.76	+27%	20.90	+10.90 dB
32	32.56	+17%	24.51	+14.51 dB
64	36.95	+13%	27.81	+17.81 dB

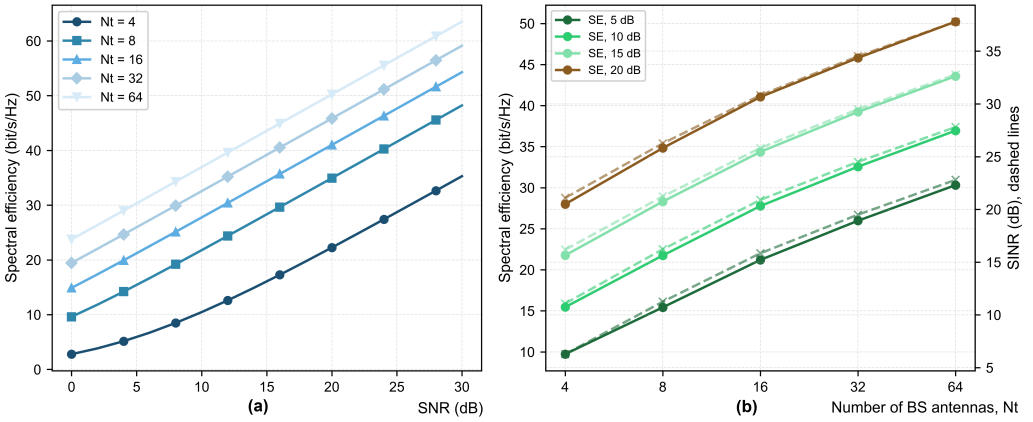


Figure 8. MU-MIMO spectral efficiency and channel hardening: (**a**) SE vs. SNR for increasing BS antenna count; (**b**) SE (solid) and SINR (dashed) vs. N_t.

This antenna-count scaling assumes ideal, uncorrelated antenna elements; the next subsection examines how relaxing that assumption—finite physical antenna spacing—affects the same spectral-efficiency metric.

3.1.5. Impact of Spatial Antenna Correlation

Table 6 and Figure 9 quantify the effect of antenna-element spacing on MU-MIMO spectral efficiency ($N_t = 16$, $\gamma = 15$ dB, $K = 4$).

Table 6. Effect of antenna-element spacing on MU-MIMO spectral efficiency ($N_t = 16$, $\gamma = 15$ dB, $K = 4$).

d/λ	$\lvert\rho\rvert$	SE, bit/s/Hz	Δ vs. $d = 0.5\lambda$
0.10	0.904	26.48	−21.9%
0.25	0.472	32.66	−3.7%
0.50	0.304	33.91	0.0% (baseline)
1.00	0.220	34.15	+0.7%
2.00	0.158	34.27	+1.1%

This behavior is consistent with the classical fading-correlation-capacity relationship established by Shiu et al. [32]: below $d \approx 0.25\lambda$, correlated fading collapses the effective spatial degrees of freedom available to the ZF precoder, whereas above this spacing the array behaves as an effectively independent-fading set of branches. Optimizing hardware positioning and spatial alignment of transceivers plays a crucial role in mitigating path loss and enhancing physical-layer signal quality, as demonstrated in automated control models for radio signal transmission and reception devices [43].

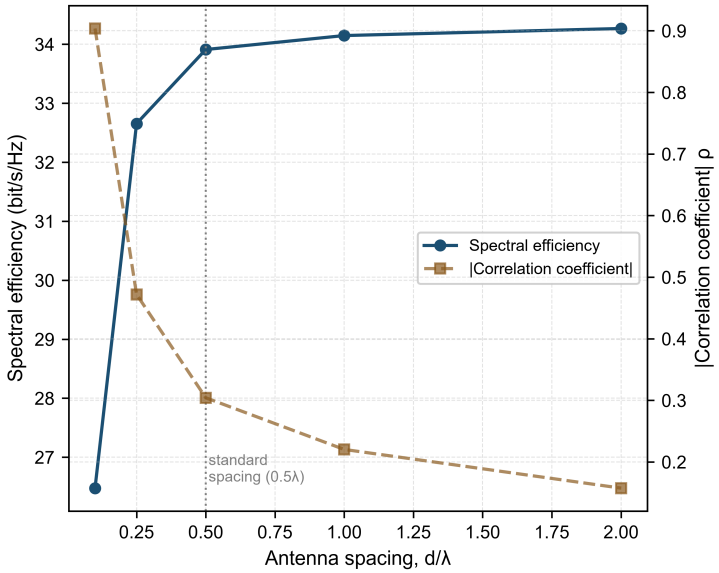


Figure 9. Effect of antenna spacing on spectral efficiency ($N_t = 16$, SNR = 15 dB, $K = 4$).

The following subsection consolidates these diversity, multiplexing, antenna-count, and correlation findings into the boundary conditions carried forward to the deployment, traffic, and SON layers examined in the rest of Section 3.

3.1.6. Summary of Physical-Layer Boundary Conditions

Three quantitative relationships established in this section recur, in modified form, at every subsequent layer of the framework: (1) a fading penalty exists and is removable by diversity, but at the cost of forgoing spatial multiplexing gain; (2) multiplexing/array gains follow a law of diminishing returns; (3) physical implementation constraints impose hard thresholds (antenna correlation below $d = 0.25\lambda$ causes a sharp SE collapse).

3.2. eMBB Deployment Scenarios: Dense Indoor Versus Sparse Rural

3.2.1. Rationale and Scope

eMBB deployments span an extremely wide range of physical environments. This section positions dense indoor and sparse rural deployment as the two opposite poles

of a single density–capacity spectrum, both governed by the physical-layer boundary conditions established in Section 3.1.6.

3.2.2. Dense Indoor Deployment

Channel Characterization

Table 7 summarizes path loss and delay-spread behavior for two 3GPP InH environments across link distances of 5–100 m.

Table 7. Indoor channel parameters at selected distances.

Scenario	Distance, m	PL, dB	τ_{DS}, ns
Mixed office	5	55.4	67.7
Mixed office	100	127.9	80.6
Open office	5	55.4	24.5
Open office	100	107.4	15.1

Figure 10 plots the path-loss and delay-spread values underlying Table 7 across the full 5–100 m distance range for both InH environments: path loss increases monotonically with distance in both scenarios, consistent with the free-space-plus-clutter behavior expected indoors, while delay spread varies non-monotonically with distance and differs substantially in character between the partitioned (Mixed) and open-plan office layouts, reflecting their different multipath geometries.

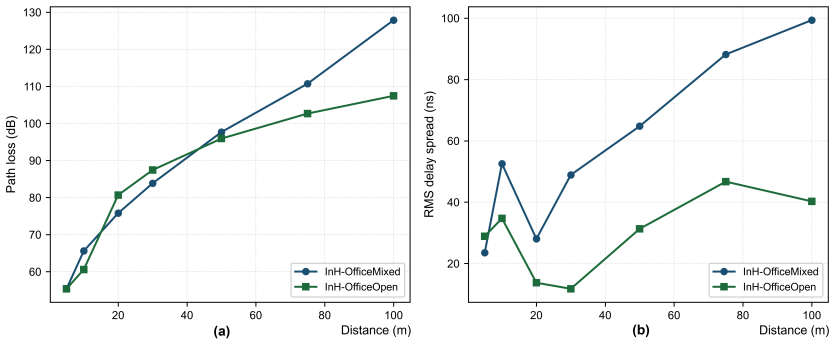


Figure 10. Indoor channel parameters: (**a**) path loss vs. distance; (**b**) delay spread vs. distance, for InH-OfficeMixed vs. InH-OfficeOpen.

Link-Level Performance: The BLER Waterfall

Table 8 and Figure 11 show the sharp BLER transition characteristic of LDPC-coded transmission.

Table 8. PHY-layer performance vs. SNR, Rician channel $K = 10$, 16-QAM.

γ, dB	BLER	Throughput, Mbit/s
$\leq$6	1.00	0
6.5	0.89	2.09
7.0	0.14	16.74
7.5	0.015	19.17
$\geq$8.0	$\approx$0	19.464 (saturation)

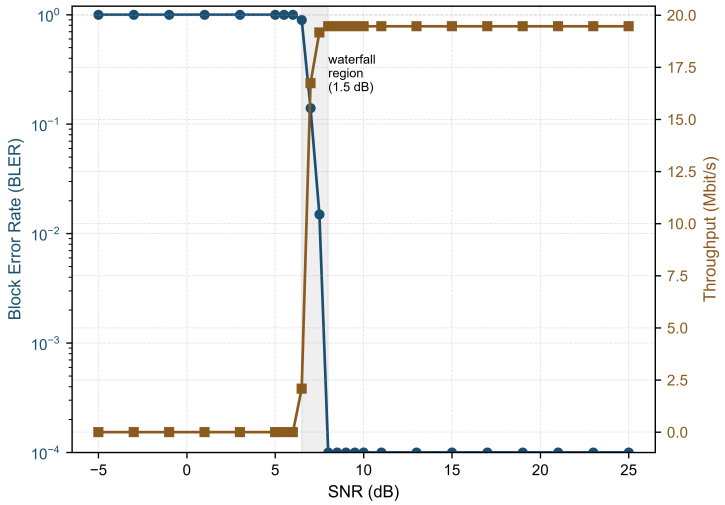


Figure 11. BLER and throughput vs. SNR (Rician channel, $K = 10$, 16-QAM).

System-Level Capacity Scaling

Table 9 reports aggregate KPIs across the four indoor scenarios. Upgrading from 2 × 2 to 4 × 4 MIMO consistently doubles both aggregate and per-user throughput at fixed user count (Section 4), while increasing the user population from 10 to 50 reduces per-user throughput roughly in proportion to the added load, with aggregate throughput only mildly affected due to the fixed pool of physical resource blocks being shared across more users. Jain's fairness index remains at or near 1.0 across most scenarios, indicating near-equal resource allocation among users; the Shopping Mall scenario shows a modest reduction (0.983), consistent with its larger average cell radius (30 m) and higher shadow-fading variability (8 dB standard deviation) causing a wider spread of per-user channel quality.

Table 9. Aggregate KPIs across indoor scenarios.

Scenario	Users	MIMO	R_{total}, Mbit/s	R_{avg}, Mbit/s	SE, bit/s/Hz	Jain
Mixed office	10	2 × 2	7931.5	793.2	80.70	1.000
Mixed office	10	4 × 4	15,863.0	1586.3	161.41	1.000
Mixed office	50	4 × 4	14,688.0	293.8	149.45	1.000
Open office	20	2 × 2	7637.8	381.9	77.71	1.000
Shopping mall	50	2 × 2	7043.8	140.9	71.67	0.983

3.2.3. Sparse Rural Deployment

Coverage–Frequency Trade-Off

Table 10 summarizes coverage radius by frequency band and terrain type, drawn from the 16 base stations of deployment Scenarios 1–4.

Table 10. Coverage radius by frequency and terrain.

Band	Terrain	R_{cov}, km
700 MHz	rural	6.77 (mean; BS_Farm_Center, BS_Warehouse, BS_Main_Low)
700 MHz	open	4.61–5.45
1800 MHz	forest	1.18 (min., BS_Forest_Edge station)
3500 MHz	open	2.59 (mean)
3500 MHz	suburban	1.21 (min.)

The frequency effect alone accounts for a 2–4× reduction in coverage radius moving from 700 MHz to 3.5 GHz. Terrain effects are of comparable magnitude to the frequency

effect: within the same 1800 MHz band, coverage radius varies by a factor of 1.6× between the greenhouse-hub station in rural terrain (1.93 km) and the forest-edge station (1.18 km); a starker terrain contrast appears across bands, where the 700 MHz open-field station (5.45 km) covers 4.6× the radius of the 1800 MHz forest-edge station (1.18 km)—a combined effect of both frequency and vegetation attenuation rather than terrain alone [44]. Figure 12 plots coverage radius against carrier frequency (log scale) for all 16 base stations across deployment Scenarios 1–4, distinguishing stations by terrain type (open, rural, suburban, forest); the resulting spread confirms that frequency and terrain act as largely independent factors shaping the achievable coverage radius.

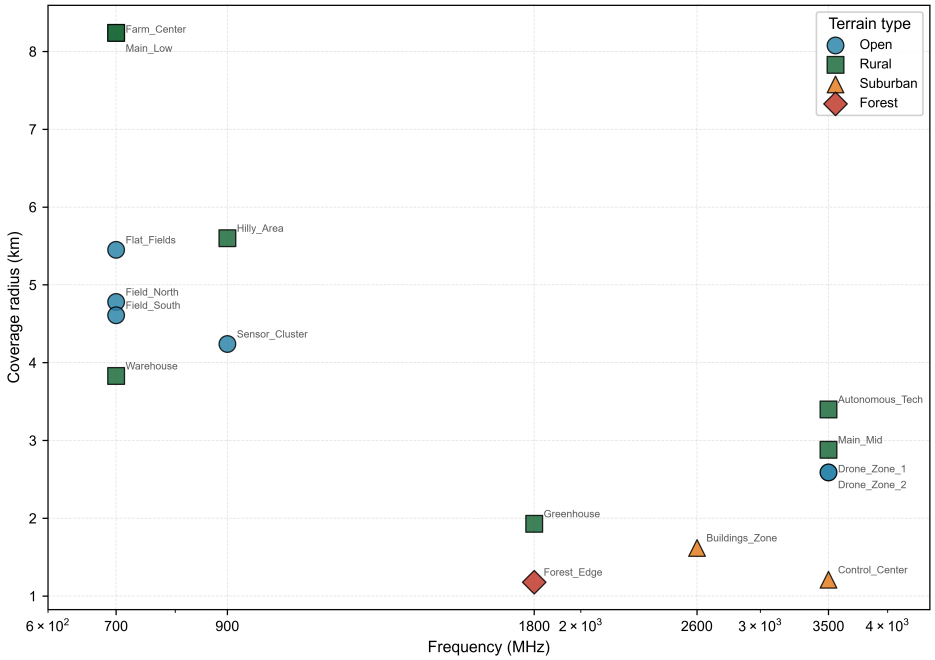


Figure 12. Coverage radius vs. frequency and terrain type (16 base stations, deployment scenarios 1–4).

Optimized Hybrid Deployment

Table 11 lists the four-station configuration that combines all four frequency bands within a single service area, sized for the 1000 ha (10 km^2) core farmland footprint, rather than the single-band per-scenario stations compared in Table 10.

Table 11. Optimized hybrid deployment (Scenario 5).

Station	f, MHz	R_{cov}, km	Area, km^2	Capacity, Mbit/s
BS_Central_Tower	700	8.24	213.08	250
BS_Perimeter_Node	900	4.24	56.55	250
BS_Operations_Center	3500	2.88	26.15	500
BS_Greenhouse_Hub	1800	1.93	11.69	350

Figure 13 shows the spatial layout of this four-station deployment superimposed on the 10 km × 10 km map extent, with each station's coverage disk drawn at its computed radius and shaded by operating band (700 MHz and 900 MHz macro cells, 1800 MHz pico cell, 3500 MHz micro cell); the overlapping low-band disks around the central and perimeter stations provide redundant wide-area coverage, while the operations-center and greenhouse-hub stations deliver higher-capacity coverage to their respective higher-demand zones.

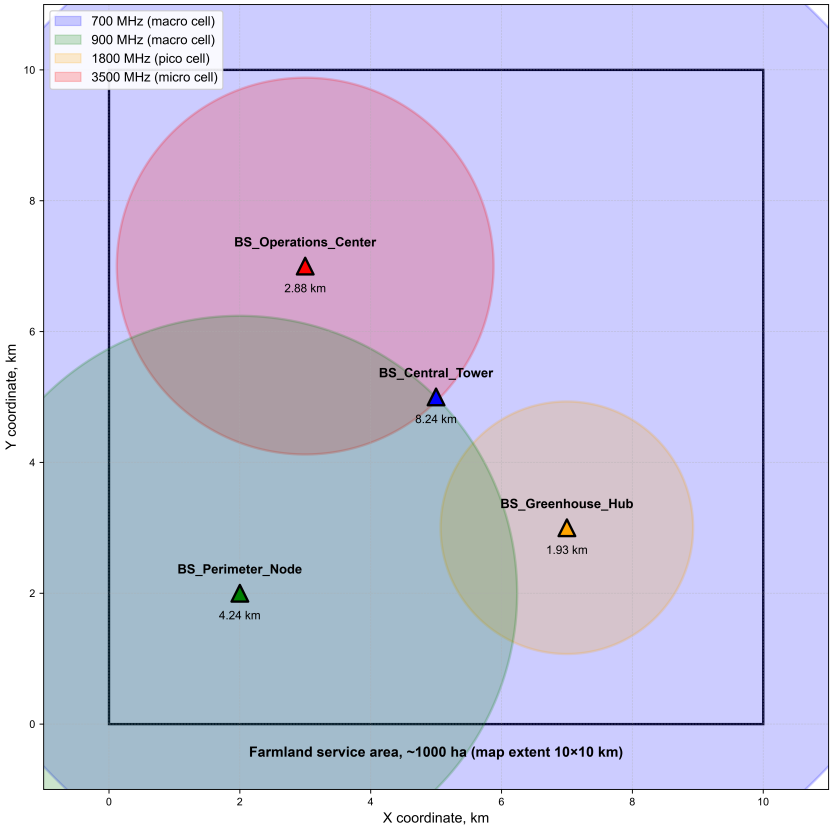


Figure 13. Base station placement and coverage zones (optimized hybrid scenario).

Having examined the dense indoor and sparse rural regimes separately, the following subsection compares them directly along a common set of dimensions.

3.2.4. Comparative Synthesis: The Density–Capacity Trade-Off

Table 12 juxtaposes the dense indoor and sparse rural deployment results from the two preceding subsections along five comparison dimensions, highlighting that the two regimes are limited by different physical mechanisms and consequently respond to different scaling levers.

Table 12. Density–capacity trade-off: dense indoor vs. sparse rural deployment.

Dimension	Dense Indoor	Sparse Rural
Typical inter-site distance	5–100 m	1.2–8.2 km
Dominant limiting factor	Link-level SNR (BLER waterfall)	Path loss/coverage radius
Primary capacity lever	MIMO order ($2 \times 2 \rightarrow 4 \times 4$ doubles throughput)	Frequency-differentiated layering
Peak throughput	Up to 15,863 Mbit/s (aggregate, 10 users, 4×4)	Up to 500 Mbit/s (single station, 3.5 GHz)
Infrastructure scaling response	Add spatial streams within existing footprint	Add base stations/frequency layers

Having established the physical- and deployment-layer boundary conditions, the next subsection turns to a third, largely orthogonal resource-management question: how a fixed pool of logical channels is shared between heterogeneous mMTC and URLLC traffic classes.

3.3. Heterogeneous Traffic Scenario: mMTC/URLLC Coexistence

3.3.1. Rationale and Scope

This section examines resource-contention behavior: how a fixed pool of logical channels is shared between mMTC and URLLC traffic classes with fundamentally incompatible service requirements.

Figure 14 summarizes this architecture end to end: the four mMTC traffic generators (Poisson, Beta, On/Off, and Periodic) together with the sparse high-priority uRLLC generator feed a common packet-formation stage that tags each packet with a priority flag; packets then undergo priority-aware channel selection informed by recent channel-load history, followed by a collision-and-SINR-based success model; accepted packets are logged as `PacketEvent` records and aggregated into the success-rate, latency, and channel-utilization statistics reported in the following sections.

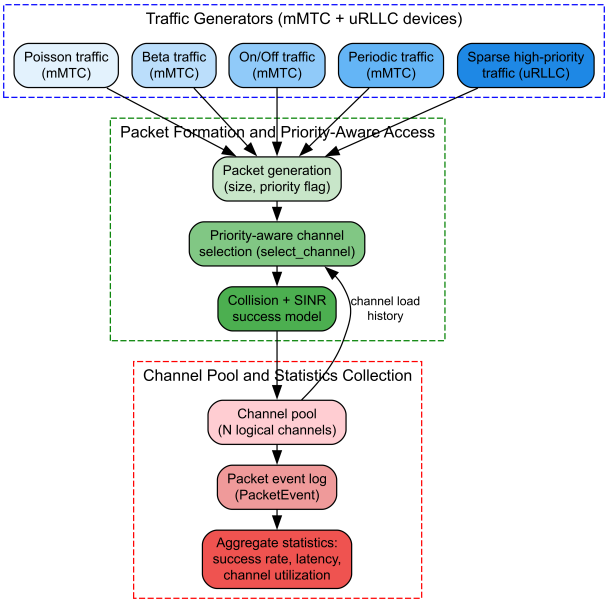


Figure 14. Architecture of the multichannel mMTC/URLLC access system.

3.3.2. Effect of Channel Count on System Performance

Table 13 reports system performance for a fixed population of 80 devices (65 mMTC + 15 URLLC, Poisson arrivals at $\lambda = 0.15$ packets/s, ALOHA access) as the number of logical channels is varied from 1 to 10.

Table 13. System performance vs. number of channels (80 devices, fixed).

Channels	SR, %	D_{avg}, ms	D_{P95}, ms	Util, %	mMTC SR, %	URLLC SR, %
1	97.07	6.66	25.62	4.54	68.68	81.82
4	100.00	4.61	8.89	1.13	93.19	94.32
6	100.00	4.55	8.78	0.76	94.32	93.10
8	100.00	4.49	4.44	0.57	95.33	94.38
10	100.00	4.51	5.18	0.45	94.67	95.70

Figure 15 plots the success-rate and latency trends from Table 13: overall success rate rises sharply from 97.07% to 100% between 1 and 4 channels and remains saturated thereafter, while mMTC-specific success rate improves more gradually (from 68.68% to 93–95%) and P95 latency drops from 25.62 ms to single-digit values, indicating that channel starvation—not raw traffic load—is the dominant bottleneck at low channel counts.

Overall SR and the class-specific mMTC/URLLC SR in Table 13 are computed differently and are not directly comparable in magnitude: Overall SR is a per-packet metric, counting a packet as successful if it is eventually delivered within up to $R_{max} = 5$ retransmission attempts, whereas mMTC/URLLC SR is a per-attempt metric, computed over every individual channel-access attempt including failed retries. Overall SR can therefore substantially exceed the class-specific values without any contradiction.

It is also important to distinguish two related but distinct notions of URLLC performance reported here. Latency compliance—the fraction of successfully delivered URLLC packets meeting the $\leq$10 ms 3GPP latency bound (Section 2.6)—is essentially always satisfied once $N_{ch} \geq 4$ (100.00% across all tested channel counts and device populations except the 1-channel case, at 96.30%). URLLC success rate (SR), by contrast, measures a different and stricter quantity: the fraction of individual URLLC channel-access attempts, including failed retries due to contention, that succeeded at all, and remains in the 91–96% range even at $N_{ch} \geq 4$—well below the 99.999% reliability target cited in Section 2.6. This gap indicates that the bottleneck identified in this study is channel-access contention (a MAC-layer effect), not end-to-end latency performance; closing the remaining gap to the 3GPP reliability target would require addressing this contention (e.g., via priority-preemption or channel reservation), which lies beyond the scope of the ALOHA-based access model implemented here.

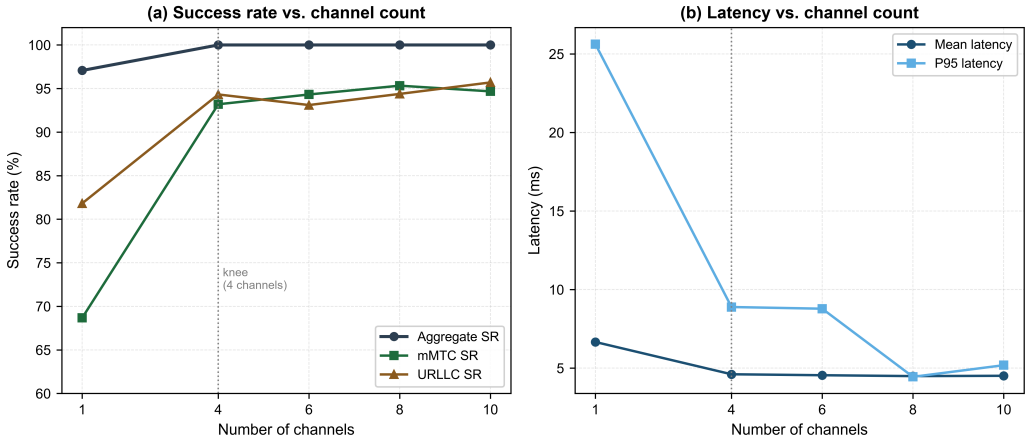


Figure 15. Effect of channel count on mMTC/URLLC coexistence performance: (**a**) success rate; (**b**) latency.

3.3.3. Effect of Device Population (Scalability)

Table 14 reports the same performance metrics as the device population scales from 50 to 150, at a fixed 6-channel pool, with the URLLC share held at roughly 13–20% of devices.

Table 14. System performance vs. device population (6 channels, fixed).

Devices	Packets	SR, %	D_{avg}, ms	D_{P95}, ms	mMTC SR, %	URLLC SR, %	D_{URLLC}, ms
50 (40 + 10)	601	100.00	4.50	4.44	95.03	94.00	2.36
80 (65 + 15)	1060	100.00	4.55	8.78	94.32	93.10	2.39
150 (130 + 20)	2059	100.00	4.67	8.89	93.67	91.38	2.43

Figure 16 plots the corresponding trends from Table 14: overall and mMTC success rates remain above 93% throughout the tested range, while URLLC success rate declines modestly from 94.00% to 91.38% and P95 latency increases from 4.44 ms to 8.89 ms as the device population—and hence channel contention—grows, indicating graceful rather than catastrophic degradation under increasing load.

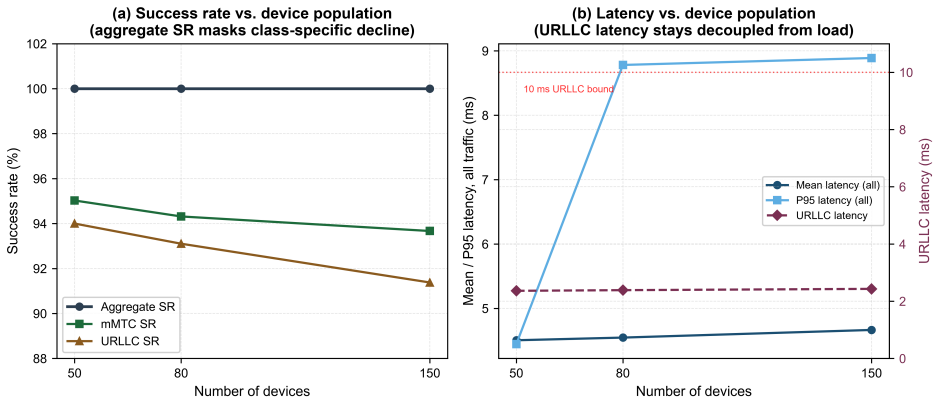


Figure 16. Scalability: effect of device population on QoS: (**a**) success rate; (**b**) latency.

Having characterized the physical, deployment, and traffic layers, the final layer examined in this study is autonomous network management.

3.4. Autonomous Network Optimization: Self-Organizing Network (SON) Layer

3.4.1. Self-Configuration: Automated Deployment

Table 15 shows that total configuration success rate reached 100%, and total self-configuration time (17.7 min for 27 stations) is dominated by the software/configuration download phase (39.5% of total time) rather than by the ANR radio-relation-discovery procedure itself.

Table 15. Self-configuration process: time breakdown and outcomes.

Phase	Duration, min	Share, %
Platform boot/initialization	4.0	22.6
Software/configuration download	7.0	39.5
ANR procedure	4.7	26.6
OSS registration	2.0	11.3
Total	17.7	100

The resulting configuration comprises 326 neighbor relations across 27 base stations (mean neighbor-table size: 12.14 for macro cells, 12.05 for small cells), and zero PCI conflicts were detected. Figure 17 visualizes both aspects: panel (a) shows the neighbor-relation-table size distribution separately for macro and small cells, confirming the near-identical mean table sizes despite the different cell counts (7 macro vs. 20 small), and panel (b) shows the time breakdown across the four configuration phases.

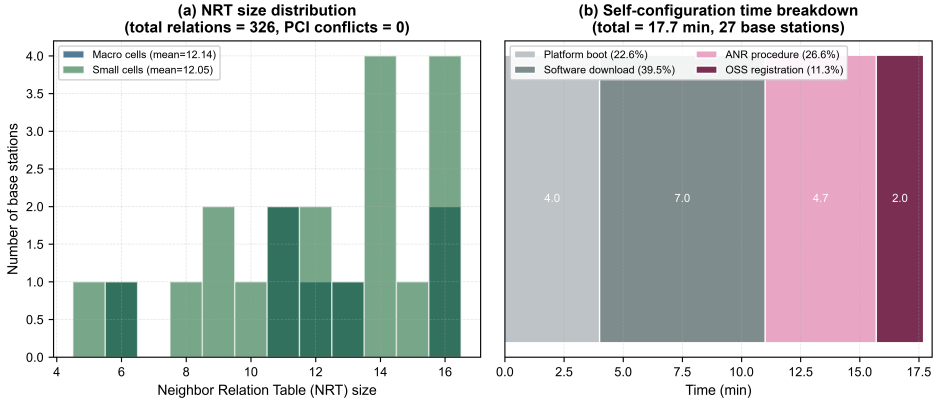


Figure 17. Self-configuration: (**a**) neighbor relation table (NRT) size distribution; (**b**) time breakdown.

3.4.2. Self-Optimization: Multi-Agent Q-Learning Formulation

Figure 18 illustrates one simulation step of this loop: each agent extracts its local four-feature state slice from the shared global state vector and selects an action via its own ε-greedy policy; the environment centrally applies all agents' actions, moves the UEs, recomputes the SINR map and global KPI set, and derives the single shared reward; each agent then updates its own Q-table from that shared reward via the Bellman update, the environment advances to the next state, and ε is decayed once per completed episode.

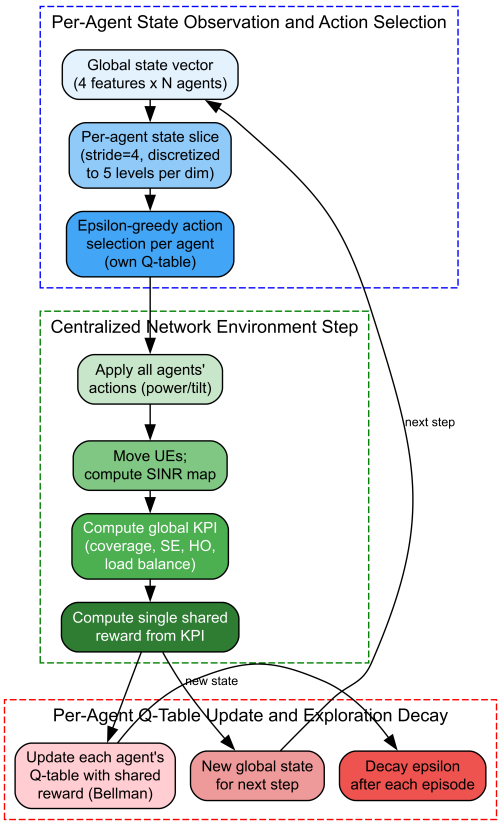


Figure 18. Agent-environment interaction loop for multi-agent Q-learning.

3.4.3. Optimization Outcomes: Baseline Versus Trained Policy

Table 16 compares the pre-training baseline configuration (averaged over the first 30 simulation steps) against the trained policy (averaged over the final training episodes) across the network KPIs entering the reward function.

Table 16. Network KPIs: baseline vs. post-optimization.

Indicator	Baseline	After SON	Target
Coverage, %	100.00	100.00	$\geq$95.0 (met)
$SINR_{avg}$, dB	2.772	3.492	–
$SINR_{median}$, dB	1.474	2.018	–
$SINR_{p5}$, dB	−3.678	−3.315	–
SE_{avg}, bit/s/Hz	1.741	1.915	$\geq$2.5 (not met)
Handover success, %	97.23	97.84	$\geq$98.0 (not met)
NQI	0.883	0.898	$\geq$0.82 (exceeded at baseline)

Figure 19 tracks these dynamics over all 200 training episodes: panel (a) shows that the worst-case (P5) SINR improves proportionally more than the mean SINR, indicating that the learned policy disproportionately benefits poorly served users; panel (b) shows

the ε-greedy exploration schedule reaching its floor of 0.05 by episode 119, alongside the corresponding decay in reward volatility as the policy stabilizes; and panel (c) shows the load-balancing standard deviation decreasing over the course of training, despite starting above its pre-training baseline value during the initial high-exploration phase.

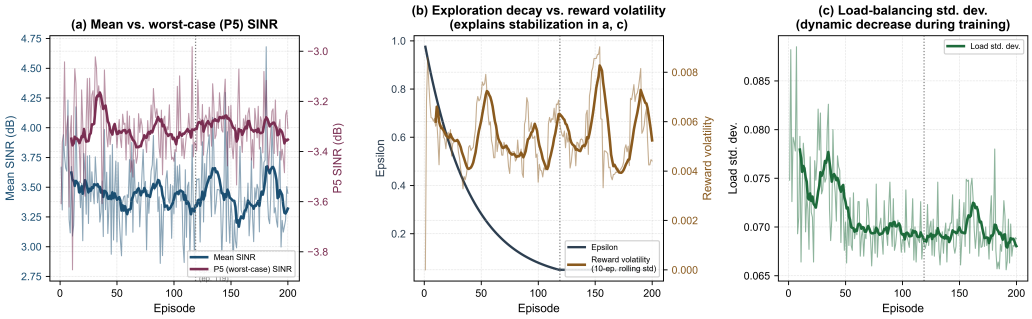


Figure 19. Divergent KPI dynamics during multi-agent Q-learning training: (**a**) mean vs. worst-case (P5) SINR; (**b**) exploration decay vs. reward volatility; (**c**) load-balancing standard deviation.

This concludes the layer-by-layer presentation of results; the following section synthesizes the physical, deployment, traffic, and SON findings into a single cross-layer pattern and a set of design guidelines.

4. Discussion

4.1. The Recurring Threshold–Saturation Pattern

Table 17 assembles the threshold/saturation evidence from the four preceding subsections into a single comparative view.

Table 17. Threshold–saturation behavior across the four framework layers.

Layer	Resource Variable	Response Variable	Threshold	Behavior Beyond Threshold
PHY	BS antenna count N_t (4→64)	SE, SINR gain	$N_t \approx$ 8–16	Marginal SE gain falls 108% → 13%
PHY	Antenna spacing d/λ	SE	$d \approx 0.25\lambda$	SE loss collapses −21.9% → −3.7%
Deployment	SNR, γ (indoor)	BLER, throughput	$\gamma \approx$ 6.5–8.0 dB	BLER falls 1.00 →≈0 within 1.5 dB
Deployment	BS count/freq. layering (rural)	Coverage + capacity	N/A—not swept	Qualitative geometric argument only; not directly measured
Traffic	Channel count (1 → 10)	SR, latency	$N_{ch} \approx 4$	Mean latency falls 31% then only 2.6% further
SON	Training episodes (ε-decay)	NQI, SINR_{avg}	Episode ≈ 118	Rapid early gain, then stabilization

Figure 20 overlays five of these six resource–response pairs—all except the rural base-station-count row, which does not admit a directly comparable single-curve treatment—on common normalized axes, with each resource and response variable min–max scaled to $[0, 1]$ (latency additionally inverted so that 1.0 always denotes best observed performance); the visual clustering of all five curves' knees within a similar region of the normalized resource axis is the graphical basis for describing this as a single recurring pattern rather than five independent, unrelated saturation effects.

The following subsection examines why four otherwise-unrelated physical mechanisms produce this same qualitative shape.

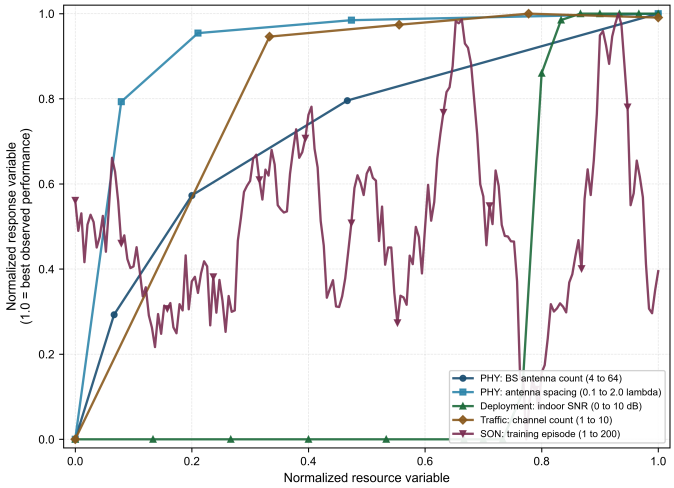


Figure 20. Recurring threshold-saturation pattern across the four framework layers, shown as normalized resource vs. response curves.

4.2. A Structural Explanation

The recurrence admits a unifying, if partial, explanation: each layer's response variable depends on its resource variable through a mechanism with an inherent nonlinear threshold. At the physical layer, the BLER waterfall is a consequence of channel-coding threshold behavior [45]; the antenna-count saturation reflects Massive MIMO channel-hardening theory [46]. At the traffic layer, the channel-count saturation reflects Erlang-type queueing behavior [47]. At the deployment layer, the rural layer's inclusion in this pattern rests on a general geometric argument rather than a directly swept station-count series: for non-overlapping circular cells, each additional station's marginal coverage-area contribution necessarily shrinks as the service area fills in [48]; this study's deployment data (Section 3.2.3) evaluates only a single four-station optimized configuration and does not itself trace this curve. At the SON layer, the episode-count saturation is a direct consequence of the ε-greedy exploration schedule [49].

These four underlying mechanisms are mathematically unrelated to one another—an information-theoretic coding-threshold effect, an array-gain/channel-hardening phenomenon, a queueing-theoretic blocking effect, and a stochastic-approximation exploration-decay effect—and none is derived from, or reducible to, any of the others. The recurring threshold–saturation shape observed here is therefore best understood as a qualitative structural regularity that these distinct mechanisms happen to share, not as evidence of a single underlying physical law governing all four layers.

4.3. Cross-Layer Dependency Chain

Figure 21 visually confirms three dependencies directly evidenced by the component results: (1) PHY $\rightarrow$ Deployment—the doubling of aggregate indoor throughput from 2×2 to 4×4 MIMO matches the ideal spatial-multiplexing gain established at the physical layer; (2) Deployment $\rightarrow$ Traffic—the channel-count sufficiency demonstrated for mMTC/URLLC coexistence is conditioned on the deployment layer provisioning that many logical channels per cell; (3) Deployment/Traffic $\rightarrow$ SON—the SON layer's power/tilt-only action space cannot compensate for insufficient base-station density or insufficient channel provisioning.

These three dependencies were evidenced only in a one-directional, post-hoc sense: a result identified in one layer's isolated simulation was shown to be numerically consistent with a subsequently examined layer's isolated simulation. What the decoupled architecture could not capture is any dynamic, closed-loop feedback between layers within a single running simulation. Specifically: (i) PHY $\rightarrow$ Deployment feedback is absent

because the deployment-layer models do not re-evaluate coverage/capacity under the antenna-correlation degradation quantified at the physical layer (Section 3.1.5); (ii) Deployment → Traffic feedback is absent because the traffic-layer channel pool is a fixed exogenous input rather than an output of the deployment-layer planning model; (iii) Traffic → SON feedback is absent because the SON agent's reward function does not observe mMTC/URLLC contention state, and so cannot learn to compensate for traffic-layer congestion; (iv) SON → PHY/Deployment feedback is absent because the trained SON policy's actions are not fed back into the physical- or deployment-layer models to re-evaluate whether the initial provisioning assumptions still hold after optimization. Each of these omitted loops is a plausible source of divergence between the guidelines derived in Section 4.4 and the behavior a genuinely co-simulated, closed-loop system might exhibit.

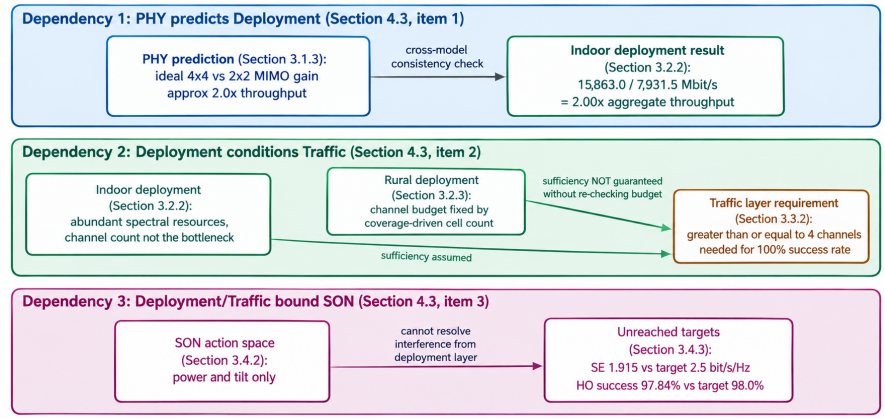


Figure 21. Cross-layer dependency evidence: quantitative confirmation of the three dependencies identified across the subsections of Section 3.

4.4. Unified Design Guidelines

Table 18 translates the three cross-layer dependencies evidenced in Figure 21 and the four-layer threshold–saturation pattern of Section 4.1 into concrete, cross-scenario planning guidance for each of the four network layers examined in this study, plus one cross-layer sequencing rule.

Table 18. Cross-scenario design guidelines.

Decision	Guideline
BS antenna array sizing	Provision antennas up to $N_t \geq$ 2–4 $\times K$; beyond that, prioritize array-spacing correctness ($\geq 0.5\lambda$)
Indoor small-cell MIMO order	Favor higher-order MIMO over additional site density when link SNR margin is already sufficient
Rural/low-density coverage planning	Use frequency-differentiated layering rather than uniform single-frequency deployment
mMTC/URLLC channel provisioning	Size the channel pool to comfortably exceed the empirically observed knee ($\approx$4 channels)
SON deployment expectation-setting	Expect the largest SON gains for worst-case links rather than already-adequate aggregate metrics
Cross-layer sequencing	Validate deployment-layer provisioning before relying on traffic- or SON-layer mitigation

These guidelines are cross-scenario in form, but their applicability is bounded by the parameter regimes and modeling choices of the underlying studies, discussed next.

4.5. Limitations

Five methodological limitations qualify the generality of this synthesis. First, the five component models were not co-simulated; Section 4.3 enumerates the specific cross-layer feedback loops this absence omits. Second, the threshold locations are specific to the parameter regimes studied in each component thesis. Third, indoor MIMO throughput scaling did not exhibit clear saturation within the narrow two-point range evaluated (2×2 vs. 4×4). Fourth, unlike the other five threshold–saturation observations in Table 17, the rural deployment layer's inclusion is based on a general geometric argument rather than a directly measured station-count sweep; the underlying rural-deployment model evaluates a fixed four-station configuration and was not run across a range of station counts. Fifth, the SON self-optimization results (Section 3.4) come from a single training run under a single configuration; an alternative configuration explored during development (differing in target spectral efficiency, episode count, and exploration-decay rate) is not reported, and no repeated-seed variance estimate is available.

5. Conclusions

This study compared five independently developed simulation studies—physical-layer MIMO characterization, dense indoor eMBB deployment, sparse rural eMBB deployment, mixed mMTC/URLLC traffic, and SON-based autonomous optimization—not by co-simulating them, but on the basis of a single, unified metric taxonomy for 5G NR performance evaluation. The unified methodology mapped five heterogeneous simulation tools, spanning four network layers, onto one consistent metric taxonomy, enabling direct comparison across layers that would otherwise remain methodologically incommensurable.

At the physical layer, spatial diversity was shown to remove the ≈11 dB Rayleigh fading penalty inherent to SISO transmission, while Massive MIMO scaling delivered a 3.5× spectral-efficiency gain with sharply diminishing marginal returns. At the deployment layer, dense indoor and sparse rural configurations were shown to represent opposite poles of a single density–capacity spectrum. At the traffic layer, mMTC/URLLC coexistence was shown to be governed by a similar threshold effect, with a fixed six-channel pool under which the URLLC latency requirement ($\leq$10 ms) remained satisfied for all successfully delivered URLLC packets across more than a threefold increase in device population, although the URLLC access success rate declined from 94.0% to 91.4%. At the autonomous-management layer, self-configuration reliably automated initial deployment, while multi-agent Q-learning-based self-optimization produced measurable but topology-dependent gains, most pronounced for worst-case link quality.

The principal original contribution of this article is the cross-scenario synthesis: none of the five constituent studies, taken individually, identifies the recurring threshold–saturation pattern, the three evidenced cross-layer dependencies, or the resulting design guidelines; each addressed only its own layer in isolation. Together, these results support a set of cross-scenario design guidelines whose applicability is bounded by the parameter regimes and modeling choices of the underlying studies.

Future Work

Three directions follow directly from the limitations identified above: (1) development of a genuinely integrated, co-simulated testbed—for example using GPU-accelerated, link-to-system platforms such as NVIDIA Sionna (NVIDIA Corporation, Santa Clara, CA, USA), which would allow the dynamic cross-layer feedback loops identified as absent in Section 4.3 to be captured directly rather than reconstructed post hoc; (2) extending the physical-layer MIMO sweep to higher-order indoor configurations (8×8 and beyond); (3)

implementing the Deep Q-Network extension [29] for the SON layer under non-uniform, "hot-spot" traffic-load scenarios, including repeated-seed runs to quantify training variance.

Author Contributions: Conceptualization, Y.P. and A.Y.; methodology, Y.P. and A.S.; software, Y.P.; validation, O.P. and A.S.; formal analysis, Y.P. and Y.Z.; investigation, A.Y., Y.Z. and Y.P.; resources, A.S.; data curation, Y.P.; writing—original draft preparation, Y.P.; writing—review and editing, A.Y. and O.P.; visualization, Y.P.; supervision, Y.P. and A.S.; project administration, A.Y. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable. This study did not involve humans or animals; all reported results were obtained exclusively through computer simulation.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data, simulation source code, and configuration files supporting the findings of this study are presented comprehensively within the article (via Tables 3–18 and Figures 6–20). The raw simulation output files (JSON/CSV/XLSX) are archived by the authors and are available from the corresponding author upon reasonable request.

Acknowledgments: During the preparation of this manuscript, the authors used Claude (Anthropic, San Francisco, CA, USA; Claude Sonnet 4.5 and Claude Sonnet 5) and ChatGPT Images 2.5 (OpenAI, San Francisco, CA, USA; free tier; used to generate the structural diagrams in Figures 1 and 21) for the following purposes: (1) drafting and language-editing of manuscript text; (2) generating Python (Matplotlib) and MATLAB-adjacent plotting scripts to visualize simulation results produced independently by the authors; (3) generating diagram source code (PlantUML 1.2026.9 and Graphviz via Viz.js 3.17.0) for architecture and workflow illustrations based on the authors' own model implementations; (4) cross-referencing and synthesizing quantitative findings across the five constituent simulation studies into the unified framework presented in Section 3. The underlying simulation models, experiments, and all numerical results were designed, implemented, and executed by the authors; the AI tools were not used to generate, alter, or interpret raw simulation data. The authors have reviewed and verified all AI-assisted outputs against the original source data and code, and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

3GPP	3rd Generation Partnership Project
5G NR	Fifth-Generation New Radio
AWGN	Additive White Gaussian Noise
BER	Bit Error Rate
BLER	Block Error Rate
BPSK	Binary Phase-Shift Keying
BS	Base Station
eMBB	Enhanced Mobile Broadband
InH	Indoor Hotspot
KPI	Key Performance Indicator
mMTC	Massive Machine-Type Communications
MRC	Maximal-Ratio Combining
NLOS	Non-Line-of-Sight
NQI	Network Quality Index
NR	New Radio
QAM	Quadrature Amplitude Modulation

SE	Spectral Efficiency
SIMO	Single-Input Multiple-Output
SINR	Signal-to-Interference-plus-Noise Ratio
SISO	Single-Input Single-Output
SNR	Signal-to-Noise Ratio
SON	Self-Organizing Network
SR	Success Rate
UMa	Urban Macro
URLLC	Ultra-Reliable Low-Latency Communications
ZF	Zero-Forcing

References

1. 3GPP. *TR 38.913: Study on Scenarios and Requirements for Next Generation Access Technologies*; Technical Report; 3GPP: Sophia Antipolis, France, 2022.
2. 3GPP. *TS 32.500: Telecommunication Management; Self-Organizing Networks (SON); Concepts and Requirements*; Technical Report; 3GPP: Sophia Antipolis, France, 2022.
3. Dala Pegorara Souto, V.; Dester, P.S.; Soares Pereira Facina, M.; Gomes Silva, D.; de Figueiredo, F.A.P.; Rodrigues de Lima Tejerina, G.; Silveira Santos Filho, J.C.; Silveira Ferreira, J.; Mendes, L.L.; Souza, R.D.; et al. Emerging MIMO Technologies for 6G Networks. *Sensors* **2023**, *23*, 1921. https://doi.org/10.3390/s23041921.
4. Al-Hilfi, H.M.T.; Kadhim, R.H.; Ioannou, I. Survey and Comparison of MIMO and Massive MIMO Evolution for 5G and 6G. In Proceedings of the 2025 Asian Conference on Communication and Networks (ASIANComNet), Zarqa, Jordan, 29–31 December 2025; pp. 1–8. https://doi.org/10.1109/ASIANComNet68615.2025.11579698.
5. Ibrahim, S.K.; Singh, M.J.; Al-Bawri, S.S.; Ibrahim, H.H.; Islam, M.T.; Islam, M.S.; Alzamil, A.; Abdulkawi, W.M. Design, Challenges and Developments for 5G Massive MIMO Antenna Systems at Sub-6 GHz Band: A Review. *Nanomaterials* **2023**, *13*, 520. https://doi.org/10.3390/nano13030520.
6. Ali, A.; Pourrostam, J. A survey on cell-free massive MIMO system with hardware impairments. *Phys. Commun.* **2025**, *72*, 102743. https://doi.org/10.1016/j.phycom.2025.102743.
7. Shi, E.; Zhang, J.; Du, H.; Ai, B.; Yuen, C.; Niyato, D.; Letaief, K.B.; Shen, X. RIS-Aided Cell-Free Massive MIMO Systems for 6G: Fundamentals, System Design, and Applications. *Proc. IEEE* **2024**, *112*, 331–364. https://doi.org/10.1109/JPROC.2024.3404491.
8. Jin, H.; Liu, K.; Lee, G.; Farag, E.J.; Zhang, M.; Zhu, D.; Zhang, L.; Onggosanusi, E.; Shafi, M.; Tataria, H. Massive MIMO evolution towards 3GPP Release 18. *IEEE J. Sel. Areas Commun.* **2023**, *41*, 1635–1654 https://doi.org/10.1109/JSAC.2023.3273768.
9. Dreifuerst, R.M.; Heath, R.W. Massive MIMO in 5G: How Beamforming, Codebooks, and Feedback Enable Larger Arrays. *IEEE Commun. Mag.* **2023**, *61*, 18–23. https://doi.org/10.1109/mcom.001.2300064.
10. Cebecioglu, B.B.; Mo, Y.K.; Dinh-van, S.; Evans, A.; Mi, D.; Higgins, M.D.; Abozariba, R.; Aneiba, A. Experimental Analysis of 5G NR for Indoor Industrial Environments. *IEEE Access* **2024**, *12*, 89310–89321. https://doi.org/10.1109/ACCESS.2024.3419011.
11. Sudhamani, C.; Roslee, M.; Tiang, J.J.; Rehman, A.U. A Survey on 5G Coverage Improvement Techniques: Issues and Future Challenges. *Sensors* **2023**, *23*, 2356. https://doi.org/10.3390/s23042356.
12. FAO; ITU. *Status of Digital Agriculture in 18 Countries of Europe and Central Asia*; Technical Report; Food and Agriculture Organization: Rome, Italy; International Telecommunication Union: Geneva, Switzerland, 2022.
13. Imam-Fulani, Y.O.; Faruk, N.; Sowande, O.A.; Abdulkarim, A.; Alozie, E.; Usman, A.D.; Adewole, K.S.; Oloyede, A.A.; Chiroma, H.; Garba, S.; et al. 5G Frequency Standardization, Technologies, Channel Models, and Network Deployment: Advances, Challenges, and Future Directions. *Sustainability* **2023**, *15*, 5173. https://doi.org/10.3390/su15065173.
14. Dorset Council. *5G RuralDorset*; Dorset Council: Dorchester, UK, 2025.
15. Djomadji, E.M.D.; Kabiena, I.B.; Mandengue, J.T.; Watching, F.; Tonye, E. Okumura Hata Propagation Model Optimization in 400 MHz Band Based on Differential Evolution Algorithm: Application to the City of Bertoua. *J. Comput. Commun.* **2023**, *11*, 52–69. https://doi.org/10.4236/jcc.2023.115005.
16. Federal Communications Commission. *Establishing a 5G Fund for Rural America*; Technical Report; Federal Register: Washington, DC, USA, 2024.
17. Sarkar, T.K.; Ji, Z.; Kim, K.; Medouri, A.; Salazar-Palma, M. A Survey of Various Propagation Models for Mobile Communication. *IEEE Antennas Propag. Mag.* **2003**, *45*, 51–82. https://doi.org/10.1109/map.2003.1232163.
18. Popovski, P.; Trillingsgaard, K.F.; Simeone, O.; Durisi, G. 5G Wireless Network Slicing for eMBB, URLLC, and mMTC: A Communication-Theoretic View. *IEEE Access* **2018**, *6*, 55765–55779. https://doi.org/10.1109/ACCESS.2018.2872781.

19. Alsenwi, M.; Tran, N.H.; Bennis, M.; Pandey, S.R.; Bairagi, A.K.; Hong, C.S. Intelligent Resource Slicing for eMBB and URLLC Coexistence in 5G and Beyond: A Deep Reinforcement Learning Based Approach. *IEEE Trans. Wirel. Commun.* **2021**, *20*, 4585–4600. https://doi.org/10.1109/TWC.2021.3060514.
20. Chandra Sekhar, R.; Singh, P. Optimization of resource allocation in 5G networks: A network slicing approach with hybrid NOMA for enhanced uRLLC and eMBB coexistence. *Int. J. Commun. Syst.* **2024**, *37*, e5928. https://doi.org/10.1002/dac.5928.
21. dos Santos, H.L.; Inacio de Souza, J.H.; Filho, J.C.M.; Abrão, T. LSTM-ACB-Based Random Access for Mixed Traffic IoT Networks. In *Proceedings of the 2022 IEEE 8th World Forum on Internet of Things (WF-IoT), Yokohama, Japan, 26 October–11 November 2022*; IEEE: Piscataway, NJ, USA, 2022; pp. 1–6. https://doi.org/10.1109/WF-IoT54382.2022.10152141.
22. Malta, S.; Pinto, P.; Fernández-Veiga, M. Optimizing 5G network slicing with DRL: Balancing eMBB, URLLC, and mMTC with OMA, NOMA, and RSMA. *J. Netw. Comput. Appl.* **2024**, *234*, 104068. https://doi.org/10.1016/j.jnca.2024.104068.
23. Klaine, P.V.; Imran, M.A.; Onireti, O.; Souza, R.D. A Survey of Machine Learning Techniques Applied to Self-Organizing Cellular Networks. *IEEE Commun. Surv. Tutor.* **2017**, 19, 2392–2431. https://doi.org/10.1109/comst.2017.2727878.
24. Moysen, J.; Giupponi, L. From 4G to 5G: Self-organized network management meets machine learning. *Comput. Commun.* **2018**, *129*, 248–268. https://doi.org/10.1016/j.comcom.2018.07.015.
25. Yanko, A.; Ponochovnyi, Y.; Bulba, E.; Hozenko, E. Influence of Diagnostics Errors on Safety: Indicators and Requirements. In Proceedings of the 2018 IEEE 9th International Conference on Dependable Systems, Services and Technologies (DESSERT), Kyiv, Ukraine, 24–27 May 2018; pp. 53–57. https://doi.org/10.1109/DESSERT.2018.8409098.
26. Mahamod, U.; Mohamad, H.; Shayea, I.; Othman, M.; Asuhaimi, F.A. Handover parameter for self-optimisation in 6G mobile networks: A survey. *Alex. Eng. J.* **2023**, *78*, 104–119. https://doi.org/10.1016/j.aej.2023.07.015.
27. Aliu, O.G.; Imran, A.; Imran, M.A.; Evans, B. A Survey of Self Organisation in Future Cellular Networks. *IEEE Commun. Surv. Tutorials* **2013**, *15*, 336–361. https://doi.org/10.1109/surv.2012.021312.00116.
28. O-RAN Alliance. *O-RAN Use Cases and Deployment Scenarios*; O-RAN Alliance, Alfter, Germany, 2023.
29. Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A.A.; Veness, J.; Bellemare, M.G.; Graves, A.; Riedmiller, M.; Fidjeland, A.K.; Ostrovski, G.; et al. Human-level control through deep reinforcement learning. *Nature* **2015**, *518*, 529–533. https://doi.org/10.1038/nature14236.
30. 3GPP. *TR 38.901: Study on Channel Model for Frequencies from 0.5 to 100 GHz*; Technical Report; 3GPP: Sophia Antipolis, France, 2022.
31. Jakes, W.C. *Microwave Mobile Communications*; Wiley-IEEE Press: New York, NY, USA, 1974.
32. Shiu, D.S.; Foschini, G.J.; Gans, M.J.; Kahn, J.M. Fading correlation and its effect on the capacity of multielement antenna systems. *IEEE Trans. Commun.* **2000**, *48*, 502–513. https://doi.org/10.1109/26.837052.
33. MathWorks. *5G Toolbox Documentation*; The MathWorks, Inc.: Natick, MA, USA, 2024.
34. 3GPP. *TS 38.214: NR; Physical Layer Procedures for Data*; Technical Report; 3GPP: Sophia Antipolis, France, 2023.
35. Hata, M. Empirical formula for propagation loss in land mobile radio services. *IEEE Trans. Veh. Technol.* **1980**, *29*, 317–325. https://doi.org/10.1109/T-VT.1980.23859.
36. Matloff, N. *Introduction to Discrete-Event Simulation and the SimPy Language*; Technical Report; University of California: Davis, CA, USA, 2008.
37. 3GPP. *TS 22.261: Service Requirements for the 5G System*; Technical Report; 3GPP: Sophia Antipolis, France, 2024.
38. 3GPP. *TS 32.501: Telecommunication Management; Self-Configuration of Network Elements*; Technical Report; 3GPP: Sophia Antipolis, France, 2022.
39. 3GPP. *TS 36.423: X2 Application Protocol (X2AP)*; Technical Report; 3GPP: Sophia Antipolis, France, 2023.
40. Buşoniu, L.; Babuška, R.; De Schutter, B. Multi-agent reinforcement learning: An overview. In *Innovations in Multi-Agent Systems and Applications*; Springer: Berlin/Heidelberg, Germany, 2010. https://doi.org/10.1007/978-3-642-14435-6_7.
41. Lohmüller, S.; Schmelz, L.C.; Hahn, S. Adaptive SON management using KPI measurements. In Proceedings of the 2016 IEEE/IFIP Network Operations and Management Symposium (NOMS), Istanbul, Turkey, 25–29 April 2016; pp. 625–631. https://doi.org/10.1109/NOMS.2016.7502868.
42. Tse, D.; Viswanath, P. *Fundamentals of Wireless Communication*; Cambridge University Press: Cambridge, UK, 2005.
43. Boriak, B.; Yanko, A.; Laktionov, O. Model of an Automated Control System for the Positioning of Radio Signal Transmission/Reception Devices. *Radioelectron. Comput. Syst.* **2024**, *2024*, 156–167. https://doi.org/10.32620/reks.2024.4.13.
44. ITU-R. *Recommendation P.833-9: Attenuation in Vegetation*; Technical Report; International Telecommunication Union: Geneva, Switzerland, 2021.
45. Richardson, T.; Urbanke, R. *Modern Coding Theory*; Cambridge University Press: Cambridge, UK, 2008.
46. Larsson, E.G.; Edfors, O.; Tufvesson, F.; Marzetta, T.L. Massive MIMO for next generation wireless systems. *IEEE Commun. Mag.* **2014**, *52*, 186–195. https://doi.org/10.1109/mcom.2014.6736761.
47. Kleinrock, L. *Queueing Systems, Volume 1: Theory*; Wiley-Interscience: New York, NY, USA, 1975.

48. MacDonald, V.H. Advanced mobile phone service: The cellular concept. *Bell Syst. Tech. J.* **1979**, *58*, 15–41. https://doi.org/10.1002/j.1538-7305.1979.tb02209.x.
49. Sutton, R.S.; Barto, A.G. *Reinforcement Learning: An Introduction*, 2nd ed.; MIT Press: Cambridge, MA, USA, 2018.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.