telecom

MDPI

Article

Joint Energy–Spectral Efficiency Trade-Offs in a Multi-Slice 5G-NR Uplink: A Link-Level Evaluation with Per-Slice PUSCH Configuration

Yahya Saeed [1,*] and Lway Abdulrazak [2]

[1] Department of Electrical Engineering, University of Sulaimani, As-Sulaymaniyah P.O. Box 334, Iraq
[2] Electrical Engineering Technical College, Middle Technical University, Baghdad 10074, Iraq; lway.abdulrazak@mtu.edu.iq
* Correspondence: yahya.saeed@univsul.edu.iq

Abstract

Network slicing allows a single 5G New Radio (NR) carrier to serve services with significantly different reliability targets. However, the cost of each slice in spectral and energy terms, once it is realized as a concrete uplink, is rarely measured. This paper reports a link-level evaluation of the joint energy efficiency (EE) and spectral efficiency (SE) trade-off in a multi-slice 5G-NR uplink in which enhanced Mobile Broadband (eMBB), Ultra-Reliable Low-Latency Communication (URLLC), and massive Machine-Type Communication (mMTC) each receive a distinct Physical Uplink Shared Channel (PUSCH) configuration matched to a block error rate (BLER) target of 10^{-3}, 10^{-5}, and 10^{-1}, respectively. Using a 3GPP-compliant simulator built on the MATLAB 5G Toolbox (R2026a), 1440 operating points across different frequency bands, FR1 and FR2, propagation delay profiles, transmit power levels, and user distances from the gNB were evaluated. Each point is evaluated twice, with and without a co-channel fixed-service (FS) interferer. Every slice exhibits an interior energy-optimal transmit power; the three slices occupy clearly separated regions of the EE–SE plane, ordered eMBB above mMTC above URLLC. Finally, FS interference reshapes the trade-off through one mechanism with very different consequences per slice: eMBB, which earns its rate from 256-QAM, loses approximately a quarter of its peak throughput and roughly half of its peak EE under line-of-sight (LoS), while mMTC is the most resilient and URLLC shows the sharpest qualitative change, with its energy optimum migrating by approximately 16 dB. Therefore, the energy-optimal operating point is not a fixed property of the band, distance, and slice, but a function of the interference that the link actually sees; thus, slice-aware uplink power control must also be interference-aware.

Keywords: 5G New Radio; network slicing; PUSCH; energy efficiency; spectral efficiency; link-level simulation; eMBB; URLLC; mMTC; interference; uplink

Check for updates

Academic Editor: Sarun Duangsuwan

Received: 5 July 2026
Revised: 18 August 2026
Accepted: 25 August 2026
Published: 1 September 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

Network slicing is the organizing principle of 5G New Radio (NR), in which a single physical network is partitioned into logical slices, each tuned to a service class with its own performance contract. IMT-2020 defines three canonical classes: eMBB, URLLC, and mMTC, which are distinguished by their reliability targets that differ by several orders of magnitude [1,2]. Multiple studies have investigated how these slices coexist; however, they have concentrated on the system and orchestration level, where a slice is an allocation of

abstract resources rather than a concrete radio link [3–7]. Far less is known about the cost of each slice at the link itself, once a real Physical Uplink Shared Channel (PUSCH) has to meet its reliability target, and energy is seldom measured alongside the interference.

This study measured these costs. A 3GPP-compliant link-level uplink simulator was built on the MATLAB 5G Toolbox, in which eMBB, URLLC, and mMTC each received a distinct PUSCH configuration, that is, modulation and coding table, time–frequency footprint, reference-signal density, and hybrid-ARQ (HARQ) policy, matched to block-error-rate (BLER) targets of 10^{-3}, 10^{-5}, and 10^{-1}, respectively. The joint EE and SE trade-off is then traced across 1440 operating points covering two carrier bands (FR1 at 3.5 GHz and FR2 at 24.8 GHz), two propagation delay profiles (CDL-C for NLoS and CDL-D for LoS), 12 transmit power levels, and 10 gNB–UE distances. Each point was evaluated twice, with and without a co-channel fixed service (FS) interferer.

Two considerations motivate the evaluation of energy and interference. First, the bands of interest are shared; fixed-service incumbents occupy parts of the C-band and mmWave range; therefore, a realistic uplink study cannot treat the link as noise-limited. Second, optimizations that chase spectral efficiency alone tend to ignore interference and the energy required to overcome [8], yet the bit-per-Joule structure of the slicing energy model [9,10] makes the transmission power a first-class design variable, rather than a quantity to be maximized.

This study makes four contributions to the literature. First, all three IMT-2020 slices are instantiated with differentiated, theoretically justified per-slice PUSCH configurations that encode the reliability diversity premise of [1] at the link level. Second, the per-slice joint EE-SE frontier is characterized, and each slice is shown to have an interior energy-optimal transmit power. Third, the manner in which a fixed-service interferer reshapes each frontier and shifts each energy optimum is quantified, and the shift is shown to be slice-dependent. Fourth, the effects of band and delay profiles were isolated using a fixed numerology, which was validated by the band-matched check in Section 4.6 and seed-variability analysis in Section 4.5. The remainder of the paper reviews related work (Section 2), describes the system model and methodology (Section 3), presents and discusses the results (Section 4), and concludes (Section 5).

2. Related Work

The work relevant to this study belongs to a few overlapping areas. One concern is the purpose of network slicing and the definitions of the three service classes. The second concern is how slices share radio resources when their numerologies and reliability targets are in conflict. The third treats energy as an explicit objective, rather than something to be spent freely. The fourth is to build detailed link-level models for the NR uplink. The fifth is the reliability and latency of URLLC. The present study draws on all five, and the gap it fills lies where they fail to meet these requirements.

2.1. Network Slicing and the Three Service Classes

The idea of treating eMBB, URLLC, and mMTC as separate slices was presented by Popovski et al. [1], who modelled the three as an uplink multiple-access problem and asked when they could share a channel without one degrading the other. Their answer turns on reliability diversity: because the services differ in target error probability by several orders of magnitude—on the order of 10^{-3}, 10^{-5}, and 10^{-1}—a non-orthogonal allocation can still keep them apart, at times outperforming a clean orthogonal split. The model is deliberately abstract; it says nothing about how a 10^{-5} target becomes a concrete choice of modulation, coding, pilot density, or retransmission policy, and energy never enters the analysis. This silence is the point of departure for the present study. The three reliability targets are

taken as given—they are the per-slice BLER targets that drive the link adaptation—and the question is what they cost in spectral and energy terms once a real PUSCH must meet them.

Three recent surveys have mapped the territory of this model. Novanana et al. [11] catalogue how access technologies such as non-orthogonal multiple access combine with spatial and frequency diversity to serve the three classes, a multi-user uplink gain formalized in the massive-MIMO and interleave-division framework of Xu et al. [12], and make a point that bears directly on the present experiments: most slicing studies assume NLoS Rayleigh fading, so the LoS regimes typical of elevated or rural deployments are comparatively neglected. Metre et al. [5] approach slicing from the resource-management side, comparing ten bandwidth-allocation algorithms for a thousand users and finding that priority-based dynamic allocation gives the best trade-off between user satisfaction and channel use; their study stops at the bandwidth-abstraction layer, with no physical-layer chain behind the allocation. Rafique et al. [6], surveying beyond-5G slicing for smart cities, reached the same boundary from another direction: their taxonomy of frameworks, orchestration, and machine-learning-driven management is thorough on architecture but treats the realization of per-slice Quality of Service (QoS) at the radio link, especially in higher bands, only in the abstract. The common thread is that the slicing literature is rich at the system and orchestration levels and thin at the level where a slice's QoS is actually won or lost: the link itself.

2.2. Slice Coexistence and Mixed Numerology

The practical sharing of slices in the spectrum has been studied most closely for the eMBB–URLLC pair, where a latency-bound service must be fitted alongside a rate-hungry service. Esmaeily et al. [13] cast this as a mixed-numerology mini-slot allocation problem with a puncturing threshold that bounds how much an intruding URLLC transmission may damage eMBB throughput; their scheduler holds eMBB users above a minimum rate while still admitting URLLC. Al-Ali and Yaacoub [14] extend the same puncturing logic with a knapsack-based scheme that selects the eMBB resource blocks least harmed by URLLC pre-emption. The link is abstracted and energy is not measured, but the design logic—short mini-slots for the latency-critical slice and a longer footprint for the broadband slice—is precisely what is encoded here at the PUSCH level through mapping type and symbol allocation, in the same spirit as the coordinated eMBB/URLLC resource allocation of Prathyusha and Sheu [15].

The closest prior study is that of Engin et al. [10]. They built a preemptive scheduler for coexisting eMBB and URLLC slices in the MATLAB 5G Toolbox: a single gNB over 10 MHz (52 resource blocks at 15 kHz spacing) serving four eMBB users and one URLLC user, a CDL-C channel at 30 ns delay spread, link adaptation, and a URLLC-specific modulation-and-coding table aimed at a 10^{-5} BLER. The toolbox, channel model, link-adaptation idea, and reliability target all overlapped with the present setup. However, they do not address the three aspects that this study is organized around: they treat two services rather than three, work mainly in the downlink rather than the uplink, and do not report energy efficiency or any contrast across frequency bands or propagation conditions. The present study can be read as the uplink, three-slice, and energy-aware extension of that simulation, with band and channel sweeps.

2.3. Energy Efficiency in Network Slicing

The clearest statement of energy as a first-class metric in slicing was made by Matthiesen et al. [9], who optimized throughput and energy efficiency together for a sliced radio access network and traced the Pareto boundary between them under per-slice QoS constraints. Their definition of energy efficiency, bits per Joule, with a denominator

that adds a power amplifier term scaled by the inverse of amplifier efficiency to a fixed circuit power, is the model adopted almost verbatim here at the link level. The total power expression used here, radiated power divided by amplifier efficiency plus circuit, baseband, and RF-chain power, is its concrete instantiation, and the interior power optimum reported later is a direct consequence of this structure. For the present purposes, the analysis lacks a real waveform; its rates are capacity expressions, not the output of a fading channel or a finite-BLER decoder. Liu et al. [16] carry the same bit-per-Joule objective into a distributed massive-MIMO system that slices eMBB and URLLC, jointly tuning the transmit power and the number of active antennas per slice.

Two surveys placed this gap in context. Donatti et al. [7] review thirty-six energy-efficiency techniques for slicing and sort them by optimization level—infrastructure, path or route, and slice operation—observing that spectrum efficiency and amplifier management recur as infrastructure-level levers, yet few of the surveyed works actually measure energy efficiency at the per-slice radio link. Liwen et al. [8], surveyed resource optimization for 6G Internet-of-Things networks, tabulated the trade-offs between energy efficiency and spectral efficiency, bandwidth, reliability, and rate, and noted that studies chasing spectral efficiency alone tend to ignore interference and latency. Collectively, the two surveys locate the present contribution fairly precisely: the underpopulated cell is the one where the per-slice spectral efficiency and amplifier inefficiency actually meet, and are measured rather than assumed. On the standards side, the 3GPP energy efficiency key performance indicator treatment in [17] supplies the units reported here—bits per Joule, with energy saving framed around amplifier activity—and argues for keeping amplifier efficiency inside the power model rather than idealizing it away.

2.4. Link-Level Modeling of the NR Uplink

Against this system-level backdrop is a smaller and more concrete literature that models PUSCH itself. The cell-free MIMO link-level study of Rahmani Ghourtani et al. [18] represents careful uplink modelling: a standard frame and slot structure, demodulation-reference-signal (DM-RS)-based channel estimation, minimum-mean-square-error equalization, and a move toward centralized estimation across distributed radio units. This shows how much the DM-RS configuration moves the channel estimation quality and, therefore, link performance, which is the reason the eMBB slice carries a denser pilot pattern than the other two slices. A complementary parametric study by Pateriya et al. [19] sweeps the PUSCH and downlink-shared-channel throughput and BLER across subcarrier spacings, modulation orders, and CDL/TDL channels, confirming the toolbox methodology on which the present simulations rely. Kabalcı and Ali [20] report a closely related PUSCH throughput-versus-SNR analysis across subcarrier spacings, modulation orders up to 256-QAM, antenna counts, and CDL/TDL channels, again in the 5G Toolbox. The interference-focused PUSCH study in [21] adds the piece that the two leave out, intra- and inter-cell interference, and quantifies how higher-order modulation buys throughput at the cost of interference sensitivity. Its configuration style is adopted for the interferer here, with the methodology making explicit that the FS contribution enters the engine as a raised effective noise floor that the decoder actually sees, through the SINR substitution of Section 3.8, so that the BLER, throughput, SE, and EE all respond to the same SINR-substitution view of multicell uplink interference, as analyzed by Wei et al. [22]. The uplink coverage enhancement schemes of Yi et al. [23] sit alongside these as link-level evidence of how the transmit power and configuration set the reachable distance. The frequency-domain spectral-shaping scheme of Nasarre et al. [24] adds the power amplifier angle, showing how a lower peak-to-average power ratio raises the feasible uplink transmit power across FR1 and FR2, the same amplifier limit that bounds the energy model used in this

study. Underlying all these choices is the standard NR reference text [25], whose account of the scalable numerology—subcarrier spacing as a power-of-two multiple of 15 kHz, with the band dictating the sensible choice—frames the FR1-versus-FR2 comparison and, by the same logic, motivates the controlled choice, made here, of holding the spacing fixed at 60 kHz in both bands, the one spacing that TS 38.104 defines for both frequency ranges (Section 3.2).

2.5. URLLC Reliability and HARQ

The reliability and latency mechanics of URLLC form a literature of their own, and three results shape the URLLC slice model used in this study. Sun et al. [26] showed that energy efficiency and URLLC reliability are not inherently at odds; by jointly setting the transmit power, bandwidth, and number of active antennas under finite block length rate expressions, they obtained the globally energy-optimal allocation subject to delay and packet loss constraints. This result makes it possible to include URLLC in an energy efficiency study, rather than treating it as a pure reliability problem. Hao et al. [27] determine how the reliability budget decomposes across control-channel detection, data decoding, and feedback, and how many HARQ retransmissions a given numerology can afford before it breaks the 1 ms latency bound—past roughly two transmissions, some numerologies simply cannot. The conservative reading of that finding is adopted here, disabling HARQ for URLLC altogether, trading coding gain for a guaranteed short transmission, which is why the URLLC slice uses a two-symbol mini-slot with no retransmission. Santos et al. [28] push in the opposite direction with a pipelined multi-user incremental redundancy HARQ scheme that recovers some of the latency retransmission costs; it marks the frontier from which the HARQ-free choice made here deliberately steps back. Varatharaajan et al. [29] reinforce this picture from the control plane, evaluating multi-TRP repetition and diversity schemes for PDCCH reliability through link-level simulation. The cross-layer URLLC survey [30] confirms the wider picture: meeting the joint reliability–latency target requires coordinated physical- and medium-access-layer design, and the techniques it catalogues are seldom evaluated alongside eMBB and mMTC in a single energy-aware framework.

2.6. Research Gap

Set side by side, the five strands above leave a specific, addressable gap: the slicing-and-energy literature is predominantly system-level and capacity-based, the link-level PUSCH literature is predominantly single-service, noise-limited, and silent on energy, and the URLLC literature deepens one slice without placing it beside the others. No prior study is known to do all of the following at once: instantiate all three IMT-2020 slices with distinct, theoretically justified PUSCH configurations; drive each to its own BLER target by adaptive modulation and coding inside a standards-compliant uplink chain; and report the resulting joint EE–SE behavior across two frequency bands and both propagation delay profiles. This study is situated at this intersection of research. Each of the four contributions listed in Section 1 addresses one element of this gap: the differentiated per-slice PUSCH configurations supply the missing three-slice instantiation, the measured EE–SE frontier supplies the missing joint metric, the interference sweep supplies the missing incumbent coexistence, and the fixed-numerology dual-band design supplies the missing band isolation. It carries the reliability targets of [1] and the bit-per-Joule model of [9] into the type of realistic PUSCH simulation exemplified by [10] and the link-level studies above, and reports per-slice EE and SE results over a 1440-point operating space. Table 1 summarizes how the reviewed studies relate to this goal.

Table 1. Positioning of the reviewed works; ✓ addressed, ○ partially addressed, ✗ not addressed.

Study	Level	3 Slices	EE–SE	Link-Level PUSCH	LoS/NLoS	FR1 + FR2	Uplink	Interf.
Popovski et al. [1]	Theory	✓	✗	✗	✗	✗	○	○
Novanana et al. [11]	Review	✓	○	✗	○	✗	✗	✗
Metre et al. [5]	System	○	✗	✗	✗	✗	✗	✗
Rafique et al. [6]	Survey	✓	✗	✗	✗	○	✗	✗
Esmaeily et al. [13]	System	○	✗	○	✗	✗	✗	○
Engin et al. [10]	Link/MAC	○	✗	✓ (DL)	✗	✗	✗	✗
Matthiesen et al. [9]	Optimal	✓	✓	✗	✗	✗	✗	✗
Donatti et al. [7]	Survey	✓	○	✗	✗	✗	✗	✗
Liwen et al. [8]	Survey	○	✓	✗	✗	✗	✗	○
3GPP-KPI study [17]	Standards	✓	○	✗	✗	✗	✗	✗
Rahmani Ghourtani et al. [18]	Link	✗	✗	✓	○	✗	✓	✗
Pateriya et al. [19]	Link	✗	✗	✓	○	○	○	✗
PUSCH interference [21]	Link	✗	✗	✓	✗	✗	✓	✓
Dahlman et al. [25]	Reference	✓	✗	✓	✗	✓	✓	○
Sun et al. [26]	Optimal	✗	✓	✗	✗	✗	○	✗
Hao et al. [27]	Analysis	✗	✗	✗	✗	✗	✗	✗
Santos et al. [28]	Scheme	✗	✗	○	✗	✗	○	✗
URLLC survey [30]	Survey	○	✗	○	✗	✗	✗	✗
This work	Link-level	✓	✓	✓ (UL)	✓	✓	✓	✓

No single prior study occupies the bottom row; the combination of three differentiated per-slice PUSCH configurations, a measured joint EE–SE frontier, and a dual-band dual-propagation uplink sweep is the contribution of this study.

3. System Model

This section describes the link-level uplink model exactly as implemented in the MATLAB 5GToolbox simulator used to produce all results. The simulator is organized as a driver that sweeps the operating space and calls a single scenario engine, which follows the structure of the 3GPP-compliant NR PUSCH throughput reference design [25]. Each subsection describes the model, lists its parameters in a table, and notes the corresponding configurations. The processing chain is summarized in Algorithm 1 at the end of this section.

3.1. Overview and Simulation Workflow

The uplink of a single 5G-NR cell is modelled, in which one gNB serves user equipment (UE) belonging to three orthogonally multiplexed network slices: eMBB, URLLC, and mMTC [1]. One-slice UE is active per engine call over its assigned resource-block (RB) set, and a single co-channel FS interferer raises the floor at the gNB.

The driver enumerates the operating space as the Cartesian product of the carrier band, slice, propagation profile, and UE transmit power; the engine then sweeps the gNB–UE distance internally and performs adaptive modulation and coding (AMC) for each distance. With two bands, three slices, two delay profiles, 12 power levels, and 10 distances, there were a total of 1440 operating points. The entire driver is executed twice to obtain the with/without interference comparison: the only change is the operating SNR fed to the decoder in the processing loop, namely the interference-free γ_{SNR} or the interference-limited γ_{SINR} defined in Section 3.8. All other parameters are held fixed, and the random-number generator is reset per operating point so that the two runs see identical channel and noise realizations and differ only through the interference term. Each accepted MCS decision is itself a Monte Carlo estimate: the block-error rate at every operating point is measured over Nf = 2 radio frames, that is, 80 consecutive slots at the 60 kHz numerology, before the loop accepts or decrements the index; for the tight eMBB and URLLC targets, this acceptance criterion amounts to requiring an error-free evaluation window, so the reliability targets are enforced through the conservative selection margin of the per-slice MCS tables rather than estimated to their nominal resolution at each point. Therefore,

holding the generator state fixed across the sweep is a common random numbers design; every pair of operating points and every with/without interference pair is compared on identical realizations, and the seed-to-seed variability of the resulting curves is quantified in Section 4.5. Table 2 lists the sweep and control parameters used in this study.

Algorithm 1 Per-slice link-level evaluation (one engine call)

Input: scenario; mode $\in$ {no-FS, FS}
Configure carrier, per-slice PUSCH, arrays, channel, path loss
Compute $\gamma_{\mathrm{SNR}}(d)$, $\gamma_{\mathrm{SINR}}(d)$ for all d
$\gamma(\cdot) \leftarrow \gamma_{\mathrm{SNR}}$ if no-FS, else γ_{SINR}
$m \leftarrow$ Starting $\mathrm{MCS}_{\mathrm{slice}}$
for each distance index $k = 1, \ldots, N_d$ **do**
  repeat
    set modulation/code rate from MCS m; reset RNG and channel
    $\sigma^2 \leftarrow$ normalized noise at $\gamma(k)$
    err $\leftarrow 0$; tot $\leftarrow 0$; $T_b \leftarrow 0$
    for $n_{slot} = 0, \ldots, N_{slot} - 1$ **do**
      update HARQ process; generate TB if new data (URLLC: single process, no retransmission)
      UL-SCH encode; PUSCH + DM-RS modulate; map ($F = I$)
      OFDM modulate; pass through CDL channel; add AWGN σ^2
      synchronize; OFDM demodulate
      channel estimate (perfect/practical); extract PUSCH REs
      MMSE equalize; PUSCH decode; UL-SCH decode $\rightarrow$ blkerr
      accumulate T_b, latency, err, tot
    end for
    $\mathrm{BLER}_k \leftarrow$ err/tot
    If $\mathrm{BLER}_k >$ terget **and** m > 1 **then** $m \leftarrow m - 1$
    end if
  until $\mathrm{BLER}_k \leq$ terget **or** m < 1
  record m, BLER_k, $\eta_{\mathrm{SE},k}$, $\eta_{\mathrm{EE},k}$, TP_k, $\bar{L}_k$
end for
Output: per-distance summary table

Table 2. Simulation Sweep and Control Parameters.

Parameter	Value
Carrier bands f_c	3.5, 24.8 GHz (FR1, FR2)
Slices	eMBB, URLLC, mMTC
Delay profiles	CDL-C (NLoS), CDL-D (LoS)
UE transmit power P_{tx}	1:2:23 dBm (12 levels)
gNB–UE distance d	25:25:250 m (10 points)
Frames per point N_f	2 (10 ms each)
Operating points per run	1440
Interference runs	γ_{SNR} (no FS)/γ_{SINR} (FS)
Reproducibility	RNG reset per operating point

3.2. Sliced Carrier and Resource-Block Partition

The carrier was configured using a normal cyclic prefix and a fixed subcarrier spacing of Δf = 60 kHz in both bands. The grid size is $N_{RB,tot}$ = 65 RBs at 3.5 GHz (FR1) and 66 RBs at 24.8 GHz (FR2), each spanning 12 subcarriers, with a slot of 14 OFDM symbols. Holding Δf constant across bands isolates the path loss and propagation effects from any change in numerology. The choice is standards-compliant rather than a simplification: 60 kHz is the only subcarrier spacing that TS 38.104 defines for both frequency ranges, and over the 50 MHz channel it yields the 65- and 66-RB grids above (Tables 5.3.2-1 and 5.3.2-2 of [31]), so the two bands occupy nearly identical bandwidths of 46.8 and 47.5 MHz. It is also a robust operating point for the delay spreads studied here: wider spacings shorten the OFDM symbol and lose block-error robustness in delay-spread, low-mobility channels, whereas narrower ones sacrifice the bandwidth scaling that FR2 exists to provide [19]. A band-matched 120 kHz check confirming that the reported conclusions are insensitive to

this choice is presented in Section 4.6, and the deployment implications are revisited in the Results and Limitations sections, respectively.

The grid is divided among the slices by the allocation vector a = [0.70, 0.20, 0.10] for [eMBB, URLLC, mMTC], with the remainder assigned to mMTC to conserve the total. Equation (1) is an internal definition of the model, chosen to conserve the total RB under integer rounding.

$$N_{\mathrm{RB}}^{\mathrm{eMBB}} = \lfloor a_1 N_{\mathrm{RB,tot}} \rfloor, \quad N_{\mathrm{RB}}^{\mathrm{URLLC}} = \lfloor a_2 N_{\mathrm{RB,tot}} \rfloor,$$
$$N_{\mathrm{RB}}^{\mathrm{mMTC}} = N_{\mathrm{RB,tot}} - N_{\mathrm{RB}}^{\mathrm{eMBB}} - N_{\mathrm{RB}}^{\mathrm{URLLC}}. \quad (1)$$

The resulting PRB sets were contiguous and disjoint, forming a single full-grid bandwidth part for each slice. Table 3 lists the carrier configuration and per-slice PRB allocation for each band, respectively.

Table 3. Carrier Configuration and Per-Slice RB Partitioning.

Parameter	FR1 (3.5 GHz)	FR2 (24.8 GHz)
Subcarrier spacing Δf	60 kHz	60 kHz
Cyclic prefix	Normal	Normal
Grid size $N_{RB,tot}$	65 RB	66 RB
Equivalent Bandwidth	50 MHz	50 MHz
Symbols per slot	14	14
Cell ID	0	0
eMBB RBs (PRB set)	45 (0–44)	46 (0–45)
URLLC RBs (PRB set)	13 (45–57)	13 (46–58)
mMTC RBs (PRB set)	7 (58–64)	7 (59–65)

3.3. Transmitter (UE) and Receiver (gNB) Configuration

The UE transmits with N_T = 2 antennas at a height of 1.5 m; the gNB receives with N_R = 16 antennas at a height of 25 m, with a noise figure of 10 dB and an antenna temperature of 290 K. The UE transmit power was swept over 1–23 dBm and represents the power delivered to the fully allocated antenna element grid. The array broadside orientations [bearing, downtilt; slant] in degrees were [0; 0; 0] at the UE and [0; 10; 0] at the gNB (a 10° downtilt). Table 4 summarizes these settings.

Table 4. Transmitter (UE) and Receiver (gNB) parameters.

Parameter	UE (Tx)	gNB (Rx)
Antennas	N_T = 2	N_R = 16
Height	1.5 m	25 m
Transmit power	1–23 dBm	–
Noise figure F	–	10 dB
Antenna temperature T_a	–	290 K
Array orientation [°]	[0; 0; 0]	[0; 10; 0]

3.4. gNB and UE Antenna Array Geometry

Because the fading channel is a CDL model, the scalar antenna counts are expanded into physical uniform planar panel arrays, invoked in the uplink direction [25]. A panel is described by the arrangement [M, N, P, M_g, N_g], where M rows and N columns of cross-polarized elements per panel ($P!\in!1$, 2 polarizations), arranged in an $M_g \times N_g$ grid of panels. The gNB (base-station) side selects the layout from a lookup indexed by

$$\mathrm{row} = \min\left(1 + \lceil \log_2 N_{\mathrm{BS}} \rceil,\ 11\right), \quad (2)$$

The panel-sizing rule in (2) follows the antenna-array conventions of the NR link-level examples [25].

Therefore, the total element count is $M!{\cdot}!N!{\cdot}!P!{\cdot}!M_g!{\cdot}!N_g$. For $N_{BS} = 16$, these yields are [2, 4, 2, 1, 1], that is, a single dual-polarized 2 × 4 panel of 16 elements. The UE side uses $[\lceil N_{UE}/2 \rceil, 1, 2, 1, 1]$ for $N_{UE} > 1$, providing a single cross-polarized pair (two elements) for $N_{UE} = 2$. The same UE-side rule builds a 4-element FS interferer array. The element spacing defaults to 0.5λ, and the gNB panel spacings are set to $M!{\cdot}!d_{row}$ and $N!{\cdot}!d_{col}$ to prevent panel overlap. Because the CDL default is a downlink, the transmit and receive ends are swapped such that the UE is the transmit array and gNB is the receive array. Table 5 lists the resulting geometry.

Table 5. Uplink Antenna Array Geometry.

Node	Requested	Layout [*M*, *N*, *P*, M_g, N_g]	Elements
gNB (Rx)	16	[2, 4, 2, 1, 1]	16
UE (Tx)	2	[1, 1, 2, 1, 1]	2
FS interferer	4	[2, 1, 2, 1, 1]	4

3.5. Per-Slice PUSCH Configuration

All slices shared a common PUSCH/UL-SCH baseline configuration: single-layer ($\nu = 1$), non-codebook transmission with one antenna port, no transform precoding, group- and sequence-hopping enabled, a single front-loaded DM-RS symbol (the demodulation reference signal, i.e., the pilot from which the receiver estimates the channel), 16 HARQ processes [32], transport-block (non-CBG) transmission, and LDPC decoding [33] using the normalized min-sum algorithm with at most six iterations. With HARQ enabled, the redundancy version (RV, the segment of the encoder output sent in a given transmission) sequence is [0, 2, 3, 1]; otherwise, a single transmission at RV = 0 is used. Table 6 lists the shared baselines.

Table 6. Shared PUSCH/UL-SCH Baseline (all slices).

Parameter	Value
Transmission scheme	non-codebook
Layers ν/antenna ports	1/1
Transform precoding	disabled
DM-RS length (front-loaded)	1 symbol
Group/sequence hopping	enabled
HARQ processes	16 (eMBB, mMTC); 1, retransmissions disabled (URLLC)
RV sequence (HARQ on/off)	[0, 2, 3, 1]/0 (eMBB, mMTC); RV = 0 single shot (URLLC)
CBG transmission	disabled (TB-based)
Rate-matching overhead X_{oh}	0
LDPC decoder	normalized min-sum, 6 iterations
RNTI/NID	1/0

The slices then differ through a switch that sets the reliability target, MCS table, time–frequency footprint, reference signal density, HARQ policy, and channel estimation mode (Table 7). The three configurations encode the reliability diversity principle [1] at the link level. eMBB target rate: the 256-QAM table, a full 14-symbol slot (mapping Type A), PT-RS, and two additional time-domain DM-RS positions, with HARQ and the largest RB share; its richer pilot footprint comes from the additional time positions and PT-RS rather than from the DM-RS configuration type, since configuration Type 1 carries the higher per-RB frequency-domain pilot density than Type 2 [18]. eMBB is also the only slice that uses practical (DM-RS-based) channel estimation, whereas URLLC and mMTC use perfect estimation. URLLC targets low latency at high reliability: a two-symbol mini-slot (mapping Type—a mini-slot occupies only a few OFDM symbols rather than the full 14-symbol slot—B), sparse DM-RS with no additional positions, and HARQ disabled to avoid retransmission delay—a deliberately conservative choice, since at 60 kHz, two or more transmissions would still fit a 1 ms budget [27].

Table 7. Per-Slice PUSCH Parameter Setting.

Parameter	eMBB	URLLC	mMTC
Target BLER	10^{-3}	10^{-5}	10^{-1}
MCS table	QAM256	QAM64	QAM64-lowSE
Starting MCS index	27	28	28
Mapping type	A (slot)	B (mini-slot)	A (slot)
Symbol allocation	[0, 14]	[0, 2]	[0, 14]
DM-RS config.\type	2	1	1
DM-RS add.\positions	2	0	0
CDM groups w/o data	2	1	1
Type-A DM-RS position	2	default	default
PT-RS	on	off	off
HARQ	on	off	on
Channel estimation	practical	perfect	perfect

In implementation terms, the URLLC slice therefore runs with a single HARQ process, and retransmissions are forbidden: every transport block is carried by one two-symbol transmission at RV = 0, and the redundancy-version sequence of Table 6 is inert for this slice. Forgoing retransmission gives up the effective SNR accumulation that incremental redundancy provides, up to approximately 3 dB for a second transmission at an equal per-transmission SNR [19]. The AMC loop absorbs this margin by settling on a lower index (Section 4.2); therefore, the cost appears as a reduced spectral efficiency rather than a missed reliability. The 35.7 µs two-symbol mini-slot keeps single-shot transmission feasible within the 1 ms budget at every operating point, and where no index meets the 10^{-5} target, the point is reported as an outage rather than silently relaxed. The mMTC targets robustness at relaxed reliability (10^{-1}) using a low-SE 64-QAM table, the smallest footprint, and HARQ.

3.6. Propagation Channel: Path Loss and Fading

The propagation model combines distance-dependent large-scale path loss with small-scale CDL fading. Large-scale attenuation uses the 3GPP TR 38.901 Urban-Macro (UMa) model [34] via nrPathLoss, with an environment height of 1 m, evaluated between the gNB at $p_R = (0, 0, hR)$ and the UE at $p_T = (d, 0, h_T)$ [25]. The LoS flag required by the path-loss model is inferred from the fading channel itself: the first-cluster Rician K-factor reported by the constructed channel determines whether the link is treated as a LoS. Small-scale fading uses the CDL profiles of TR 38.901 with a 300 ns delay spread and a 5 Hz maximum Doppler shift; CDL-C represents NLoS, and CDL-D represents LoS. The channel object (nrCDLChannel) is built with the uplink array geometry of Section 3.4, with the transmit and receive ends swapped, and the channel response output set to the OFDM response such that perfect channel and timing estimates are available when necessary. The carrier wavelength is $\lambda = c/f_c$. Table 8 lists the channel parameters used in this study.

Table 8. Propagation Channel Parameters.

Parameter	Value
Path-loss model	TR 38.901 UMa (`nrPathLoss`)
Environment height	1 m
LoS determination	first-cluster Rician K-factor
Fading model	CDL (`nrCDLChannel`)
Delay profiles	CDL-C (NLoS), CDL-D (LoS)
Delay spread	300 ns
Max.\ Doppler shift	5 Hz
Channel response output	OFDM response
Link direction	uplink (Tx/Rx swapped)

3.7. Interference and Noise Model

A single co-channel FS (fixed-service, denoting the licensed point-to-point incumbent, not "frequency-selective") interferer—representing the point-to-point fixed-wireless incum-

bents that share the C-band and mmWave ranges—transmits at 43 dBm from 1 km with a 4-element array over a 50 MHz channel, with its broadside oriented back toward the gNB at ([180; 0; 0]). The interferer reaches the gNB through an independent CDL-D channel and free-space path loss $L_I = FSPL\ (1, km, \lambda)$. The receiver noise chain converts the noise figure and antenna temperature to an equivalent temperature and per-antenna noise power as follows:

$$F = 10^{F_{\mathrm{dB}}/10}, \qquad T_{\mathrm{eq}} = T_a + 290\ (F - 1), \tag{3}$$

$$N_0 = \sqrt{1/2\ k_B\ B\ T_{\mathrm{eq}}}, \qquad P_N = 2N_0^2, \tag{4}$$

where k_B is Boltzmann's constant and B is the waveform sample rate. With F_{dB} = 10 dB and T_a = 290 K, T_{eq} = 290 K. Equations (3) and (4) are the standard receiver noise-figure and thermal noise relations, respectively [25]. The received FS interference power and the combined floor are

$$P_{I,\mathrm{dB}} = (P_{\mathrm{tx},I} - 30) - L_I, \quad P_I = 10^{P_{I,\mathrm{dB}}/10} \tag{5}$$

$$P_{N+I} = P_I + P_N \tag{6}$$

Equation (5) converts the transmit power of the incumbent through its free-space path loss, and (6) is the additive interference-plus-noise floor used by SINR-substitution link analyses [21,22]. Table 9 lists the interferer and noise parameters.

Table 9. Interferer and Receiver Noise Parameters.

Parameter	Value
Interferer transmit power	43 dBm
Interferer distance	1000 m
Interferer antennas	4
Interferer height	25 m
Interferer channel bandwidth	50 MHz
Interferer channel/orientation	CDL-D/[180; 0; 0]
Interferer path loss	FSPL at 1 km
Boltzmann constant k_B	1.38×10^{-23} J/K
Equivalent noise temp.\ T_{eq}	290 K
Noise bandwidth B	waveform sample rate

The single fixed-service transmitter, placed broadside to the gNB array at a fixed distance, is a deliberate worst-case geometry rather than an arbitrary simplification. Because the interference reaches the receiver only through the combined floor of (6), any additional co-channel incumbent adds non-negatively to the same scalar term: further interferers, or a closer or stronger one, raise PN + I monotonically and push every reported trend further in the direction already shown, without introducing a new mechanism. Therefore, the single incumbent results bound the per-slice sensitivity ordering from the floor side, and the migration of the energy optima reported in Section 4 is the worst case for this geometry. The multi-interferer extension is outlined as future work in Section 5.

3.8. Received SNR, SINR, and the Interference-Aware Decode

Using the FFT occupancy $\rho = 12, N_{RB,tot}/N_{FFT}$, the average received SNR and SINR per resource element (dB) at distance d and transmit power P_{tx} are

$$\gamma_{\mathrm{SNR}} = (P_{\mathrm{tx}} - 30) - L(d) - 10\log_{10}\rho - 10\log_{10}P_N \tag{7}$$

$$\gamma_{\mathrm{SINR}} = (P_{\mathrm{tx}} - 30) - L(d) - 10\log_{10}\rho - 10\log_{10}P_{N+I}, \tag{8}$$

where *L(d)* is the UMa path loss. These follow the 5G-Toolbox link-level SNR convention [25]. The decoder operates at a single chosen value γ: the interference-free run uses $\gamma = \gamma_{SNR}$ (7) and the FS run uses $\gamma = \gamma_{SINR}$ (8), so the FS power enters as a raised effective

noise floor, and the block-error rate, throughput, SE, and EE all respond to it, following the SINR-substitution principle of the multi-cell uplink analysis [22] and an interference-aware PUSCH study [21]. Inside the loop the chosen γ sets the normalized AWGN level applied to the received waveform and used as the perfect noise estimate,

$$N_0^{\text{loop}} = \frac{1}{\sqrt{N_R\, N_{\text{FFT}}\, 10^{\gamma/10}}}, \qquad \sigma^2 = (N_0^{\text{loop}})^2 N_{\text{FFT}}, \tag{9}$$

where N_R and N_{FFT} normalizations match those applied by the OFDM modulator and the channel model. A diagnostic table produced per configuration report, for each distance, the columns Signal (dB), Interference (dB), Noise (dB), SNR (dB), and SINR (dB) defined by (7) and (8); representative FR1/CDL-D values at 23 dBm span $\gamma_{SNR} \approx 19$ *to* 41dB and $\gamma_{SINR} \approx -7$ *to* 41 dB across 25 to 250 m, the gap quantifying the FS impact.

3.9. Efficiency and Latency Metrics

The spectral efficiency is the successfully delivered information bits T_b over the observation time $T_{obs} = N_f \cdot 10$ ms, normalized to the full carrier bandwidth $W = N_{RB,\,tot} \cdot 12 \cdot \Delta f$:

$$\eta_{\text{SE}} = \frac{T_b}{T_{\text{obs}}\, W} \quad [\text{b/s/Hz}]. \tag{10}$$

Equation (10) is the standard delivered-bits normalization of the spectral efficiency [25]. Normalizing to the full grid means that a slice's reported SE is its share of the total carrier SE, which is why the narrow URLLC and mMTC footprints yield a small absolute SE. The total consumed power follows the bit-per-Joule model [9,10], adding a PA term scaled by the inverse amplifier efficiency to the static circuit, baseband, and RF-chain power; the bandwidth-normalized energy efficiency is

$$P_{\text{tot}} = \frac{P_{\text{tx,lin}}}{\eta_{\text{PA}}} + P_{\text{circ}} + P_{\text{bb}} + P_{\text{rf}}, \tag{11}$$

$$\eta_{\text{EE}} = \frac{\eta_{\text{SE}}}{P_{\text{tot}}} \quad [\text{b/J/Hz}], \tag{12}$$

with $P_{tx,lin} = 10^{(P_{tx}-30)/10}$ W. The conventional bit-per-Joule figure is recovered by multiplying η_{EE} by W. Numerically, ηPA = 0.35 and Pcirc + Pbb + Prf = 0.5 W (Table 10), and W is the full carrier grid of Table 3, 46.8 MHz in FR1 and 47.5 MHz in FR2, rather than the RBs allocated to the slice; therefore, (12) prices all three slices in a common bits-per-Joule-per-Hertz currency. The shape of P_{tot} produces an interior EE optimum in the transmit power, as analyzed in the Results section. The slot duration and latency metrics are

$$\tau = \frac{1\ \text{ms}}{\Delta f / 15\ \text{kHz}}, \quad \bar{L} = \frac{\sum_i \updownarrow_i}{N_{\text{succ}}}\, \tau, \quad L_{\text{eff}} = \frac{\bar{L}}{1 - \text{BLER}}, \tag{13}$$

where ell_i is the per-block latency in slots; at 60 kHz, τ = 0.25 ms. The slot-duration rule in (13) is the NR scalable-numerology relation [25], and the latency terms are defined directly over the decoded blocks. Table 10 lists the power model constants.

3.10. Adaptive Modulation and Coding, and the Processing Loop

During each simulation sweep, the simulator executed an AMC loop [35] for every slice, employing various parameters, such as the carrier frequency, transmit power, UE distance, and delay profile at each operating point. The MCS index is initialized once to the slice starting index and adapts monotonically downward across the distance sweep. Because increasing distance only worsens the channel, the accepted index is non-increasing,

and at each distance, the loop decrements the index until the measured BLER meets the slice target or the minimum index is reached. The inner per-slot chain is the standard NR PUSCH transmit–receive pipeline [25], with the interference-aware noise of (9) injected at the channel output as follows: Algorithm 1 summarizes the engine; the per-distance outputs (MCS, BLER, throughput, SE, EE, and latency) are assembled into a summary table, and the driver aggregates the full operating space.

Table 10. Circuit-Power Model and Metric Constants.

Parameter	Value
PA efficiency η_{PA}	0.35
Circuit power P_{circ}	0.1 W
Baseband power P_{bb}	0.2 W
RF-chain power P_{rf}	0.2 W
Static power $P_0 = P_{circ} + P_{bb} + P_{rf}$	0.5 W
Bandwidth W (FR1/FR2)	46.8/47.5 MHz
Slot duration τ (60 kHz)	0.25 ms

4. Results and Discussion

The link-level performance of the three slices is reported over the full operating space defined in Section 3. Each slice is examined on its own terms, and the section closes with a cross-slice comparison of the findings. Throughout, two runs are contrasted that differ only in whether the FS interferer is folded into the decoded waveform: the interference-free run, decoded at γ_{SNR}, and the interference-limited run, decoded at γ_{SINR}. Four results were extracted from each slice. The throughput against distance shows raw coverage; the AMC index against distance shows which modulation order the link can actually sustain; energy efficiency and spectral efficiency against transmit power expose the interior energy optimum; and the EE–SE frontier collects every operating point into the joint trade-off that the study sets out to characterize. In all transmit power plots, the marker labelled *optEE* indicates the power that maximizes EE for the corresponding distance, and CDL-C and CDL-D denote the NLoS and LoS profiles, respectively.

Two features of the metric definitions help interpret the figures. The spectral efficiency is normalized to the full carrier grid (Section 3.9); therefore, each slice reports its share of the total carrier SE rather than a per-allocation figure. This is why the absolute SE and EE of URLLC and mMTC are well below those of eMBB, even when their links are healthy. The energy efficiency follows the bit-per-Joule structure of [9,10], whose fixed circuit term and linear power amplifier term together produce an interior EE optimum in the transmit power. The optimum and the way interference moves it is the recurring thread of this study.

4.1. eMBB Slice

Figures 1–8 report the eMBB results: throughput (Figures 1 and 2), the AMC-selected modulation order (Figures 3 and 4), the energy and spectral efficiency against transmit power (Figures 5 and 6), and the EE–SE frontier (Figures 7 and 8), in each case without and with FS interference. The eMBB slice carries the broadband load: the 256-QAM table, a full 14-symbol slot, and the largest resource block share. Correspondingly, its throughput was the highest of the three. Without interference (Figure 1), the LoS link at full power is close to its ceiling across the entire distance range in FR1, reaching approximately 184 Mbps at short ranges, whereas the NLoS curves decrease steeply with distance, and the FR2 band lags behind FR1 because of its higher path loss. The AMC view (Figure 3) explains the shape: at a short range, the loop sustains the 256-QAM region, and the selected order steps down with distance as the received SNR decreases, eventually reaching QPSK or an outage.

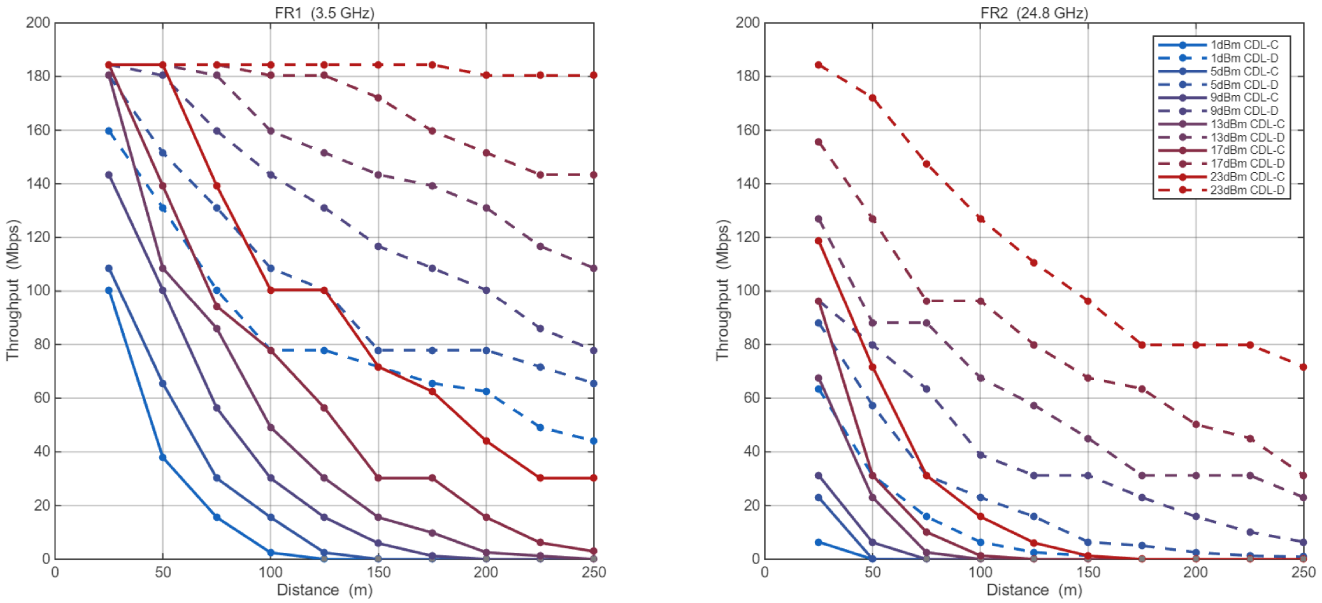


Figure 1. eMBB throughput versus distance in FR1 (3.5 GHz) and FR2 (24.8 GHz) for all transmit power and delay profiles, without FS interference. The LoS curve holds near its peak of ≈ 184 Mbps across the distance range.

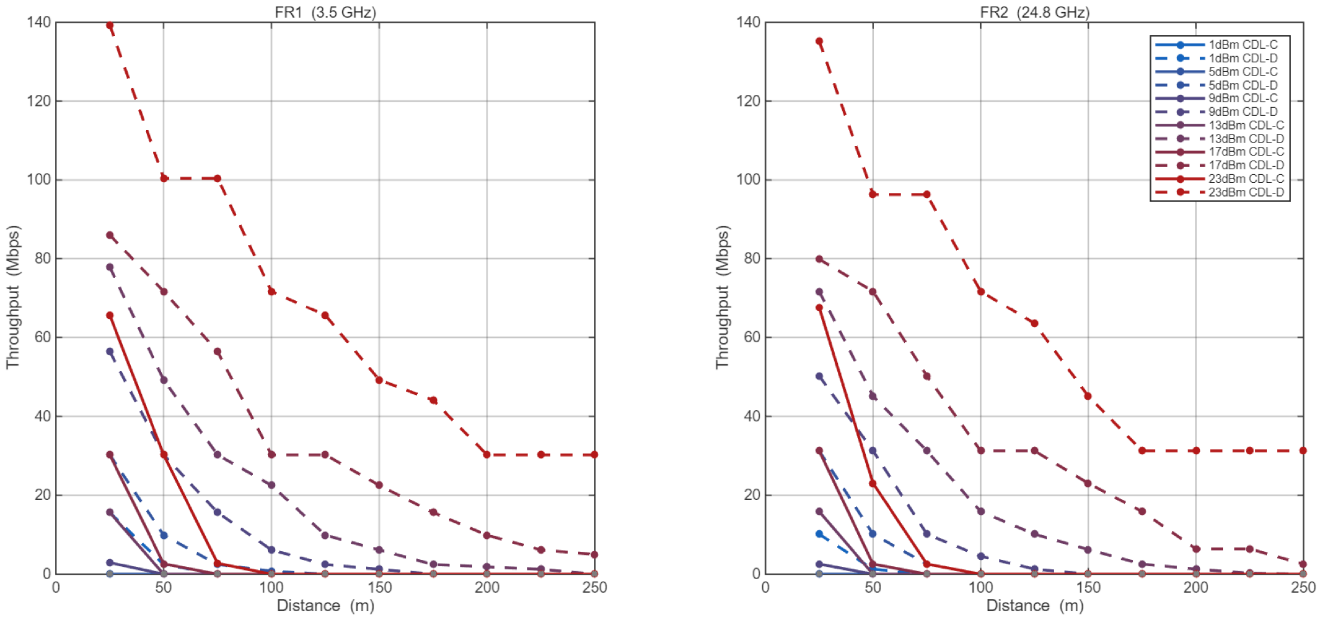


Figure 2. eMBB throughput versus distance in FR1 (3.5 GHz) and FR2 (24.8 GHz) for all transmit powers and delay profiles, with FS interference. The peak falls to ≈139 Mbps and the LoS plateau was removed.

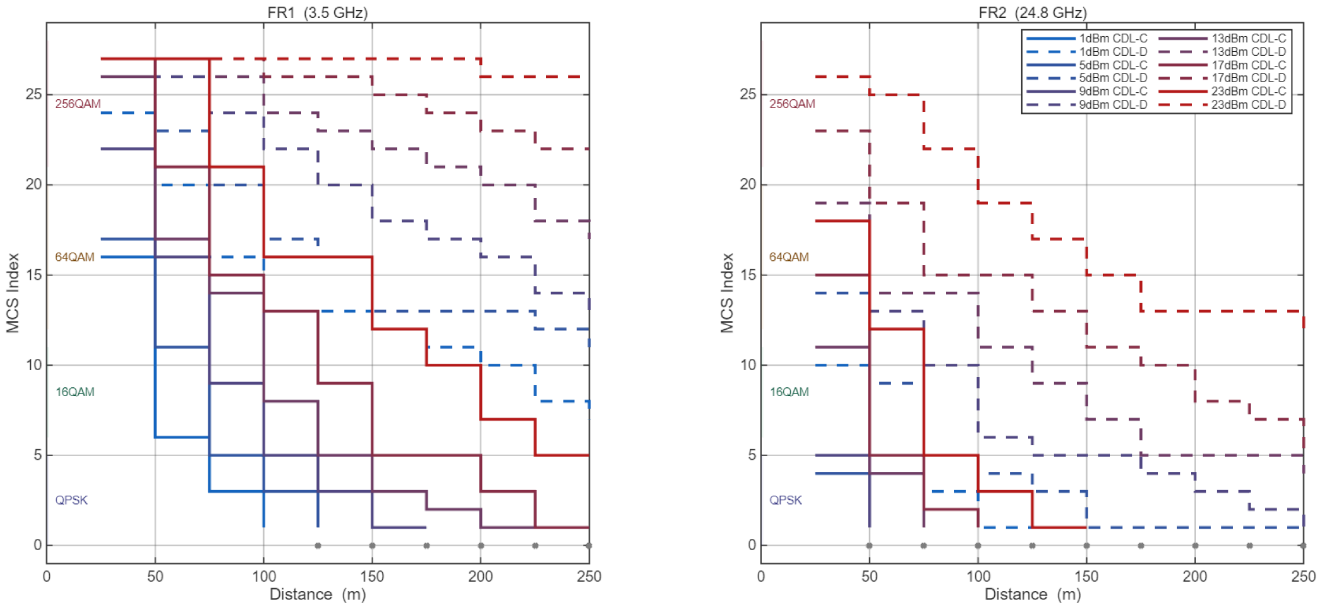


Figure 3. eMBB AMC-selected MCS index versus distance, without FS interference. The loop holds the 256-QAM region at a short range. Gray dots at MCS index 0 denote operating points in outage, where no MCS index met the target BLER.

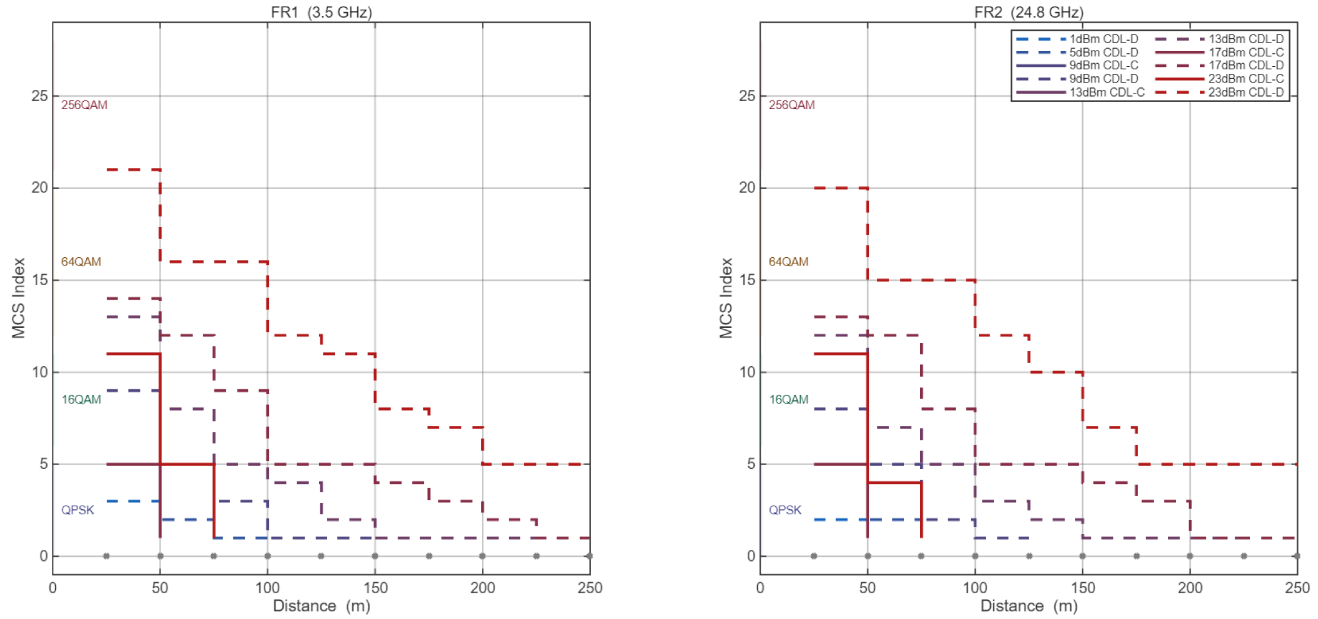


Figure 4. eMBB AMC-selected MCS index versus distance, with FS interference. The sustainable order is capped in the 64-QAM region and falls to QPSK sooner. Gray dots at MCS index 0 denote operating points in outage, where no MCS index met the target BLER.

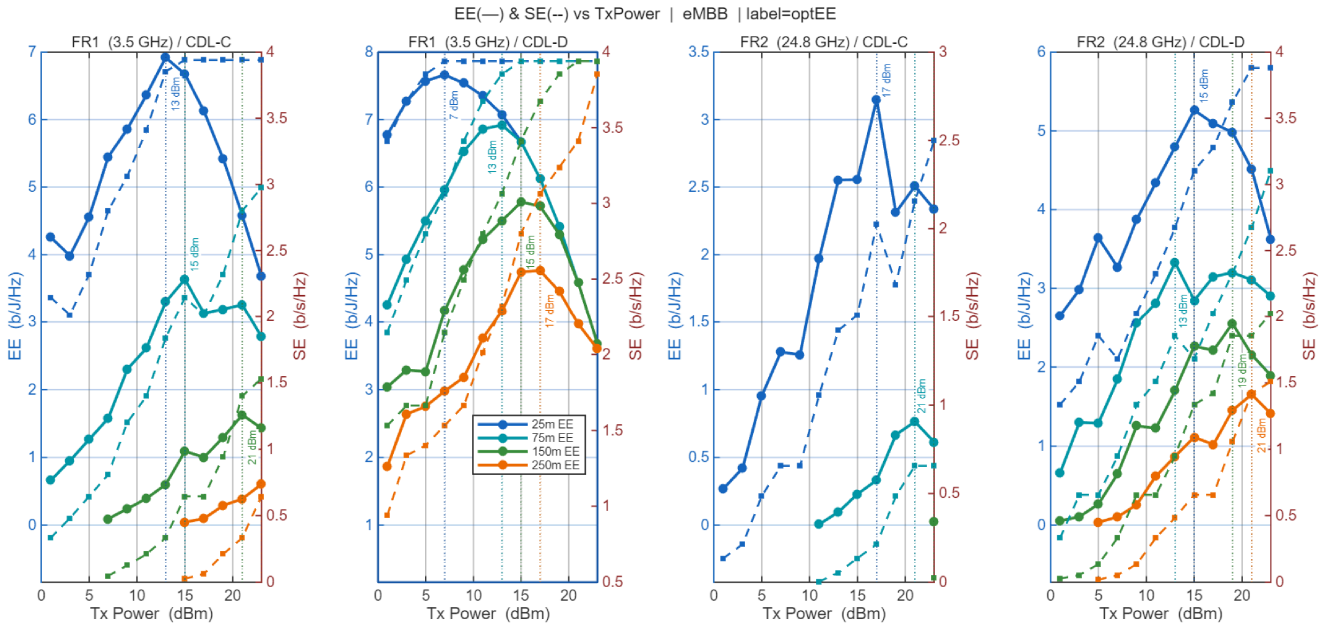


Figure 5. eMBB energy efficiency (solid, left axis) and spectral efficiency (dashed, right axis) versus UE transmit power, per band and delay profile, without FS interference, with the EE-optimal power marked. The vertical dotted line in each panel marks the EE-optimal transmit power, annotated in dBm.

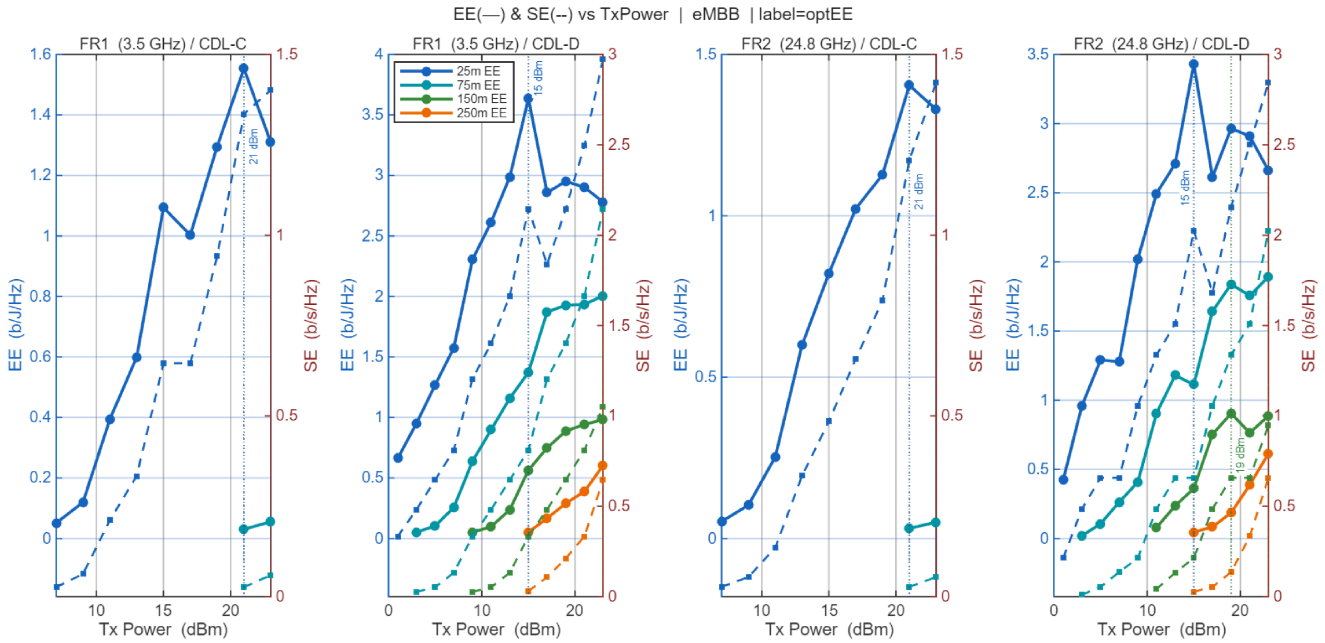


Figure 6. eMBB energy efficiency (solid, left axis) and spectral efficiency (dashed, right axis) versus UE transmit power, per band and delay profile, with FS interference. Interference roughly halves the peak EE under LoS and shifts the EE-optimal power toward the power limit. Both CDL-C panels contain fewer curves because many long-range operating points fall into outage under the raised interference floor; the reduced curve count is a physical outcome, not a plotting omission. The vertical dotted line in each panel marks the EE-optimal transmit power, annotated in dBm.

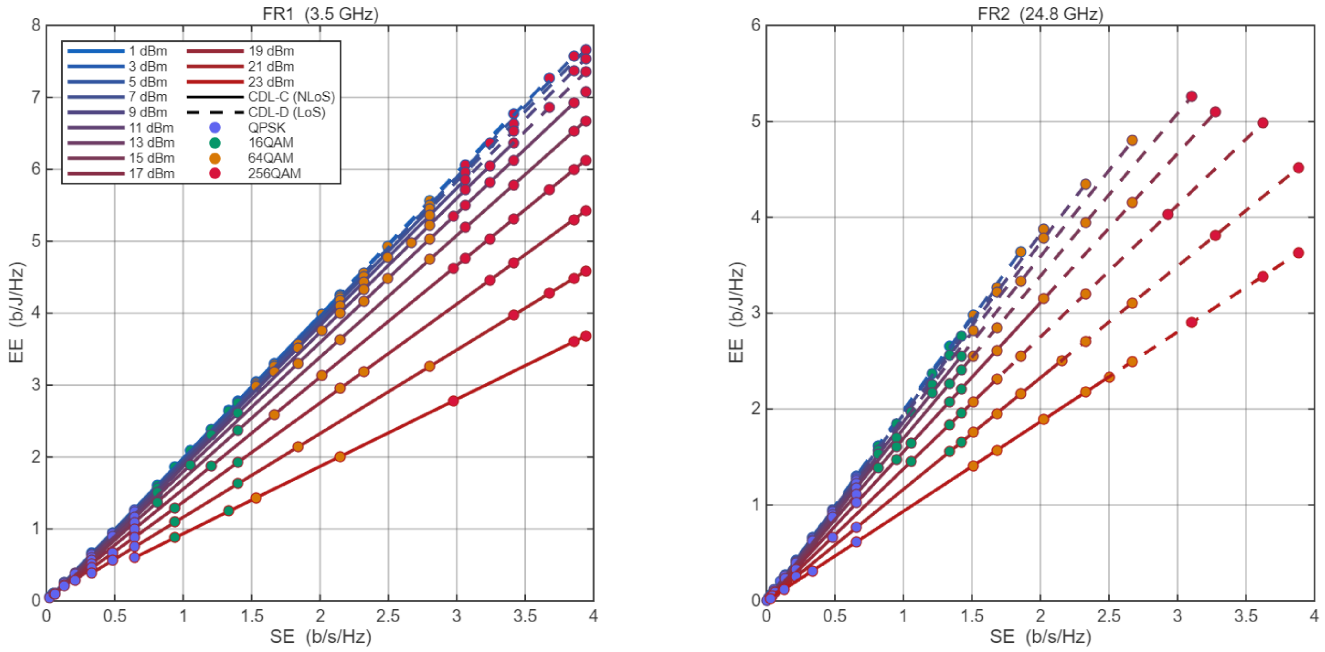


Figure 7. eMBB EE–SE frontier without FS interference, colored by transmit power and marked by modulation order.

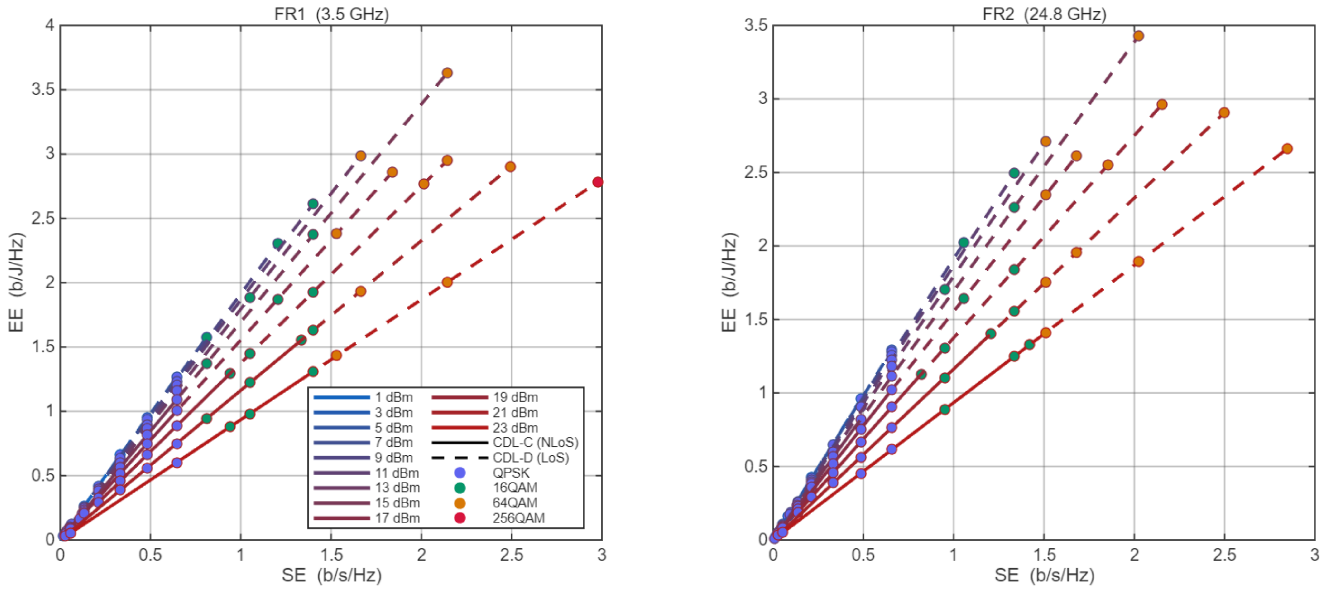


Figure 8. eMBB EE–SE frontier with FS interference, colored by transmit power and marked by modulation order. Interference contracts the region and removes the 256-QAM operating points at the high-SE point.

The folding of the FS interferer compresses every part of the results. The peak throughput falls to approximately 139 Mbps (Figure 2), a reduction of approximately one quarter, in line with the interference sensitivity of higher-order modulation reported for dense intra- and inter-cell PUSCH operations [21,22]. The LoS curve that previously held its ceiling now decays with distance, and the AMC loop is confined to the 64-QAM region, collapsing to QPSK at shorter ranges, where it carried 256-QAM (Figure 4). The mechanism is the raised noise-plus-interference floor: the effective γ_{SINR} at which each block must be decoded is lower; therefore, the highest order meeting the 10^{-3} target is lower.

The energy view further sharpens this conclusion. Without interference (Figure 5), every distance shows an interior EE optimum: under LoS in FR1, the short-range optimum

sits near 7 dBm with a peak around 7.7 b/J/Hz, migrating toward 13–17 dBm at longer range, while the NLoS and FR2 cases peak later and lower. Pushing beyond the optimum buys spectral efficiency but spends energy faster than it returns; thus, the EE turns over while the SE continues to increase. Under interference (Figure 6), the peak EE is roughly halved under LoS to approximately 3.6 b/J/Hz and decreases significantly further under NLoS, where the FR1 short-range peak drops from approximately 6.9 to 1.6 b/J/Hz. The optimum also moves outwards, to 15–23 dBm, and at longer ranges and in FR2, the interior optimum disappears as the curve climbs monotonically toward the 23 dBm ceiling. The practical reading is that the "operate below maximum power" guidance that holds in the noise-limited case weakens under interference: more radiated power is needed before the AMC saturates against the raised floor, so the energy-optimal operating point is interference-dependent.

The EE–SE frontier (Figures 7 and 8) combines these points into a single graph. Each transmit power traces a near-linear ray whose slope is the EE returned per unit of SE; the upper-left envelope, formed by the lower powers, is the efficient boundary, and the high-power rays extend to higher SE only by sacrificing the EE. Interference shrinks the attainable region on both axes and strips away the 256-QAM operating points that define the high-SE tip of the frontier, leaving a region dominated by the QPSK, 16-QAM, and 64-QAM modulation schemes. Therefore, the slice that gains the most from high-order modulation is the one that loses the most in terms of relative interference.

4.2. URLLC Slice

Figures 9–16 report the URLLC results in the same order: throughput (Figures 9 and 10), the AMC-selected modulation order (Figures 11 and 12), the energy and spectral efficiency (Figures 13 and 14), and the EE–SE frontier (Figures 15 and 16). The URLLC slice is built for latency and reliability rather than rate: a two-symbol mini-slot, a 10^{-5} target, and HARQ disabled. Therefore, its absolute throughput is small, peaking near 5 Mbps (Figure 9), and its share of the full-carrier SE is the lowest of the three because it occupies only two of the 14 symbols. Without interference, the LoS link at full power holds this peak across the entire distance range; the NLoS links fall away quickly, and FR2 lags behind FR1.

Interference affects URLLC differently from eMBB, and the difference is significant. Because the slice already operates at a modest modulation order on a short allocation, its short-range LoS peak throughput is largely retained at approximately 5 Mbps (Figure 10); the interference removed is reach, with the LoS curve now decaying steadily with distance instead of holding flat, and the NLoS links collapsing within the first few tens of meters. The energy graphs (Figures 13 and 14) show the most striking behavior in this study. Without interference, the short-range LoS optimum sits at the lowest transmit power tested, approximately 1 dBm, because the tiny allocation saturates its spectral efficiency almost immediately, and any further power only erodes EE. Under interference, the optimum point jumps to approximately 17 dBm: the link now requires substantial power simply to reach a decodable γ_{SINR}, so the EE curve rises across almost the entire power range before peaking. The peak EE falls from approximately 0.22 to 0.17 b/J/Hz under LoS and by more than half under NLoS conditions. This upward migration of the optimum by approximately 16 dB is the clearest evidence that energy-optimal uplink power control for the small-footprint slices cannot be set without knowledge of the interference environment.

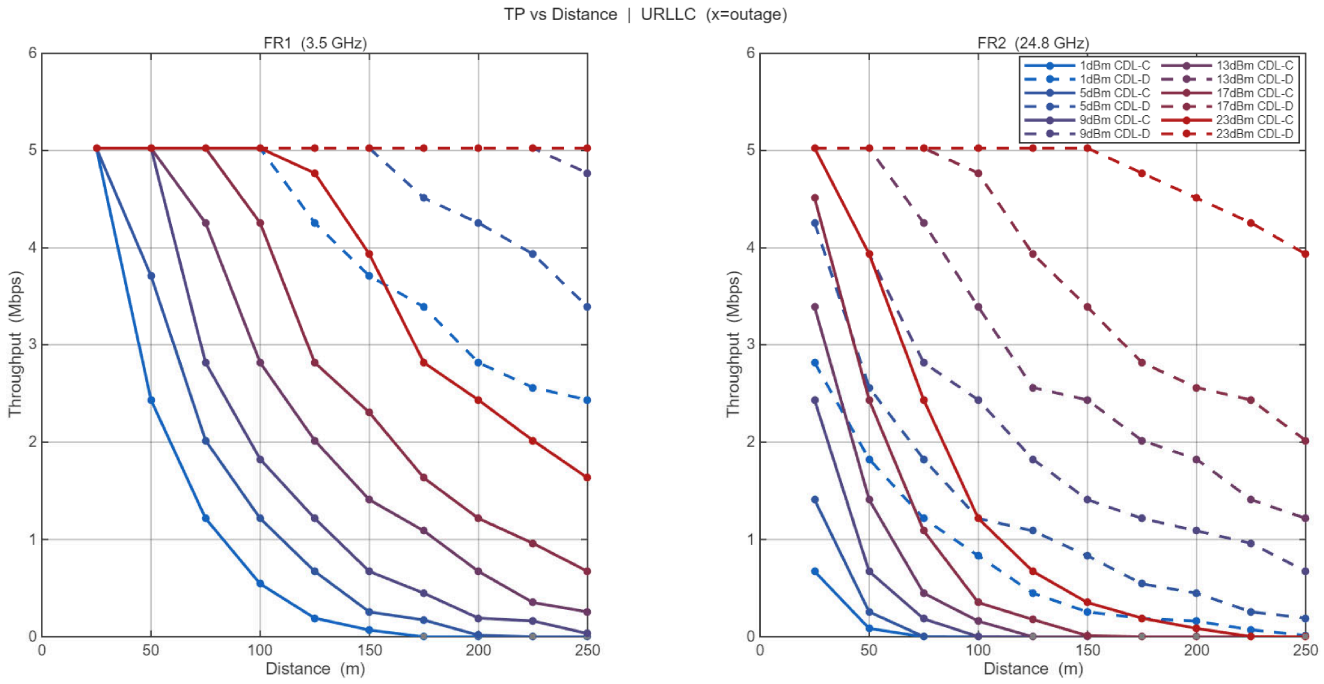


Figure 9. Throughput versus distance for URLLC without FS interference. The LoS link holds its peak of ≈ 5 Mbps across the entire distance range.

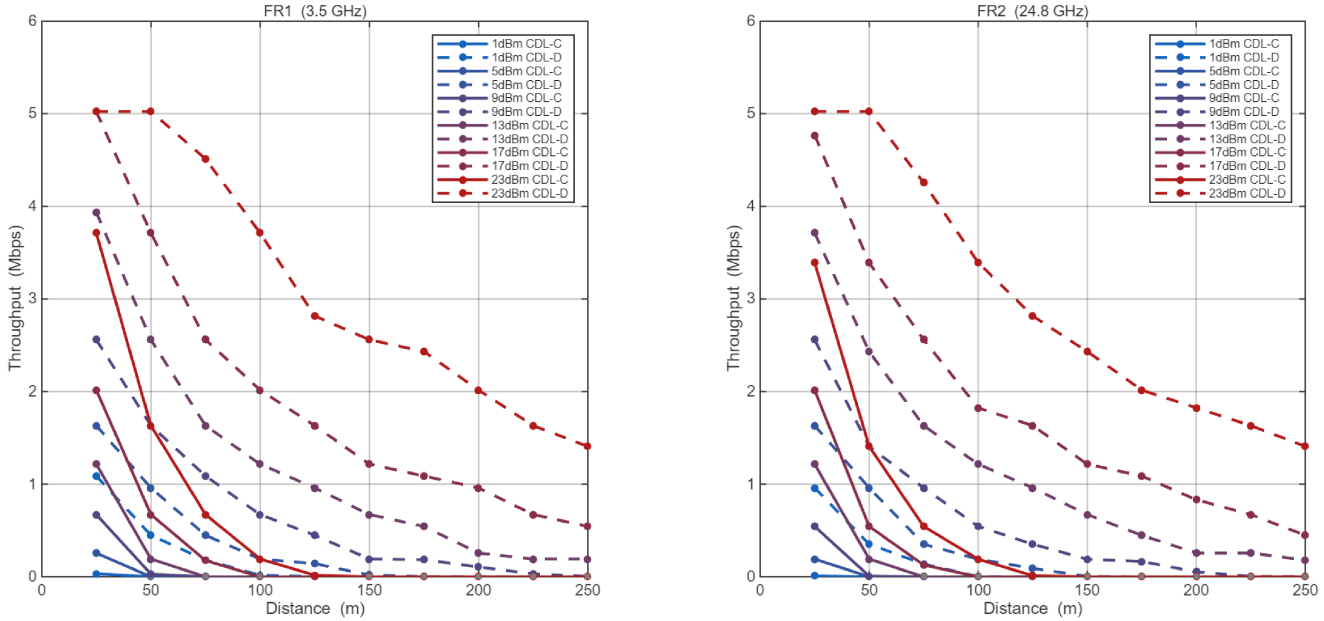


Figure 10. Throughput versus distance for URLLC with FS interference. The short-range LoS peak (≈5 Mbps) was largely retained, but the distance over which it was sustained decreased sharply.

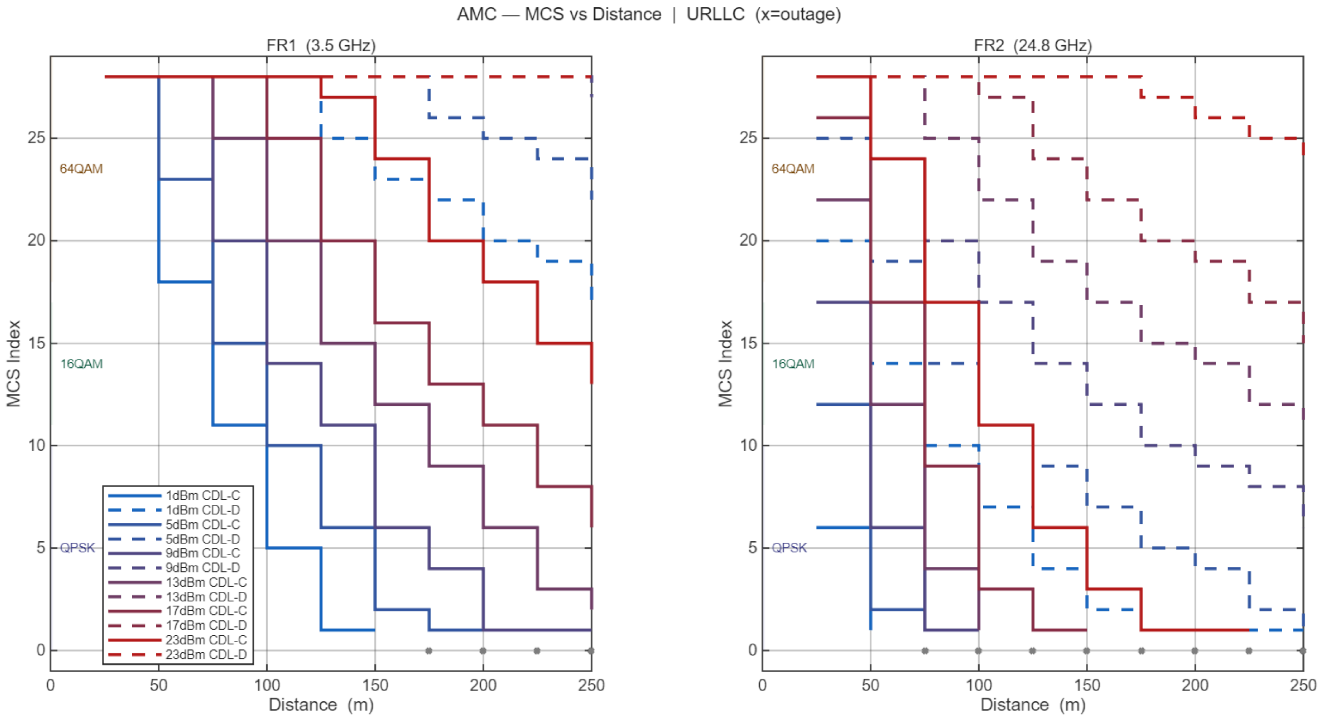


Figure 11. URLLC AMC-selected MCS index versus distance, without FS interference. The mini-slot allocation and the 10^{-5} target keep the selected order conservative (• link is in outage).

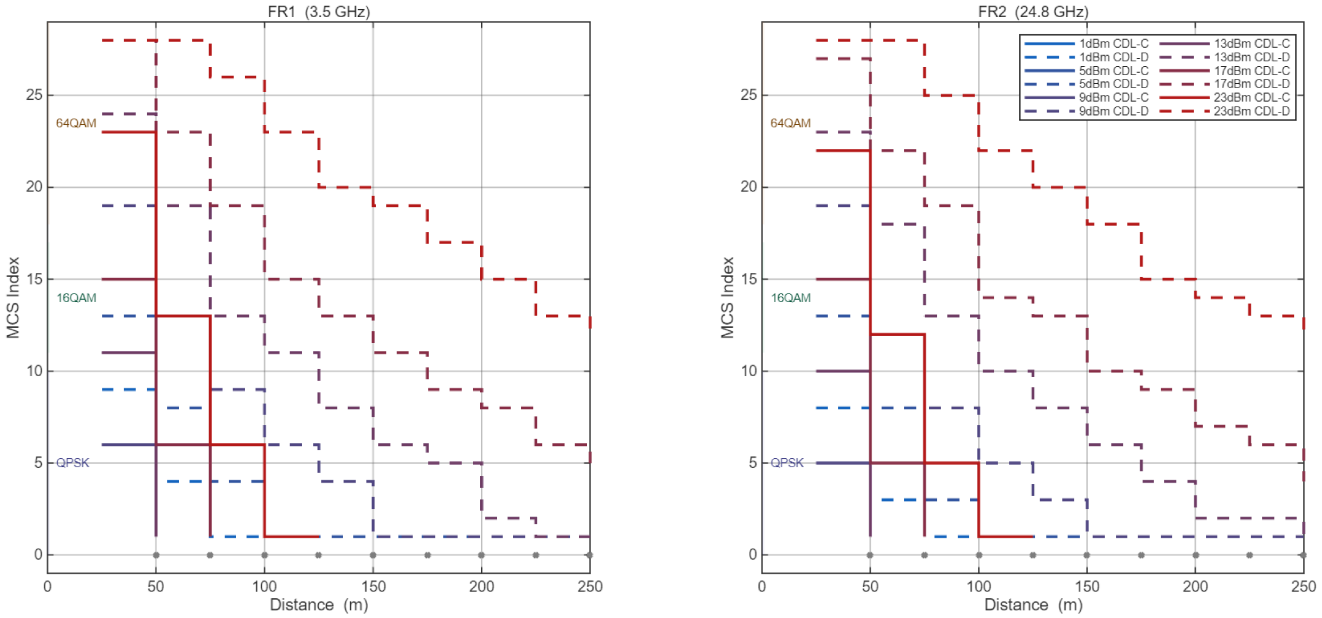


Figure 12. URLLC AMC-selected MCS index versus distance, with FS interference. The sustainable order is pushed down at a shorter range and the decodable distance is curtailed (• link is in outage).

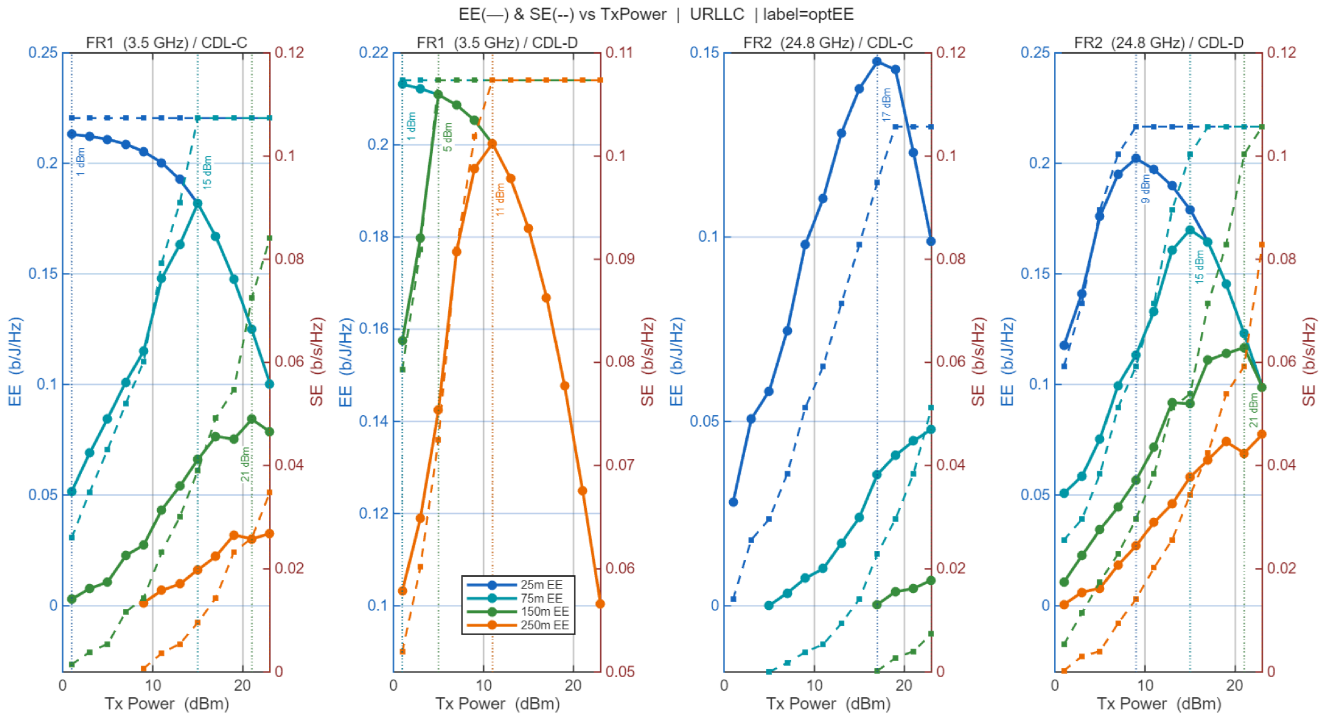


Figure 13. URLLC energy (solid, left axis) and spectral efficiency (dashed, right axis) versus transmit power, without FS interference. The short-range LoS optimum observed at the lowest power level tested. The vertical dotted line in each panel marks the EE-optimal transmit power, annotated in dBm.

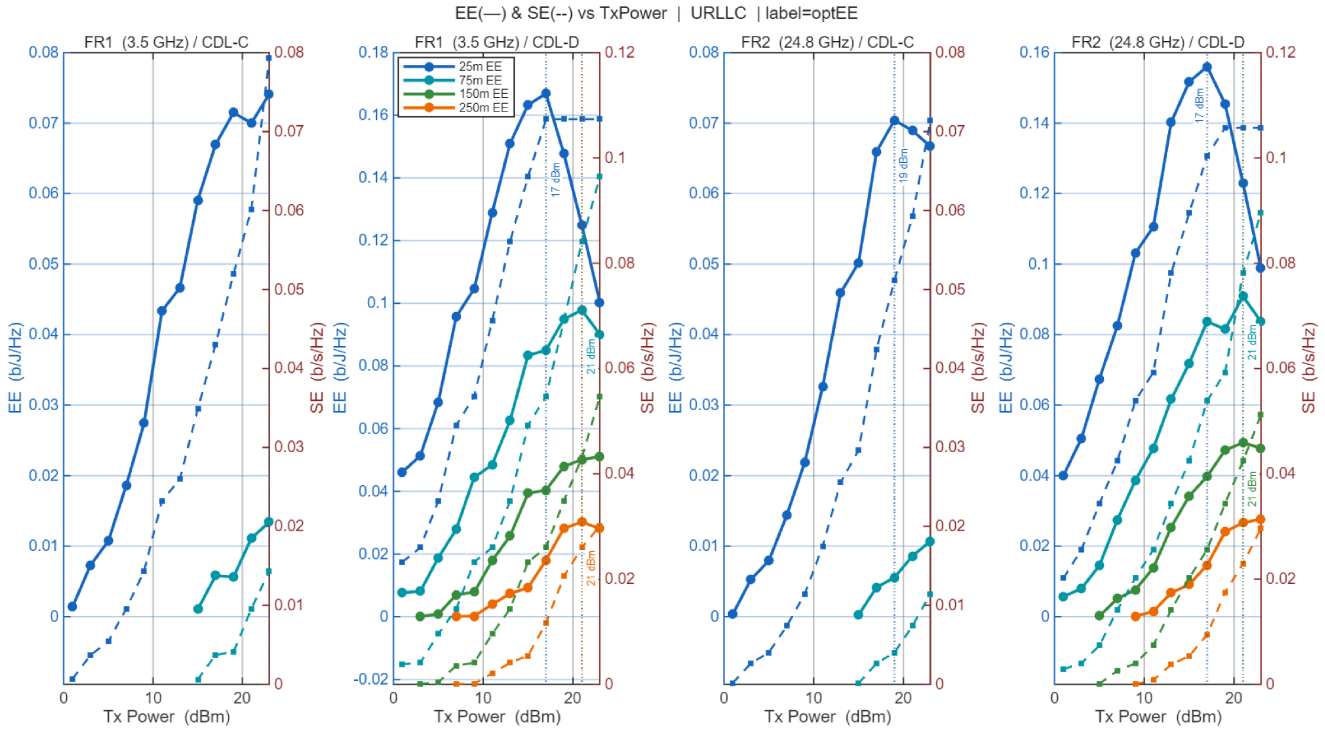


Figure 14. URLLC energy (solid, left axis) and spectral efficiency (dashed, right axis), with FS interference. The energy-optimal power migrated to ≈17 dBm. The vertical dotted line in each panel marks the EE-optimal transmit power, annotated in dBm.

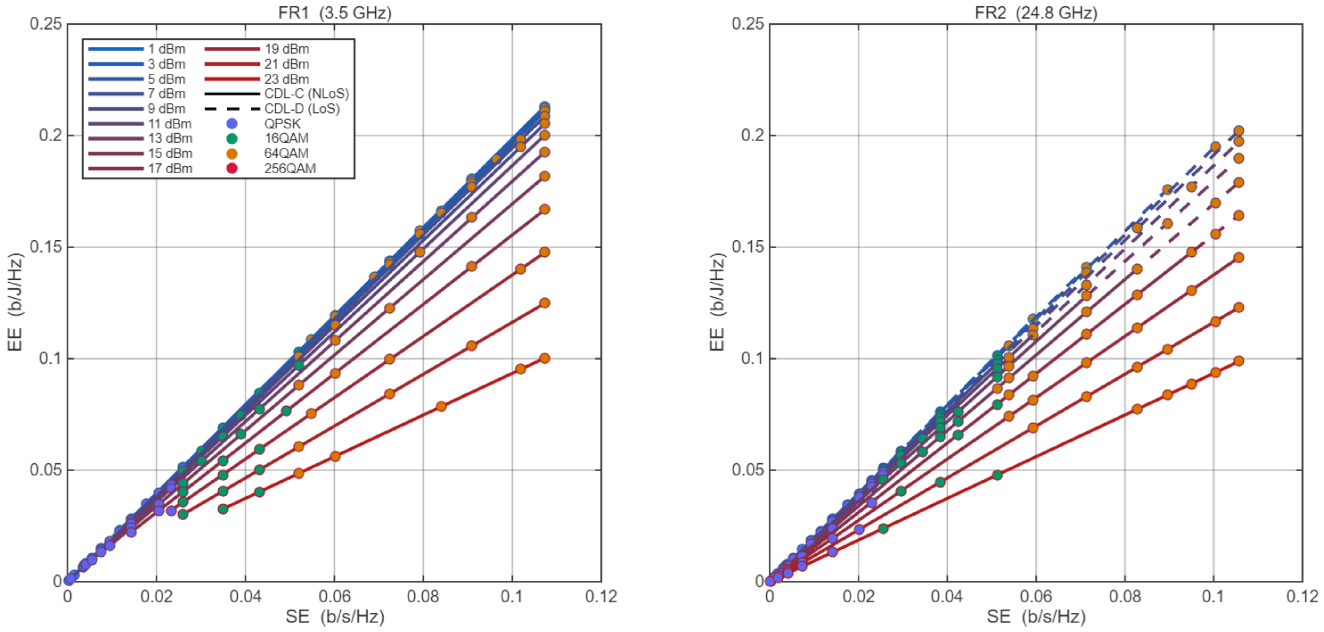


Figure 15. URLLC EE–SE frontier without FS interference. The attainable region is small in absolute terms because of the two-symbol footprint.

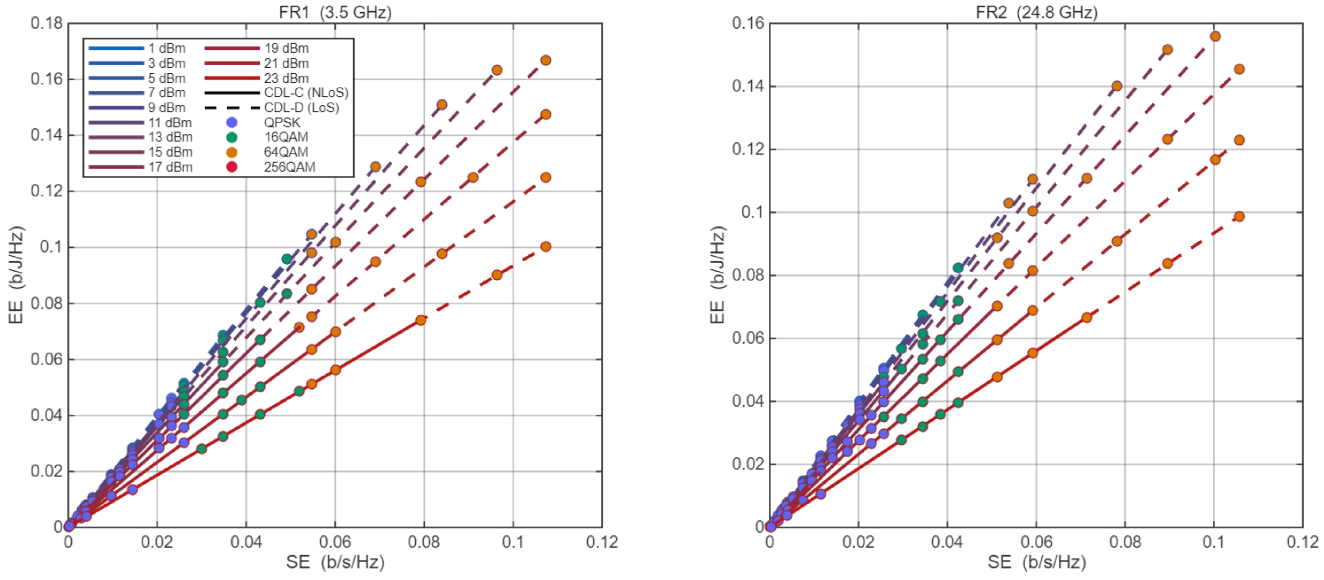


Figure 16. URLLC EE–SE frontier with FS interference. The already small attainable region contracts further.

The reliability and latency design behind these numbers is the one set in Section 3.5: a short mini-slot and a single transmission keep the user plane latency bounded, at the cost of the coding gain that HARQ would provide, in accordance with the numerology versus latency analysis in [27] and the cross-layer view in [30]. The AMC trace (Figures 11 and 12) reflects this conservative operating point, with the selected order being low and curtailed in distance when interference is present.

4.3. mMTC Slice

Figures 17–24 report the mMTC results: throughput (Figures 17 and 18), the AMC-selected modulation order (Figures 19 and 20), the energy and spectral efficiency (Figures 21 and 22), and the EE–SE frontier (Figures 23 and 24). The mMTC slice can be positioned between the other two slices. It uses the low spectral efficiency 64-QAM table and the smallest resource

block share, but it maintains the full 14-symbol slot; therefore, its share of the carrier SE is larger than that of URLLC, and its relaxed 10^{-1} reliability target makes it the most robust of the three. Throughput peaks near 18.4 Mbps, and without interference, the LoS link holds that value flat across the entire distance range in both bands (Figure 17).

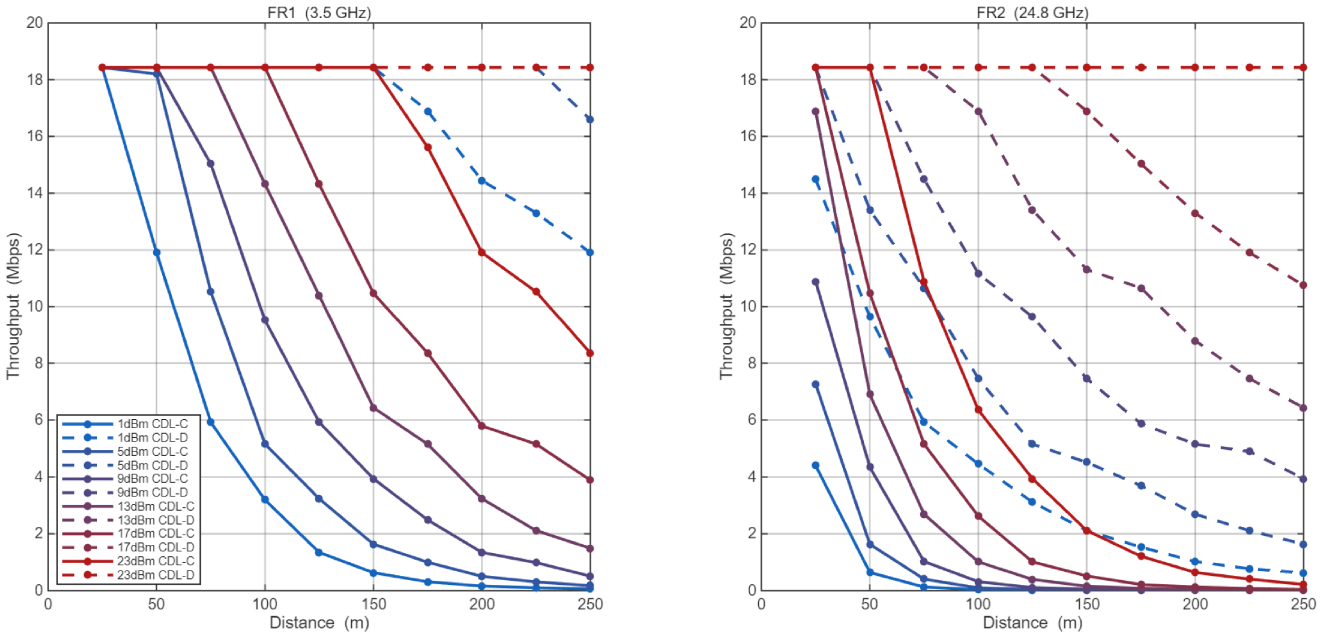


Figure 17. mMTC throughput versus distance, without FS interference. The LoS link held its peak of ≈ 18.4 Mbps flat across the full distance range.

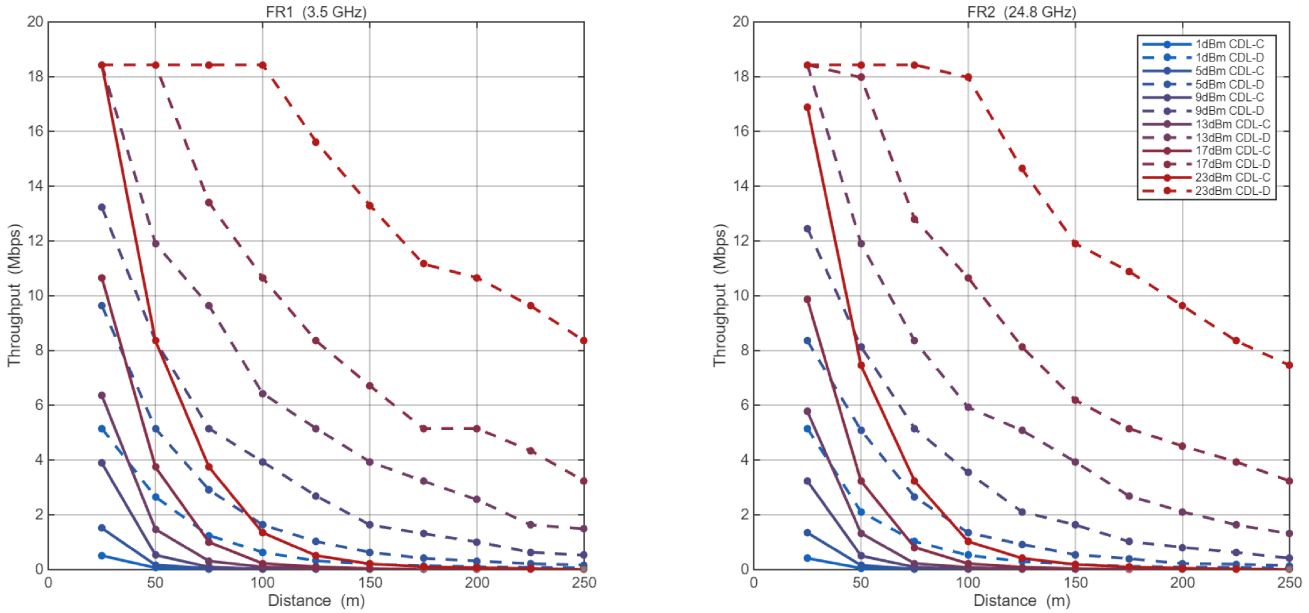


Figure 18. mMTC throughput versus distance, with FS interference. The ≈ 18.4 Mbps LoS peak is retained at a short range, but the flat plateau shortens, and the reach contracts.

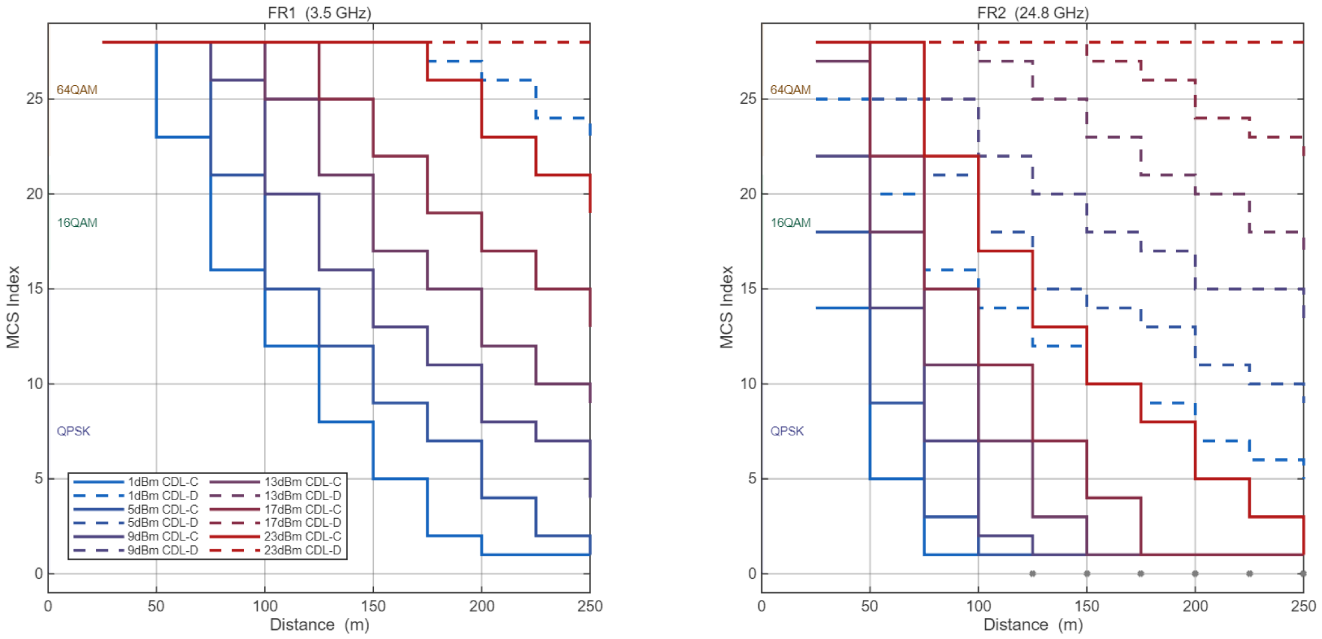


Figure 19. mMTC AMC-selected MCS index versus distance, without FS interference. The relaxed 10^{-1} target allows the loop to maintain a usable order over a wide range. Gray dots at MCS index 0 denote operating points in outage, where no MCS index met the target BLER.

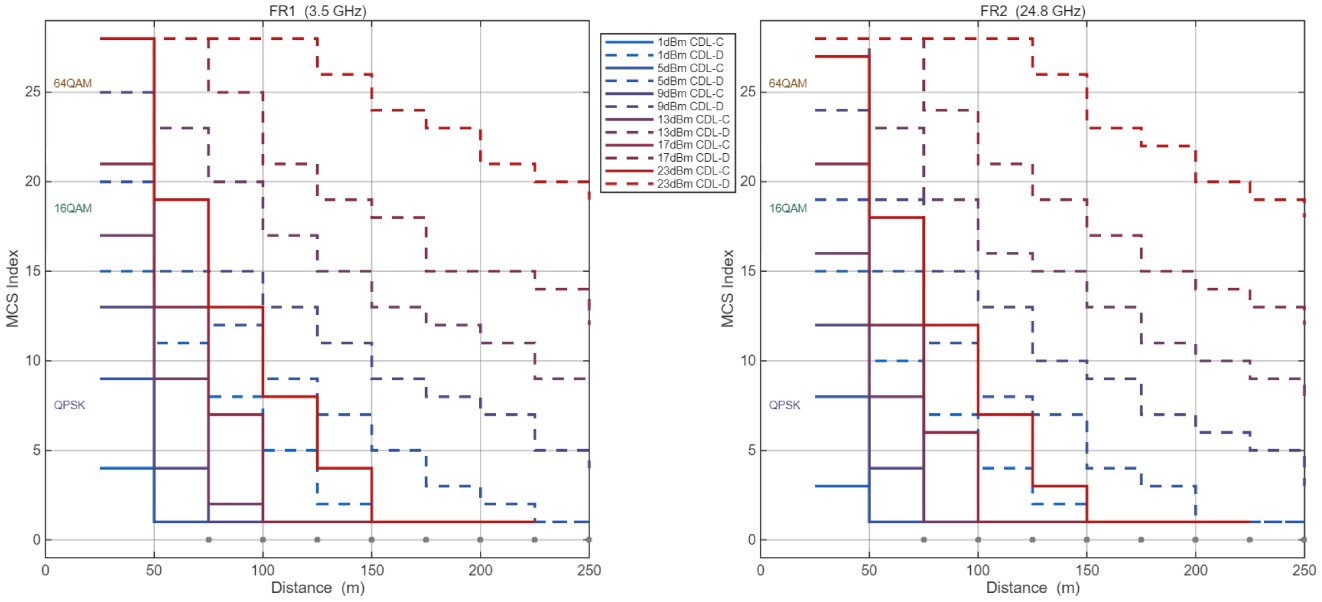


Figure 20. mMTC AMC-selected MCS index versus distance, with FS interference. The sustainable order and reach are reduced, although less severely than for eMBB. Gray dots at MCS index 0 denote operating points in outage, where no MCS index met the target BLER.

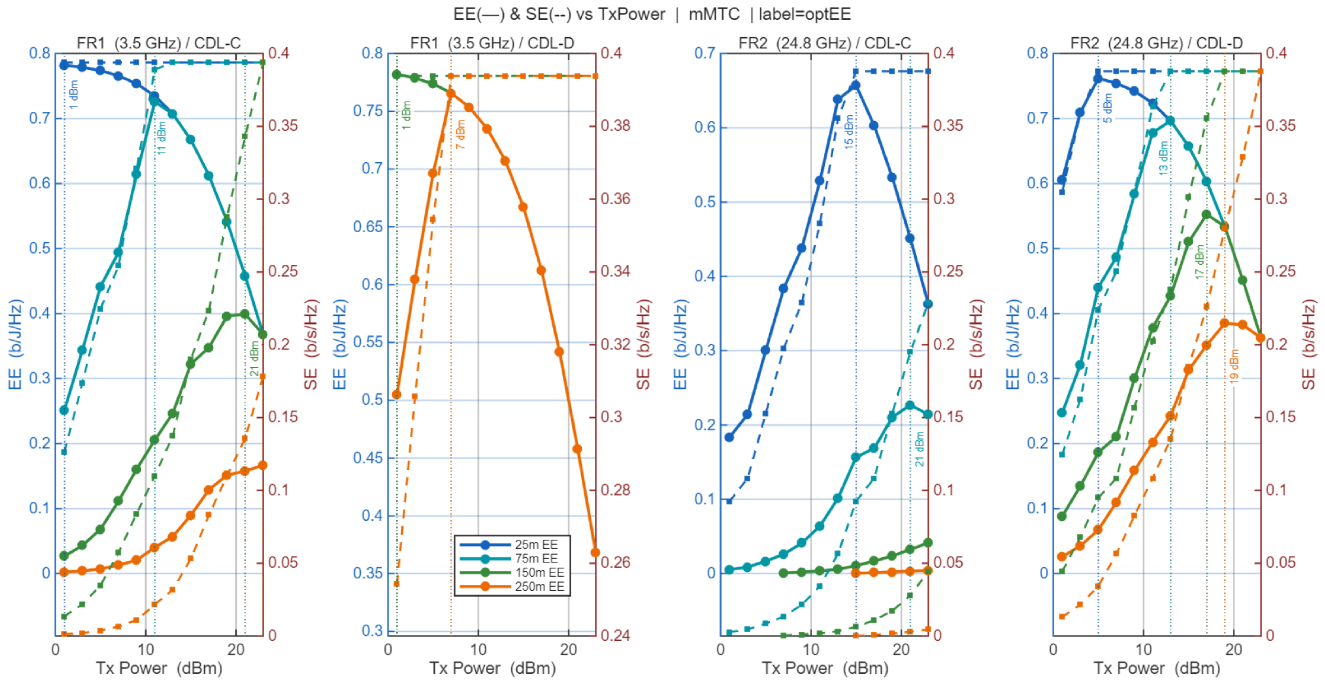


Figure 21. mMTC energy (solid, left axis) and spectral efficiency (dashed, right axis), without FS interference. The vertical dotted line in each panel marks the EE-optimal transmit power, annotated in dBm.

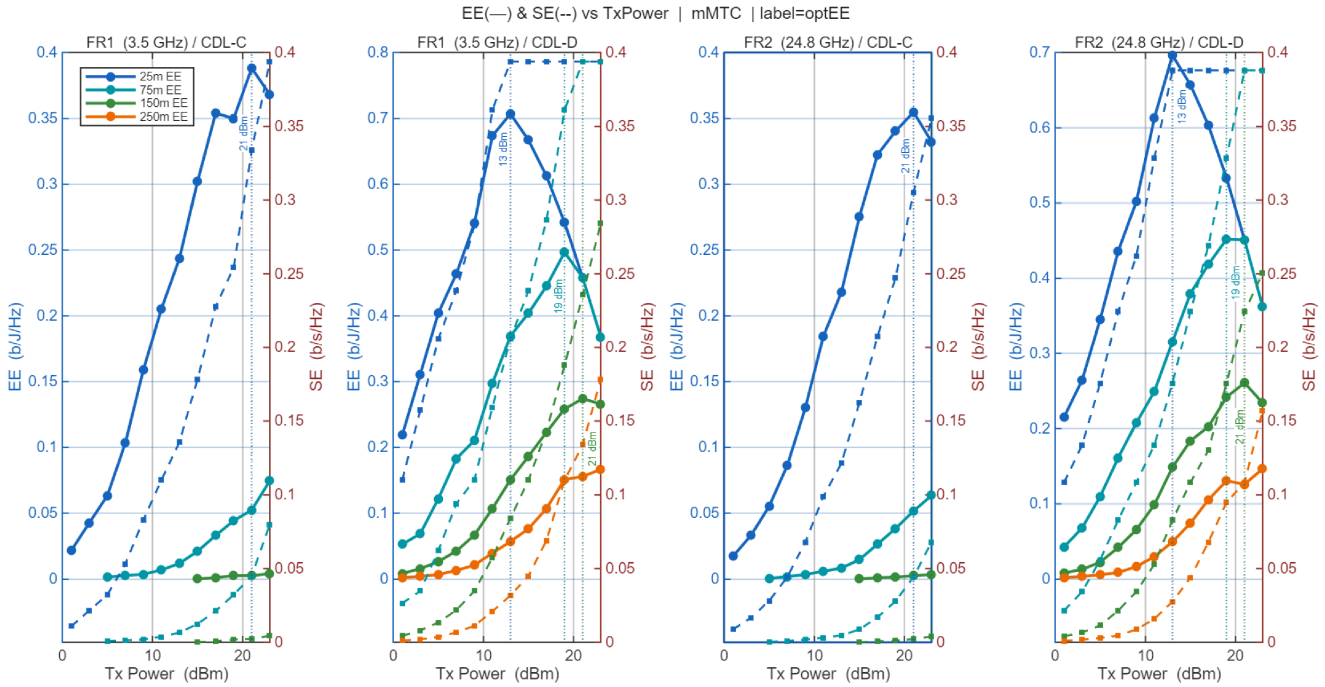


Figure 22. mMTC energy (solid, left axis) and spectral efficiency (dashed, right axis), with FS interference. The peak EE is the most resilient of the three slices to interference, although the EE-optimal power still shifts upward. The vertical dotted line in each panel marks the EE-optimal transmit power, annotated in dBm.

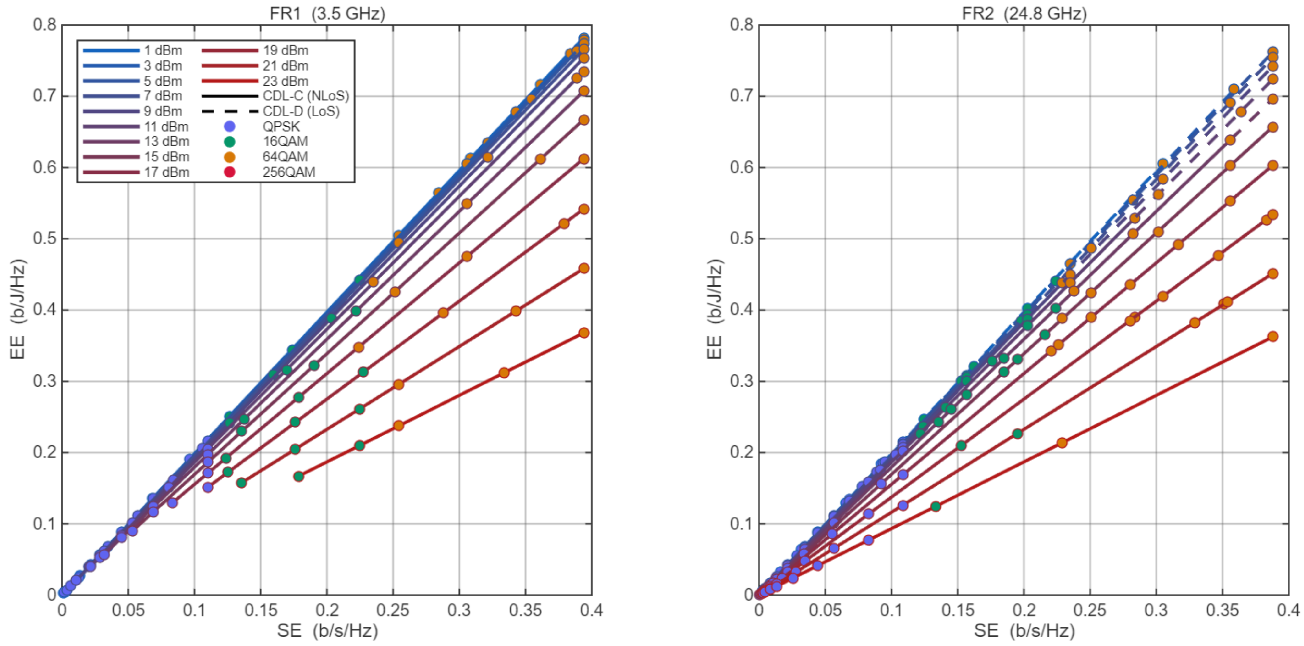


Figure 23. mMTC EE–SE frontier without FS interference. The attainable region is intermediate in scale between eMBB and URLLC.

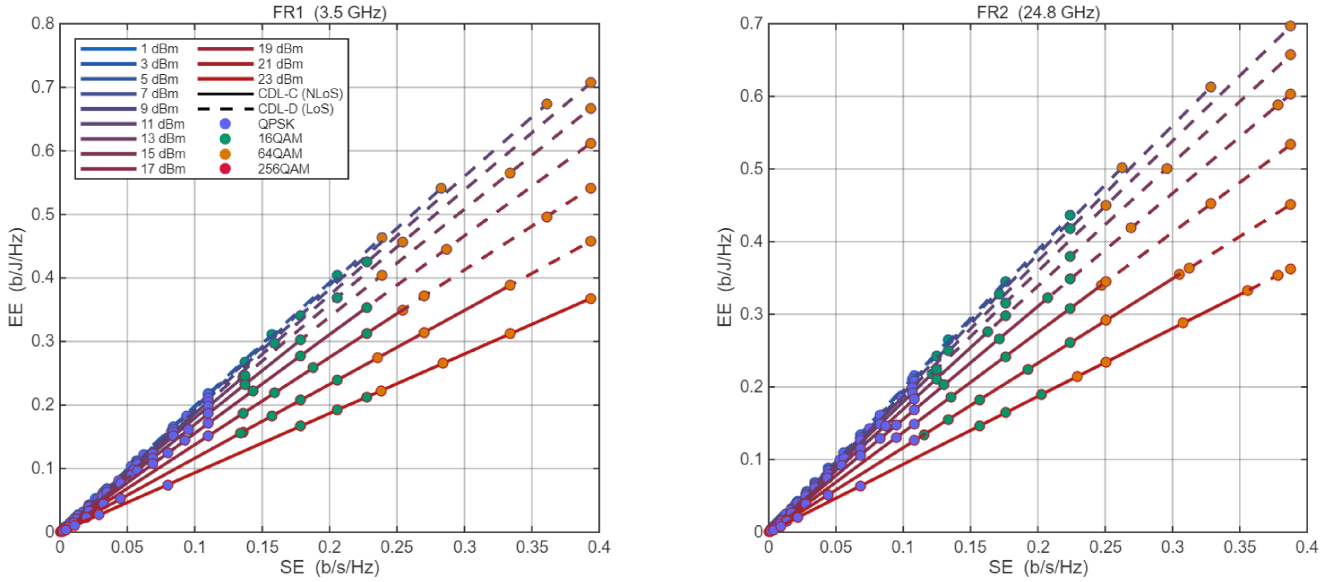


Figure 24. mMTC EE–SE frontier with FS interference. The region contracts modestly compared to the interference-free case.

Under interference (Figure 18), the short-range LoS peak is again retained at approximately 18.4 Mbps, but the flat plateau shortens, and the throughput decays with distance, whereas the NLoS links collapse. The energy plot (Figures 21 and 22) mirrors the URLLC's behavior at a larger scale: without interference, the short-range optimum sits at very low power, at approximately 1 dBm under LoS, because the small allocation saturates quickly; with interference, the optimum shifts to roughly 13 dBm. The peak EE is the most resilient of the three slices, falling only from approximately 0.78 to 0.71 b/J/Hz under LoS, although the NLoS peak still drops by approximately half. This resilience is a direct consequence of the relaxed reliability target: a slice that is content with a 10^{-1} error does not depend on the high received SINR, which is most eroded by interference. The AMC record shows the same pattern in milder form: the relaxed 10^{-1} target sustains a usable order over a wide range without interference

(Figure 19), and both the sustainable order and the reach are reduced under interference, though less severely than for eMBB (Figure 20). The attainable EE–SE region sits between those of eMBB and URLLC and contracts only modestly (Figures 23 and 24).

4.4. Cross-Slice Comparison and Discussion

Read side by side, the three slices occupy clearly separated regions of the EE–SE plane, and the separation follows directly from the per-slice PUSCH design rather than from the channel conditions. The ordering of peak energy efficiency, eMBB well above mMTC and mMTC above URLLC, tracks each slice's effective time–frequency footprint: eMBB spends a full slot over the largest band, mMTC a full slot over a small band, and URLLC only two symbols over a moderate band. Thus, when the per-slice figures are normalized to the common carrier, the broadband slice dominates the energy and spectral metrics, whereas the latency slice contributes the least to the overall performance. This is the reliability-diversity idea of [1] made visible at the link level: the same physical layer, configured in three ways, yields three distinct trade-off curves. Table 11 summarizes the peak figures and their interference impact.

Table 11. Cross-slice peak performance and interference impact at the best operating point of each configuration, across both bands (FR1, FR2) and both propagation profiles (CDL-D LoS, CDL-C NLoS).

Slice	Band	Profile	Cond.	Peak EE (b/J/Hz)	Peak SE (b/s/Hz)	Peak TP (Mbps)	EE-Opt Ptx (dBm)
eMBB	FR1	CDL-D (LoS)	no FS	7.66	3.94	184.4	7
			with FS	3.63	2.98	139.3	15
		CDL-C (NLoS)	no FS	6.92	3.94	184.4	13
			with FS	1.55	1.40	65.6	21
	FR2	CDL-D (LoS)	no FS	5.26	3.88	184.4	15
			with FS	3.43	2.85	135.3	15
		CDL-C (NLoS)	no FS	3.15	2.50	118.8	17
			with FS	1.40	1.42	67.6	21
URLLC	FR1	CDL-D (LoS)	no FS	0.213	0.11	5.0	1
			with FS	0.167	0.11	5.0	17
		CDL-C (NLoS)	no FS	0.213	0.11	5.0	1
			with FS	0.074	0.08	3.7	23
	FR2	CDL-D (LoS)	no FS	0.202	0.11	5.0	9
			with FS	0.156	0.11	5.0	17
		CDL-C (NLoS)	no FS	0.148	0.11	5.0	17
			with FS	0.070	0.07	3.4	19
mMTC	FR1	CDL-D (LoS)	no FS	0.782	0.39	18.4	1
			with FS	0.707	0.39	18.4	13
		CDL-C (NLoS)	no FS	0.782	0.39	18.4	1
			with FS	0.388	0.39	18.4	21
	FR2	CDL-D (LoS)	no FS	0.762	0.39	18.4	5
			with FS	0.696	0.39	18.4	13
		CDL-C (NLoS)	no FS	0.657	0.39	18.4	15
			with FS	0.355	0.36	16.9	21

Interference acts on the three slices through a common mechanism but has markedly different effects. In every case, the raised noise-plus-interference floor lowers the decodable γ_{SINR}, which both caps the attainable modulation order and increases the energy-optimal transmit power. The importance of this depends on how hard the slice leans on a high received SINR value. eMBB, which earns its rate from 256-QAM, loses the most: approximately a quarter of its peak throughput, approximately half of its peak EE under LoS and far more under NLoS, and the entire high-SE tip of its frontier. This exposure is designed rather than incidental: the slices sit on different MCS tables precisely because their reliability definitions differ [1], and the eMBB table is the only one that reaches 256-QAM, whose successful decoding demands the highest received SINR; when the interferer raises the floor, that high-order region is the first to become unsustainable, while the deliberately

conservative URLLC and mMTC tables lean far less on the high-SINR tail. eMBB is also the only slice that uses practical DM-RS-based channel estimation; therefore, its estimation quality degrades with the same floor. The mMTC, with a 10^{-1} target, was the most resilient, surrendering less than one-tenth of its peak EE under LoS. URLLC sits in between on energy but shows the sharpest qualitative change, with its EE-optimal power migrating by approximately 16 dB once the link becomes interference-limited. A single conclusion ties these together: the energy-optimal operating point is not a fixed property of the band, distance, and slice but a function of the interference that the link actually sees. Therefore, a slice-aware uplink power controller must also be interference-aware. This is the operating regime that noise-limited link studies and SE-only optimizations cautioned against in [8], and that such studies do not expose.

The delay profile and band moved in the same direction for all three slices. The LoS profile roughly doubles the achievable EE relative to the NLoS and allows the AMC loop to hold a higher order at the same power and distance, which quantifies the LoS regime that much of the slicing literature sets aside in favor of NLoS Rayleigh assumptions [11]. FR2 incurs a consistent penalty compared to FR1 because of its higher path loss at 24.8 GHz. Because the subcarrier spacing is held at 60 kHz in both bands, this gap isolates the path-loss effect from any change in numerology, and a band-matched numerology with wider FR2 spacing is expected to recover part of it [19,25]. The band-matched check of Section 4.6 quantifies this expectation: the 120 kHz numerology reproduces the 60 kHz energy optima to within a few percent; therefore, the FR1–FR2 gap reported here is propagation-dominated rather than a numerology artifact. Under interference, the LoS advantage narrows because the FS contribution raises the floor regardless of the desired signal geometry; however, the order of the slices, bands, and propagation conditions is maintained. Taken together, the results show that the shape of the EE–SE frontier, location of the energy optimum, and coverage envelope respond to the four factors isolated in this study: the per-slice PUSCH configuration, propagation delay profile, carrier band, and presence of fixed-service interference.

4.5. Statistical Reliability of the Reported Optima

The single-realization presentation of Sections 4.1–4.4 follows from the common random numbers design of Section 3.1: every operating point and every with/without interference pair is evaluated on the identical channel and noise realization, so the curves are strictly comparable, and any difference between two of them is attributable to the controlled factors alone. To quantify the extent to which those curves would move under a different draw, the headline condition (eMBB, FR1, CDL-D, no FS interference) was repeated over ten independent channel seeds across the full 12-level power sweep and all ten distances, 1200 additional link-level runs. Figure 25 reports the mean and $\pm 1\sigma$ envelope, and Table 12 lists the per-distance statistics. The following three observations can be made. First, the location of the energy optimum is statistically representative: the ensemble-optimal transmit power lies within one 2 dB step of the deterministic run's optimum at every distance, and where individual seeds prefer a neighbouring level (the per-seed ranges of Table 12), the surface there is flat, with the ensemble EE at any per-seed optimum staying within 6% of the peak, and within 3% at seven of the ten distances. Second, the peak energy efficiency is tightly clustered, with a per-distance coefficient of variation between 0.0% and 3.8% (Table 12) and a maximum of 4.8% at any distance–power combination of the operating space (median 1.2%). Third, the throughput at full power is seed-invariant up to 175 m, which localizes the residual variability to the modulation-switch boundaries around the optimum. Therefore, the deterministic curves reported in this study are representative of the seed ensemble, and the qualitative conclusions do not rest on a particular realization.

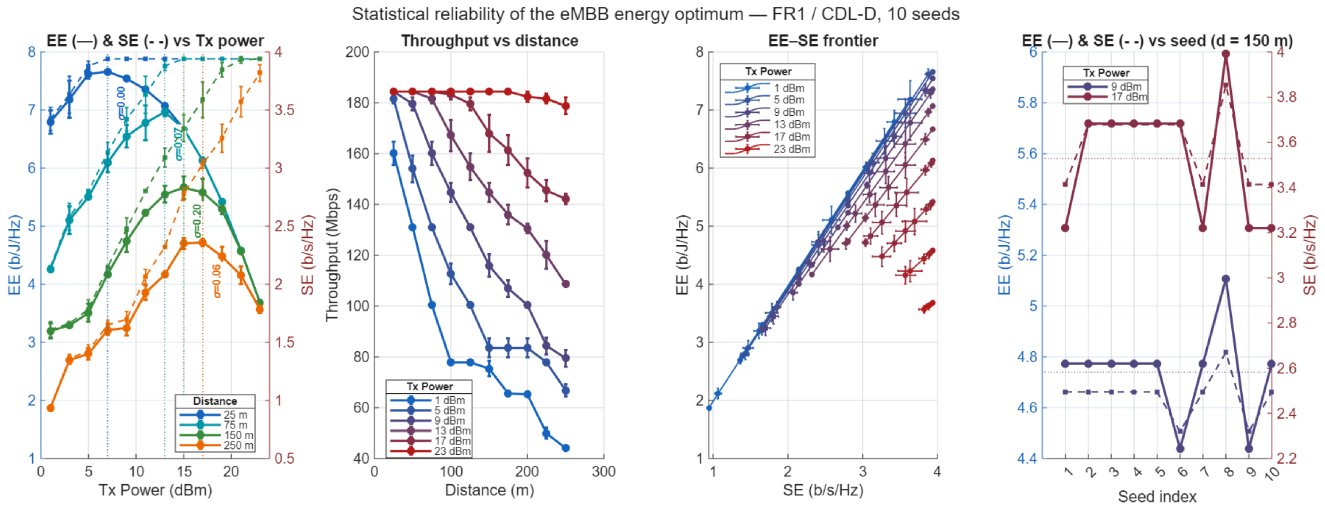


Figure 25. Statistical reliability of the eMBB energy optimum (FR1, CDL-D, no FS interference) over ten independent channel seeds. Left to right: EE (solid, left axis) and SE (dashed, right axis) versus transmit power at four distances, with ±1σ error bars and the standard deviation of the peak EE annotated at each ensemble optimum (dotted); throughput versus distance for six power levels (±1σ); EE–SE frontier with ±1σ bars on both axes; and per-seed EE and SE at d = 150 m for 9 and 17 dBm, with dotted lines marking the ten-seed means.

Table 12. Per-distance seed statistics for the eMBB slice (FR1, CDL-D, no FS interference; ten independent seeds). P* is the ensemble-optimal transmit power with the per-seed range in parentheses; peak EE is the mean ± one standard deviation at P*; throughput is at full power (23 dBm).

d (m)	P* (dBm)	Peak EE (b/J/Hz)	CV (%)	Throughput (Mbps)
25	7 (5–7)	7.66 ± 0.00	0.0	184.4 ± 0.0
50	9 (9–11)	7.34 ± 0.11	1.5	184.4 ± 0.0
75	13 (9–13)	6.97 ± 0.07	1.1	184.4 ± 0.0
100	15 (13–15)	6.50 ± 0.09	1.5	184.4 ± 0.0
125	15 (13–17)	6.05 ± 0.23	3.8	184.4 ± 0.0
150	15 (13–17)	5.66 ± 0.21	3.6	184.4 ± 0.0
175	15 (13–17)	5.37 ± 0.21	3.8	184.4 ± 0.0
200	15 (15–19)	5.15 ± 0.07	1.4	182.4 ± 2.1
225	17 (13–19)	4.83 ± 0.15	3.0	181.6 ± 1.9
250	17 (15–17)	4.72 ± 0.07	1.4	178.7 ± 3.5

4.6. Band-Matched Numerology Check at 120 kHz

The NR scales the subcarrier spacing with the band, and an FR2 deployment normally adopts 120 kHz. To verify that the fixed 60 kHz numerology of Section 3.2 does not distort the FR2 conclusions, the eMBB CDL-D sweep was repeated with the band-matched configuration of TS 38.104 Table 5.3.2-2 [31]: 120 kHz spacing with 32 RBs over the same 50 MHz channel and an occupied bandwidth of 46.08 MHz against 47.52 MHz at 60 kHz. Figure 26 overlays the two numerologies, and Table 13 compares the peak figures. The picture is one of agreement: the interior energy optimum survives at every distance, the energy-optimal transmit power is unchanged at seven of the ten distances on the matched power grid and within one step at the rest, the peak energy efficiency differs by 2.9% on average, and beyond 100 m the throughput agrees to within approximately 5%. The visible gap is confined to short range, where the peak throughput drops from 184.4 to 147.5 Mbps and the peak EE from 5.26 to 4.79 b/J/Hz: the halved symbol duration averages over less of the channel per block-error decision, and the AMC settles one MCS step below the 256-QAM region, the spacing-versus-robustness behavior reported in [19]. The slot duration halves to 0.125 ms, which is a benefit to latency-sensitive traffic that the eMBB metrics do not capture. Therefore, the FR1–FR2 contrast reported in this study is dominated

by propagation rather than numerology, and the fixed-spacing comparison isolates the band effect as intended.

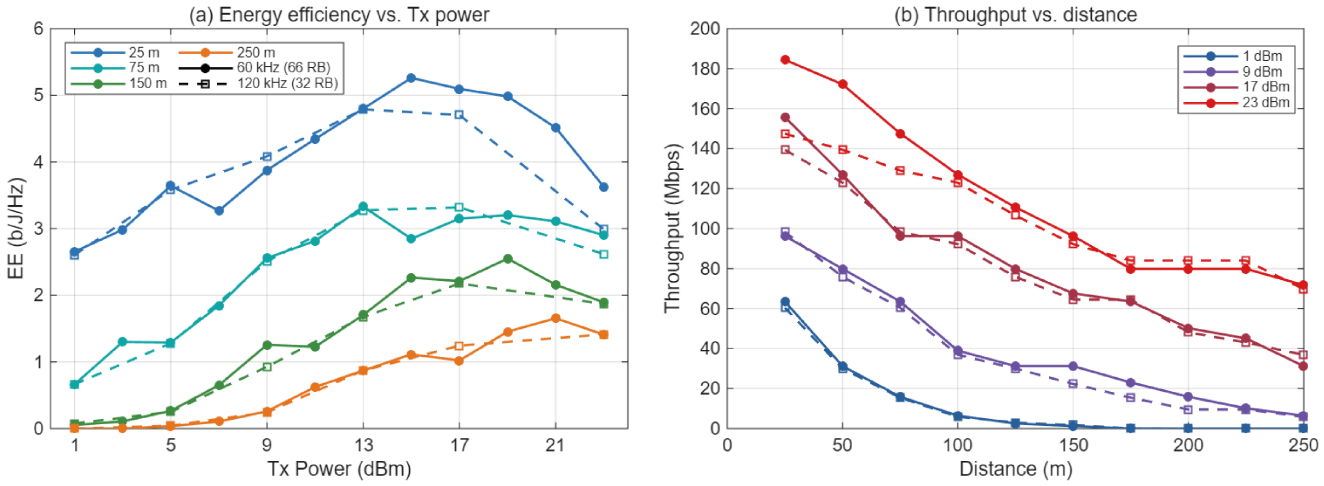


Figure 26. Band-matched numerology check for the eMBB slice in FR2 (CDL-D, no FS interference): 60 kHz with 66 RBs (solid, filled markers) against 120 kHz with 32 RBs (dashed, open markers) over the same 50 MHz channel. (**a**) EE versus transmit power at four distances; (**b**) throughput versus distance at four power levels. The 120 kHz sweep samples six of the twelve power levels.

Table 13. FR2 (CDL-D, eMBB) peak figures at 60 kHz versus 120 kHz, both restricted to the six transmit power levels sampled in the 120 kHz run (1–23 dBm in 4 dB steps) for a like-for-like comparison. Over its full 12-level grid, the 60 kHz sweep places the 25 m optimum at 15 dBm with 5.26 b/J/Hz, the sweep peak quoted in the text; throughput columns are at full power and are identical on both grids.

d (m)	P*60 (dBm)	EE60 (b/J/Hz)	TP60 (Mbps)	P*120 (dBm)	EE120 (b/J/Hz)	TP120 (Mbps)	ΔEE (%)	ΔTP (%)
25	17	5.10	184.4	13	4.79	147.5	−6.0	−20.0
50	17	4.16	172.1	17	4.15	139.4	−0.2	−19.0
75	13	3.33	147.6	17	3.32	129.1	−0.3	−12.5
100	17	3.15	127.0	17	3.11	122.9	−1.2	−3.2
125	17	2.61	110.6	17	2.56	106.6	−2.1	−3.7
150	17	2.21	96.3	17	2.18	92.2	−1.6	−4.2
175	17	2.08	79.9	17	2.18	84.0	+4.7	+5.2
200	17	1.64	79.9	23	1.70	84.0	+3.7	+5.2
225	23	1.57	79.9	23	1.70	84.0	+8.5	+5.2
250	23	1.41	71.7	23	1.41	69.7	+0.2	−2.8

5. Conclusions

This study measured, rather than assumed, the spectral and energy cost of the three IMT-2020 service classes once each is realized as a concrete uplink: a 3GPP-compliant PUSCH link-level simulator provides eMBB, URLLC, and mMTC distinct physical-layer configurations matched to BLER targets of 10^{-3}, 10^{-5}, and 10^{-1}, and traced the joint EE–SE trade-off over 1440 operating points spanning two bands, two propagation delay profiles, twelve transmit-power levels, and ten distances, each evaluated with and without a co-channel fixed-service interferer.

Four results stand out. First, every slice exhibits an interior energy-optimal transmit power set by the balance between the fixed circuit power and amplifier term; therefore, operating at maximum power is energetically wasteful once adaptive coding saturates. Second, the slices occupy clearly separated regions of the EE–SE plane, ordered eMBB above mMTC above URLLC by their effective time–frequency footprints, and reliability diversity made visible at the link level. Third, fixed-service interference reshapes the trade-off through a single mechanism, the raised noise-plus-interference floor, with sharply

slice-dependent consequences: eMBB loses roughly a quarter of its peak throughput and half of its peak EE under LoS (far more under NLoS), mMTC is the most resilient, and the URLLC energy optimum migrates by approximately 16 dB once the link becomes interference limited. Fourth, LoS propagation roughly doubles the achievable EE relative to NLoS, and FR1 sustains approximately twice the EE of FR2 at a fixed numerology, isolating the path-loss effect from any change in subcarrier spacing.

The practical message is that the energy-optimal operating point is not a fixed property of the band, distance, or slice, but a function of the interference that the link actually experiences. Therefore, a power and link-adaptation policy that is slice-aware must also be interference-aware, a regime that noise-limited link studies and spectral-efficiency-only optimizations do not expose. By reporting the per-slice energy and spectral efficiency over the full operating space, with and without fixed-service interference, this study provides concrete interference-conditioned guidance for slice-aware uplink power, modulation, and coding control.

Several extensions of this study remain open. A richer interference model with multiple incumbents and explicit antenna directivity, together with multi-user and non-orthogonal sharing, would allow inter-slice interference to reach the decoded link alongside the fixed-service contribution. A band-matched numerology with wider subcarrier spacing in FR2 would better reflect a real mmWave deployment and is expected to recover part of the reported FR2 penalty. These extensions are the planned next stage of the doctoral programme behind this study: the present engine, exercised over multiple incumbents, directions, and activity patterns, will generate a structured operating database from which a machine-learning stage learns to select the per-slice PUSCH parameters and transmit power that meet each slice's target under the observed channel and interference conditions. A closed-loop controller that is both slice-aware and interference-aware, tracking the interference-conditioned energy optimum rather than a static target, is a natural next step toward tightening the reliability–latency–energy balance of the system.

Author Contributions: Conceptualization, Y.S. and L.A.; methodology, Y.S.; software, Y.S.; validation, Y.S. and L.A.; formal analysis, Y.S.; investigation, Y.S.; resources, Y.S.; data curation, Y.S.; writing—original draft preparation, Y.S.; writing—review and editing, Y.S. and L.A.; visualization, Y.S.; supervision, L.A.; project administration, Y.S. 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.

Informed Consent Statement: Not applicable.

Data Availability Statement: The simulation scripts and datasets generated in this study are not yet publicly available because they are part of the corresponding author's ongoing PhD research. They will be made openly available upon completion of the doctoral project and are available from the corresponding author upon reasonable request in the meantime.

Acknowledgments: The authors acknowledge the use of the MATLAB NR 5G Toolbox in the development and evaluation of the simulation framework. The authors acknowledge the use of GenAI (Paperpal, Pro version) during the editing and revision phases of this manuscript to improve readability and language quality.

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	Fifth-generation mobile network
AMC	Adaptive Modulation and Coding
BLER	Block Error Rate
CDL	Clustered Delay Line
DM-RS	Demodulation Reference Signal
EE	Energy Efficiency
eMBB	Enhanced Mobile Broadband
FR1	Frequency Range 1 (sub-6 GHz)
FR2	Frequency Range 2 (millimeter-wave)
FS	Fixed Service
gNB	Next-Generation Node B (5G base station)
HARQ	Hybrid Automatic Repeat reQuest
IMT	International Mobile Telecommunications
KPI	Key Performance Indicator
LDPC	Low-Density Parity-Check
LoS	Line of Sight
MCS	Modulation and Coding Scheme
mMTC	Massive Machine-Type Communication
NLoS	Non-Line of Sight
NR	New Radio
PA	Power Amplifier
PT-RS	Phase-Tracking Reference Signal
PUSCH	Physical Uplink Shared Channel
RB	Resource Block
SE	Spectral Efficiency
SINR	Signal-to-Interference-plus-Noise Ratio
SNR	Signal-to-Noise Ratio
UE	User Equipment
UL-SCH	Uplink Shared Channel
UMa	Urban Macro
URLLC	Ultra-Reliable Low-Latency Communication

References

1. 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. [CrossRef]
2. Chintalapati, R.V.; Tripathi, N.; Reed, J. Following the Journey of Network Slicing in 5G and 5G Advanced: Essential Concepts, Key Operations, Lifecycle Management, and Future Enhancements. *IEEE Access* **2025**, *13*, 198945–198972. [CrossRef]
3. De Javel, B. Network slicing modeling and RAN physical-layer implementation in 5G. Ph.D. Thesis, Polytechnic Institute of Paris, Paris, France, 2022.
4. Zhang, S. An overview of network slicing for 5G. *IEEE Wirel. Commun.* **2019**, *26*, 111–117. [CrossRef]
5. Metre, P.B.; Kalnoor, G.; G., M.; Gowrishankar, S. A comparative study of network slicing techniques for effective utilization of channel for 5G and beyond 5G networks. *IEEE Access* **2025**, *13*, 31870–31889. [CrossRef]
6. Rafique, W.; Barai, J.R.; Fapojuwo, A.O.; Krishnamurthy, D. A survey on beyond 5G network slicing for smart cities applications. *IEEE Commun. Surv. Tutor.* **2025**, *27*, 595–646. [CrossRef]
7. Donatti, A.W.; Machado, M.C.; Lopez Martinez, M.A.; Antunes, S.R.S.; Figueiredo Souza, E.C.; Corrêa, S.L.; Ferreto, T.C.; Suruagy, J.A.; Martins, J.S.B.; Carvalho, T.C. Energy efficiency in network slicing: Survey and taxonomy. *IEEE Access* **2025**, *13*, 134570–134600. [CrossRef]
8. Liwen, Z.; Qamar, F.; Liaqat, M.; Hindia, M.N.; Zainol Ariffin, K.A. Toward efficient 6G IoT networks: A perspective on resource optimization strategies, challenges and future directions. *IEEE Access* **2024**, *12*, 76600–76640. [CrossRef]

9. Matthiesen, B.; Aydin, O.; Jorswieck, E.A. Throughput and Energy-Efficient Network Slicing. In *WSA 2018; 22nd International ITG Workshop on Smart Antennas*; VDE: Bochum, Germany, 2018; pp. 1–6.
10. Engin, E.; Hökelek, I.; Görçin, A.; Çırpan, H.A. A pre-emptive scheduling mechanism for service assurance of network slicing in next-generation cellular networks. *IEEE Access* **2025**, *13*, 23298–23315. [CrossRef]
11. Novanana, S.; Kliks, A.; Arifin, A.S.; Wibisono, G. Performance of 5G slicing with access technologies and diversity: A review and challenges. *IEEE Access* **2024**, *12*, 170782–170810. [CrossRef]
12. Xu, C.; Hu, Y.; Liang, C.; Ma, J.; Ping, L. Massive MIMO, non-orthogonal multiple access and interleave division multiple access. *IEEE Access* **2017**, *5*, 14728–14748. [CrossRef]
13. Esmaeily, A.; Mendis, H.V.K.; Mahmoodi, T.; Kralevska, K. Beyond 5G resource slicing with mixed-numerologies for mission-critical URLLC and eMBB coexistence. *IEEE Open J. Commun. Soc.* **2023**, *4*, 727–747. [CrossRef]
14. Al-Ali, M.; Yaacoub, E. Resource allocation scheme for eMBB and uRLLC coexistence in 6G networks. *Wirel. Netw.* **2023**, *29*, 2519–2538. [CrossRef]
15. Prathyusha, Y.; Sheu, T.-L. Coordinated resource allocations for eMBB and URLLC in 5G communication networks. *IEEE Trans. Veh. Technol.* **2022**, *71*, 8717–8728. [CrossRef]
16. Liu, B.; Zhu, P.; Li, J.; Wang, D.; You, X. Energy-efficient optimization in distributed massive MIMO systems for slicing eMBB and URLLC services. *IEEE Trans. Veh. Technol.* **2023**, *72*, 10473–10487. [CrossRef]
17. Lorincz, J.; Kukuruzović, A.; Blažević, Z. A Comprehensive Overview of Network Slicing for Improving the Energy Efficiency of Fifth-Generation Networks. *Sensors* **2024**, *24*, 3242. [CrossRef]
18. Rahmani Ghourtani, M.; Zhao, J.; Chu, Y.; Ahmadi, H.; Grace, D.; Maunder, R.G.; Burr, A. Link-level evaluation of uplink cell-free MIMO in 5G NR over frequency-selective channels. *IEEE Open J. Commun. Soc.* **2024**, *6*, 5903–5922. [CrossRef]
19. Pateriya, S.; Bandopadhaya, S.; Bairwa, A.K.; Jain, V. Optimizing 5G NR link layer parameters for eMBB and URLLC applications under dynamic channel and transmission configurations. *Sci. Rep.* **2025**, *15*, 34674. [CrossRef]
20. Kabalcı, Y.; Ali, M. Throughput analysis over 5G NR physical uplink shared channels. In Proceedings of the 2020 2nd Global Power, Energy and Communication Conference (GPECOM), GIzmir, Turkey, 20–23 October 2020; pp. 345–349.
21. Guel, D.; Kouraogo, P.J.; Somda, F.H.; Zerbo, B.; Sié, O. 5G NR Uplink Performance Optimization: A Comprehensive Study on PRACH and PUSCH Interference Management. *EAI Endorsed Trans. Mob. Commun. Appl.* **2024**, *9*, 1–24. [CrossRef]
22. Wei, Z.; Yang, L.; Ng, D.W.K.; Yuan, J.; Hanzo, L. On the performance gain of NOMA over OMA in uplink communication systems. *IEEE Trans. Commun.* **2020**, *68*, 536–568. [CrossRef]
23. Yi, J.; Kim, Y.; Ryu, H. Performance of uplink coverage enhancement schemes for 5G NR in 3GPP. In Proceedings of the IEEE 95th Vehicular Technology Conference (VTC2022-Spring), Helsinki, Finland, 19–22 June 2022.
24. Nasarre, I.P.; Levanen, T.; Pajukoski, K.; Lehti, A.; Tiirola, E.; Valkama, M. Enhanced uplink coverage for 5G NR: Frequency-domain spectral shaping with spectral extension. *IEEE Open J. Commun. Soc.* **2021**, *2*, 1188–1204. [CrossRef]
25. Dahlman, E.; Parkvall, S.; Sköld, J. *5G NR: The Next Generation Wireless Access Technology*; Academic Press: Cambridge, MA, USA, 2018.
26. Sun, C.; She, C.; Yang, C. Energy-efficient resource allocation for ultra-reliable and low-latency communications. In Proceedings of the IEEE Global Communications Conference (GLOBECOM), Singapore, 4–8 December 2017.
27. Hao, P.; Han, X.; Xia, S.; Ren, M.; Deng, Y. Performance evaluation of 5G ultra-reliable and low latency communications. In *2020 International Wireless Communications and Mobile Computing (IWCMC)*; IEEE: Piscataway, NJ, USA, 2020; pp. 1047–1052.
28. Santos, R.; Castanheira, D.; Silva, A.; Gameiro, A. Pipelined multi-user IR-HARQ scheme for improved latency performance in URLLC. *IEEE Access* **2024**, *12*, 33475–33490. [CrossRef]
29. Varatharaajan, S.; Grossmann, M.; Del Galdo, G. 5G New Radio physical downlink control channel reliability enhancements for multiple transmission-reception-point communications. *IEEE Access* **2022**, *10*, 97394–97407. [CrossRef]
30. Haque, M.E.; Tariq, F.; Hossain, M.S.; Khandaker, M.R.A.; Imran, M.A.; Wong, K.-K. A Comprehensive Survey of 5G URLLC and Challenges in the 6G Era. *arXiv* **2025**, arXiv:2508.20205.
31. 3rd Generation Partnership Project (3GPP). NR; Base Station (BS) Radio Transmission and Reception. In *Technical Specification TS 38.104*; 3GPP: Sophia Antipolis, France, 2024.
32. Chase, D. Code combining—A maximum-likelihood decoding approach for combining an arbitrary number of noisy packets. *IEEE Trans. Commun.* **1985**, *33*, 385–393. [CrossRef]
33. Gallager, R.G. Low-density parity-check codes. *IRE Trans. Inf. Theory* **1962**, *8*, 21–28. [CrossRef]

34. 3rd Generation Partnership Project (3GPP). Study on Channel Model for Frequencies from 0.5 to 100 GHz. In *Technical Report TR 38.901*; 3GPP: Sophia Antipolis, France, 2020.
35. Goldsmith, A.J.; Chua, S.-G. Adaptive coded modulation for fading channels. *IEEE Trans. Commun.* **1998**, *46*, 595–602. [CrossRef]

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.