Disentangling the Throughput Contributions of MIMO and Carrier Aggregation in 5G Networks

Yufei Feng[1], Phuc Dinh[1], Moinak Ghosal[1], Omar Basit[2], Sizhe Wang[1], Y. Charlie Hu[2], and Dimitrios Koutsonikolas[1]

[1] Northeastern University, USA
{feng.yuf,dinh.p,ghoshal.m,wang.sizh,d.koutsonikolas}@northeastern.edu
[2] Purdue University, USA {obasit,ychu}@purdue.edu

Abstract. Multiple-input multiple-output (MIMO) and carrier aggregation (CA) are two key MAC/PHY layer technologies employed by both user equipment and base stations to boost data throughput in 5G networks; yet, their respective contributions on real-world throughput improvements remain largely unexplored. Although both approaches conceptually rely on parallel transmissions, they differ fundamentally in their implementations: MIMO exploits spatial diversity through multiple data streams within a band, while CA aggregates spectrum across multiple bands. In this work, we present the first comparative study of MIMO and CA throughput gains in operational 5G networks. Using extensive measurements with commercial smartphones over all three major US cellular operators during a cross-country trip (from LA to Boston, 5700km+), we first present the current state of deployment of both technologies in today's 5G networks. We then disentangle their combined effects on throughput, quantifying the relative contribution of each technology to overall performance. Finally, we analyze how throughput scales across higher-order configurations of each technology, considering different MIMO transmission ranks and numbers of aggregated carriers, providing insights for future 5G deployments.

1 Introduction

The deployment of 5G networks has been driven largely by the promise of higher data rates. While 5G promises multi-gigabit rates, the actual improvements over 4G can be inconsistent, depending on spectrum availability, deployment configurations, and device capabilities. Two physical layer techniques are central to these throughput gains: multiple-input multiple-output (MIMO) communication and carrier aggregation (CA). Although both were first introduced in 4G, their role has expanded in 5G, where the availability of new spectrum bands and more advanced device capabilities make them key enablers of high throughput.

MIMO and CA, however, represent fundamentally different throughput enhancing approaches with distinct trade-offs. CA improves throughput by expanding the aggregate communication bandwidth, but requires access to licensed spectrum, which is scarce and tightly regulated, with limited availability, driving up the cost to mobile operators. In contrast, MIMO increases capacity by transmitting multiple independent data streams over the same frequency band, exploiting differences in wireless propagation paths along with advanced signal processing. Its gains, however, depend on sophisticated hardware at both the base station and the user equipment (UE), as well as favorable propagation conditions.

Consequently, operators face spectrum costs when scaling throughput with CA, while vendors face higher design complexity and manufacturing cost when scaling throughput with MIMO. Moreover, while in theory throughput scales linearly with the MIMO rank (i.e., the number of independent data streams) and the aggregate channel bandwidth, the benefits can quickly saturate in practice and may even reverse under unfavorable conditions. Ultimately, these trade-offs impose costs on different stakeholders, raising the question of whether the throughput gains of the two technologies justify their respective economic and technical burdens.

Understanding the individual impact of each of the two technologies on end user throughput, however, is challenging. In operational networks, MIMO and CA are often used jointly and the throughput observed by end users reflects their combined effects, as well as the impact of other factors, such as frequency band, channel quality, modulation and coding scheme (MCS), or overall cell load. Previous works [28–30, 36, 42, 47] have often focused on improving MIMO and CA in isolation, developing models and algorithms that advance the performance of one of the two technologies in terms of physical layer metrics rather than the overall end user throughput in operational networks. In contrast, previous measurement studies [10, 18–20, 34, 40] report only aggregate end user performance results from operational networks, where the gains of MIMO and CA are entangled. To our best knowledge, no previous work has attempted to *disentangle* the contributions of MIMO and CA to the overall end user throughput, to determine which technology contributes more and under what conditions, and to analyze how throughput scales with higher order MIMO or CA configurations in operational networks. Answering these questions is essential for operators and vendors to inform spectrum acquisition decisions, hardware design choices, and future deployment and configuration decisions.

To fill this gap, in this paper, we present the first large-scale comparative study of the individual throughput gains from MIMO and CA in operational 5G networks, using a large dataset collected through a cross-country driving trip (Los Angeles to Boston, 5700+ km). Our dataset spans all three major US operators (AT&T, Verizon, T-Mobile), multiple frequency bands (5G-low, 5G-mid, and 5G-high/mmWave), and both 5G deployment modes (non-standalone/NSA and standalone/SA). Our study makes the following contributions: (1) We characterize the current state of deployment of MIMO and CA across 5G low, mid, and mmWave bands. (2) We disentangle their combined effects and quantify the relative throughput gains attributable to each technology. (3) We analyze throughput scaling with higher-order configurations, including different transmission ranks and numbers of aggregated carriers. (4) To facilitate reproducibility and further research, we make our dataset publicly available [3].

The main findings of our study are summarized as follows:

• **Deployment–standard gap.** Halfway through the 5G life cycle, commercial deployments remain far from 3GPP-specified capabilities. While NR supports up to 8×8 downlink MIMO, 4×4 uplink MIMO, and 16 aggregated carriers, we observe at most 4×4/2×2 MIMO in the downlink/uplink direction and up to 8 CA in 5G-high bands, only half of the standardized limit, and even fewer carriers in 5G-low and 5G-mid bands.

• **Operator-specific priorities.** Throughput scaling patterns reveal clear differences in operator priorities and deployment strategies. AT&T achieves most gains through MIMO enhancements in the mid band, T-Mobile leverages broader mid band CA for consistent throughput growth, and Verizon's mmWave focused strategy yields extreme peak rates but limited coverage.

• **Technology scaling limits.** MIMO performance exhibits diminishing returns beyond two to three layers, with higher ranks often reducing throughput due to practical channel limitations, while CA generally delivers steadier, typically superlinear throughput growth with the number of aggregated carriers across bands but remains disproportional to the aggregate bandwidth. Together, these trends indicate that scaling in commercial networks is ultimately constrained by real-time operating conditions rather than merely using higher-order configurations of these technologies.

2 Background

This section provides a brief background on CA and MIMO and their implementation details specified by 3GPP standards.

2.1 MIMO

MIMO uses multiple transmit and receive antennas to improve both the capacity and reliability of a wireless link. It operates primarily in two modes. Spatial multiplexing increases throughput by transmitting multiple independent data streams simultaneously, each over a distinct spatial channel. Spatial diversity, in contrast, transmits redundant copies of the same signal across antennas to improve reception reliability under fading. The number of spatial streams that can be transmitted at once depends on the rank of the channel, which reflects how many of those spatial paths are effectively independent. High-rank channels arise in rich scattering environments with low correlation, enabling true multiplexing gain, while highly correlated channels have fewer independent paths, limiting the number of spatial streams. As the number of transmitted streams approaches the effective rank of the channel, inter-stream interference grows, leading to diminishing returns from higher-order multiplexing.

The transmitter continuously estimates or receives channel state information (CSI) feedback from the receiver to adapt its transmission strategy. This process forms a closed-loop control system: the receiver measures the downlink channel, quantizes and reports channel state information feedback such as preferred precoding matrices, rank indicators (RI), and channel quality indicators (CQI), and the transmitter updates its precoder accordingly. The precoding matrix linearly combines the data streams across antennas to align them with the strongest and most independent spatial paths of the channel. This alignment reduces interference between layers and helps the receiver separate the signals more effectively. In high-mobility environments, feedback delay causes the transmitter to base its precoding on outdated CSI, reducing the accuracy with which it can match its transmission to the true channel and leading to performance degradation.

3GPP Release 15 established the baseline MIMO framework for 5G NR, defining support for up to 8 spatial data streams in the downlink and up to 4 in the uplink [1], although in commercial networks these numbers are typically lower. Measurement and diagnostic interfaces, including XCAL and Qualcomm logs, report two MIMO-related quantities, *MIMO mode* and *MIMO layers*. The *MIMO mode* specifies the physical antenna configuration, such as 2×2, 4×4, or 8×8, representing the number of available transmit and receive antennas. The number of *MIMO layers* (or transmission rank) indicates the number of independent data streams actually transmitted at a given time, which depends on the instantaneous channel conditions and scheduling decisions.

2.2 Carrier Aggregation

CA enables a user equipment (UE) to transmit and receive data over multiple component carriers (CCs) or cells[3] at the same time. Since each CC provides its own independent channel resource, combining them effectively increases the total bandwidth available to the link, which directly translates into higher potential throughput. In theory, the maximum achievable throughput grows proportionally to the aggregate bandwidth, as a wider channel allows more symbols to be conveyed per unit time under similar channel conditions. However, in practice, the realized gains can be lower due to scheduling inefficiencies, coordination overhead between carriers, differences in channel propagation properties among different bands, and the choice of modulation and coding schemes (MCS) and number of MIMO layers in each carrier.

In CA, one carrier is designated as the Primary Cell (PCell), which serves as the anchor for control signaling and may also carry user data, while the additional Secondary Cells (SCells) provide extra bandwidth for data transmission and can be activated or released as needed.[4] CA can be *intra-band*, where the CCs belong to the same 3GPP-defined frequency band (e.g., n41+n41, with

[3] The two terms are used interchangeably.

[4] In this work, CA refers to the aggregation of different 5G cells. In this context, one of the 5G CCs, over which the 5G RRC signalling messages are transmitted along with user data, is marked as PCell, while other CCs (SCells) are added or removed dynamically for data transmission only. We do not consider CA with respect to 5G NSA Dual Connectivity, where the 5G connection acts as a secondary cell (SCell) by anchoring over a 4G LTE primary Cell (PCell).

both CCs belonging to the mid band n41 at 2.5 GHz) or *inter-band*, where the CCs belong to different bands (e.g., n71+n41, with the PCell belonging to the low band n71 at 600 MHz and the SCell belonging to the mid band n41 at 2.5 GHz). Different CCs can have different bandwidths and, in the case of inter-band CA, different subcarrier spacing (SCS) and different transmission mode (TDD vs. FDD). While inter-band CA allows for a more flexible use of resources by enabling operators to utilize scattered spectrum chunks, it also introduces design complexity for the RF hardware and requires interference-aware scheduling, along with additional signaling overhead.

3GPP Release 15, finalized in 2019, was the first complete 5G NR specification and established the foundation for CA capabilities in NR [1]. 5G NR, by standardization, may support up to 16 CCs in each direction. Each carrier must be at least 5 MHz wide. In the sub-6 GHz (low and mid) frequency bands (FR1), a single CC can be as wide as 100 MHz, while in the mmWave bands (FR2) it can be as wide as 400 MHz. This gives a theoretical maximum of 1.6 GHz total downlink or uplink bandwidth in FR1 and 6.4 GHz in FR2. In practice, commercial deployments are far more limited. Most networks activate only a few carriers in parallel, with total aggregate bandwidth usually in the range of a few tens to a few hundreds of MHz. Later releases (16-18) refined signaling, expanded uplink aggregation options, and added efficiency improvements, but did not alter the fundamental limits introduced in Release 15.

Summary. Both MIMO and CA conceptually rely on parallel transmissions, but they expand capacity along different resource dimensions of the system. MIMO increases spectral efficiency by transmitting multiple data streams within the same bandwidth, while CA increases throughput by combining additional spectrum resources. In practice, MIMO is limited by propagation conditions and hardware, while CA is limited primarily by spectrum availability and coordination overhead.

3 Methodology

Driving route. We drove 5700+ km from Los Angeles, CA to Boston, MA covering all major cities in between (Las Vegas, Salt Lake City, Denver, Omaha, Chicago, Indianapolis, Cleveland, Rochester) over 8 days – November 1–8, 2024. The route is shown in Fig. 1. Our measurement data were collected on interstate highways, in suburban areas, and within city limits. Additional measurements were conducted within several major cities encountered along the routes. The vehicle speed ranged from 5 miles per hour in cities to 90 miles per hour on interstate highways.

5G Carriers, UE, and Servers. We obtained multiple unlimited data plans from the three major US cellular operators – Verizon, T-Mobile, and AT&T – to conduct our experiments. All three operators offer substantial 5G coverage along the driving route (AT&T: 44%, Verizon: 49%, T-Mobile: 94% of the driving route). T-Mobile is the only US carrier currently offering both NSA and SA services, with 5G SA covering 60% of the driving route. The other two operators

Fig. 1: Driving route.

Table 1: Bands and number of measurement samples in our dataset.

		Number of samples		
	Band(frequency/mode)	**AT&T**	**Verizon**	**T-Mobile**
Low	n5(850 MHz/FDD)	192,243	7,947	-
	n71(600 MHz/FDD)	-	-	226,829
Mid	n77(3.7 GHz/TDD)	210,235	531,237	-
	n2(1.9 GHz/FDD)	9,585	1,208	-
	n66(1.7/2.1 GHz/FDD)	7,084	-	209
	n41(2.5 GHz/TDD)	-	-	766,978
	n25(1.9 GHz/FDD)	-	-	206,990
High	n260(39 GHz/TDD)	21,149	1,990	37
	n261(28 GHz/TDD)	-	100,895	-

offer only 5G NSA service. AT&T and Verizon offer service in the low, mid and high bands, while T-Mobile offers service primarily in the low and mid bands. Table 1 summarizes the different bands for each operator and the number of measurement samples collected in each band. In our analysis, we do not consider bands n66 and n260 for T-Mobile, due to the very small number of measurements over these two bands.

We used Samsung Galaxy S24 smartphones as our UEs, featuring the Snapdragon X75 modem. The phone supports downlink/uplink CA of up to 4/2 CCs in the low- and mid-bands and up to 8/4 CCs in the mmWave bands. All our experiments were conducted with 3 UEs, each connected to a different operator, to *concurrently* measure the performance over the three operators. Since commercial off-the-shelf smartphones do not provide access to low-level KPIs (MIMO, CA, MCS, etc.), we used Accuver XCAL Solo [2] devices, which tap into the Qualcomm diagnostic (diag) interface to capture signaling messages exchanged between the UE and the base station and extract these KPIs. For server-side infrastructure, we deployed two AWS EC2 cloud servers in California and Ohio for tests done with all three operators, and AWS EC2 Wavelength edge servers in Los Angeles, Las Vegas, Denver, Chicago, and Boston for tests done only with Verizon in those cities.

Experimental Methodology. We conducted backlogged downlink and uplink throughput tests using nuttcp [37] in a round robin fashion throughout the driving trip. The nuttcp tests lasted for 2 min each. Throughput samples were logged at intervals of 500 ms. Lower layer metrics (CA and MIMO information, MAC throughput per carrier, bandwidth, MCS, RSRP, CQI, etc.) are logged by XCAL at a granularity of 100 ms. Since higher-order MIMO and CA are typically activated only under heavy traffic to conserve energy and spectrum, we only use the data collected during download/upload tests for our downlink/uplink analysis, respectively. In the case of inter-band CA, different CCs belong to different bands, as we noted in §2; in our CA analysis, we classify each CA sample as 5G-low, 5G-mid, or 5G-mmWave based on the PCell. In contrast, in our MIMO analysis, we consider each cell separately, regardless of whether it is a PCell or an SCell.

4 Deployment Statistics

Even though 5G NR specifications define an ambitious PHY layer design, supporting up to 8 MIMO layers and 16 CCs within a single link, the extent to which these capabilities have materialized in real deployments remains unclear. In principle, such configurations would deliver multi-Gbps rates and unprecedented spectral efficiency. In practice, however, feature adoption often lags behind standardization: network rollouts are incremental, hardware support evolves across device generations, and operator policies often favor stability over aggressive feature extension. As 5G is passing the midpoint of its deployment cycle, it is timely to ask how close commercial networks have come to realizing the standard's full promise. In this section, we examine the real-world deployments of the three major US operators, analyzing the distributions of MIMO modes, MIMO layers, and CA configurations across frequency bands and link directions.

4.1 MIMO Deployment Statistics

Figs. 2a-2c show the distribution of downlink MIMO modes across operators and frequency bands during downlink throughput tests. Our first observation is a clear standard–deployment gap. Although NR permits up to 8 downlink layers, we observe no 8×8 configurations; the highest mode seen in practice is 4×4. Deployment is also strongly frequency dependent: mid-band cells operate predominantly in 4×4 mode (minimum 96% in the n2/n66 bands for AT&T, n2 band Verizon, and n41/n25 bands for T-Mobile), whereas all observed low-band and mmWave cells use exclusively 2×2 MIMO. This band-dependent behavior aligns with physical realities. The low band is provisioned for coverage and control rather than peak throughput, and hence, higher spatial multiplexing is not prioritized. On the other hand, mmWave radios face higher path loss and strong spatial correlation due to directional transmissions [16, 17], which limit the gains of higher order MIMO. The mid band thus emerges as the practical middle ground for higher-order downlink MIMO.

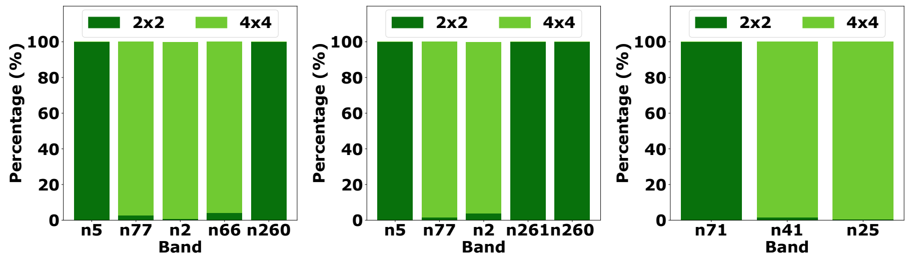

(a) AT&T (b) Verizon (c) T-Mobile

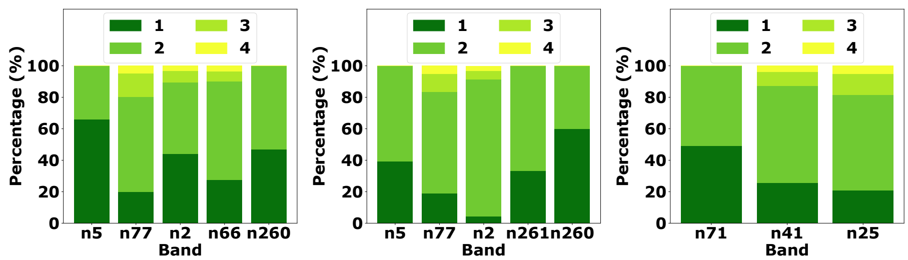

(d) AT&T. (e) Verizon. (f) T-Mobile.

Fig. 2: Distribution of MIMO modes and layers for different operators (DL).

While MIMO mode expresses the provisioned hardware configuration of a cell, MIMO layers report the number of spatial streams actually scheduled at transmission time, thus reflecting real utilization rather than configuration. The difference between the two arises from rank adaptation to instantaneous channel conditions (e.g., mobility, blockage) and scheduler policies. The layer distributions in Figs. 2d-2f show that scheduled rank is typically below the provisioned mode. In the low bands, where 2x2 MIMO is always available, AT&T uses two layers only about 34% of the time (n5), Verizon about 61% (n5), and T-Mobile about 51% of the time (n71). In the mid bands, where 4x4 MIMO is almost always available, we observe a similar trend. Two-layer operation dominates (roughly 60–90%) while three layers are used only 6-15% of the time and four layers only 3-5% of the time for different bands. Finally, in the mmWave bands, where, similar to low bands, 2x2 MIMO is always available, two-layer operation ranges from 40% (Verizon n260) to 67% (Verizon n261). Overall, our results show that actual layer usage remains conservative in practice, even when high order MIMO configurations are provisioned.

Figs. 3a-3c show that uplink MIMO configurations are far more conservative than downlink configurations. No 4×4 configurations are observed, and all FR1 bands across all three operators are exclusively SISO (single-input single output). MIMO is used only in the mmWave bands and is limited to 2x2 mode. Further, in contrast to downlink transmissions, here only a small fraction of the MIMO configurations actually support two different data streams (SMDL, two layers): around 45% in n261 for Verizon and roughly 15/35% in n260 for Verizon/AT&T.

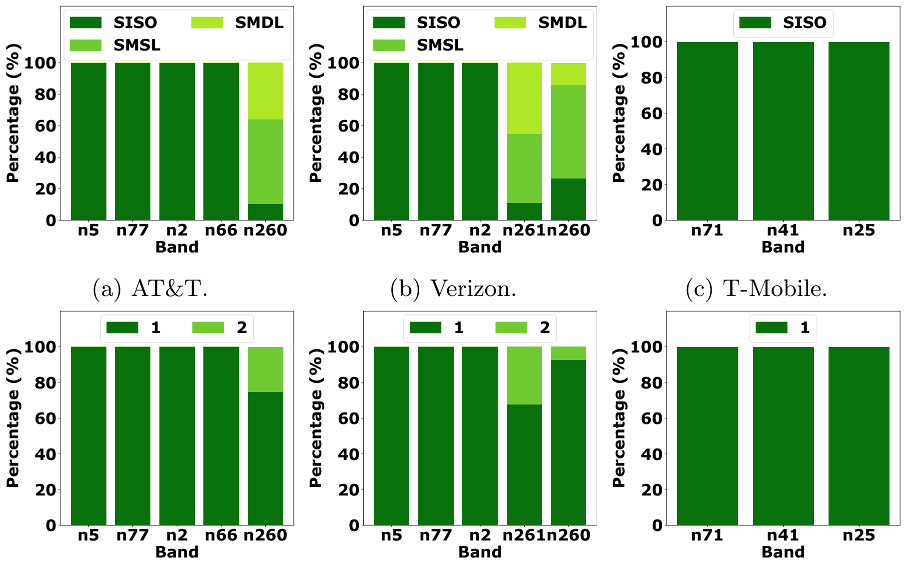

(a) AT&T. (b) Verizon. (c) T-Mobile.

(d) AT&T. (e) Verizon. (f) T-Mobile.

Fig. 3: Distribution of MIMO modes and layers for different operators (UL).

Most of the configurations support only one data stream (SMSL, one layer), which is transmitted over both antennas for diversity.

Uplink MIMO layer utilization also shows even more conservative configurations compare to downlink that mirror the distributions of uplink MIMO modes. Figs. 3d-3f show that in FR1, all operators use a single layer 100% of the time. Only mmWave occasionally schedules two layers (either two different data streams or a single stream duplicated over both antennas) as a minority share: 25% for AT&T n260, 32% for Verizon n261, 7.5% for Verizon n260.

4.2 CA Deployment Statistics

Fig. 4 summarizes the distribution of downlink CA. The highest level we observe is 8CC, appearing exclusively in the mmWave bands. Both AT&T n260 and Verizon n260 use 8CC most often (about 59% and 57%, respectively), while n261, operated exclusively by Verizon, shows a more balanced distribution among 8CC, 6CC, and 4CC. We also observe that in the mmWave bands, all three operators show preference to certain CA configurations, while other configurations are rarely used or not used at all. AT&T uses only 1CC, 4CC, and 8 CC. Verizon uses primarily 1CC, 2CC, and 8CC in n260 and rarely 5CC while in n261 it uses primarily 1CC, 4CC, 6CC, and 8CC, and rarely 2CC.

In contrast, in FR1 the maximum CA level is 4CA. AT&T and Verizon make little use of CA in low and mid bands, where most samples show use only 1 CC (i.e., no CA). The only notable exception is their mid band n77,

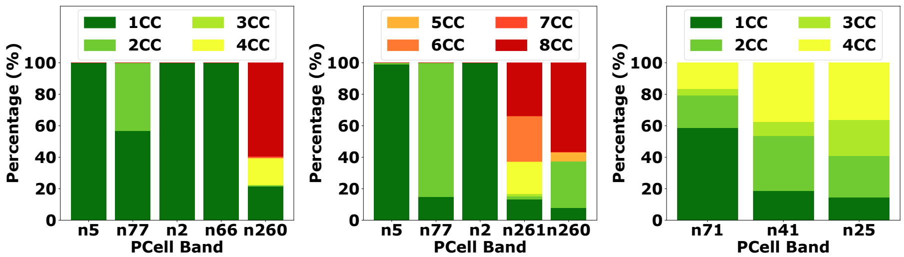

(a) AT&T. (b) Verizon. (c) T-Mobile.

Fig. 4: Distribution of CA for different operators (DL).

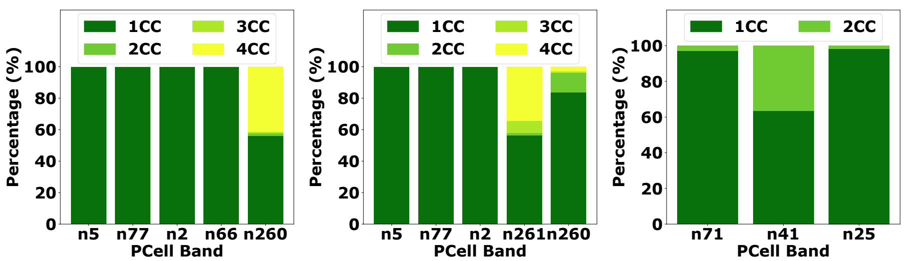

(a) AT&T. (b) Verizon. (c) T-Mobile.

Fig. 5: Distribution of CA for different operators (UL).

where AT&T uses 2CC about 43% of the time and Verizon about 86% of the time. T-Mobile, which does not deploy mmWave, but is the only operator that uses 5G SA, relies more heavily on sub-6 CA. In the mid bands n41 and n25, 3CC and 4CC (observed exclusively in SA mode [18]) account for a large share (for example, 4CC alone appears 37.5% in n41 and 36% in n25). These patterns reflect operator-specific preferences for downlink CA deployment: Verizon invests most heavily in mmWave, operating two FR2 bands with frequent high-order CA, whereas T-Mobile focuses on the mid bands, where the SA mode enables 3CC and 4CC configurations in contrast to AT&T and Verizon, which operate exclusively in NSA mode and use at most 2CC.

We also found that T-Mobile relies heavily on inter-band CA (56% of the samples), aggregating cells from up to 3 different bands. A similar observation was made in [18, 45]. In contrast, Verizon uses inter-band CA less than 0.02% of the time in FR1 but never in FR2 and AT&T uses only intra-band CA.

Fig. 5 shows uplink CA deployment statistics. Overall, uplink CA is limited compared to downlink CA. The highest CA level we observe is 4CC in the mmWave bands: AT&T n260 uses 4CC 41% of the time, Verizon n260 uses 4CC 3% of the time, and Verizon n261 uses 4CC 35% of the time. In FR1, AT&T and Verizon use no CA in the uplink, and T-Mobile shows only a modest 2CC share in n41 (37%), and a tiny 2CC share (3%) in n71 and n25.

Summary In summary, we observe a consistent standard–deployment gap. While NR allows up to 8 downlink MIMO layers, 4 uplink MIMO layers, and 16 CCs, we see at most 4x4 MIMO in the downlink and 2x2 MIMO in the uplink, and high-order CA (up to 8) appears only in the mmWave bands while FR1 CA remains modest, especially for Verizon and AT&T, which rely only on NSA. MIMO utilization is also limited: downlink transmissions rarely use more than two layers, whereas uplink transmissions use a single layer most of the time, even if higher order MIMO modes are supported. Deployment configurations also show clear asymmetry between the downlink and uplink directions, reflecting design choices driven by practical power constraints on the UE versus the base station side.

5 Disentangling Throughput Contributions of MIMO and CA

In this section, we dissect how much of the observed 5G throughput contribution originates from MIMO vs CA. In commercial networks, both technologies operate simultaneously, and their effects are intertwined at runtime. Each active CC contributes its own throughput based on its allocated bandwidth; each active CC may also independently employ a different number of MIMO layers as the wireless channel response varies across frequencies. This coupling makes it challenging to attribute the total throughput contribution to a single technology to evaluate their effectiveness. As such, to disentangle the individual contributions of the two technologies to the overall throughput, we decompose each throughput measurement sample into four distinct components, each capturing a specific contribution of the underlying technologies, as follows.

Let $T_i, L_i, i \in \{0, 1, \ldots, N\}$ denote the MAC throughput and the number of MIMO layers of CC i, respectively; $i = 0$ indexes the primary cell, and $i > 0$ indexes the $i - th$ secondary cell (recall from §4 that $N = 1$ in FR1 for AT&T and Verizon, $N = 3$ for T-Mobile in FR1, and $N = 7$ for AT&T and Verizon in FR2). We then define:

• **Baseline throughput** ($\mathbf{T}_{\text{BASE}}$): the throughput of the primary cell only, normalized by the number of MIMO layers active at runtime. It captures the contribution of a single spatial stream on the primary cell, without accounting for CA.

$$T_{\text{BASE}} = \frac{T_0}{L_0} \tag{1}$$

• **MIMO throughput** ($\mathbf{T}_{\text{MIMO}}$): the throughput under MIMO-enabled operation on the primary cell, reflecting the spatial multiplexing gain from transmitting multiple layers simultaneously.

$$T_{\text{MIMO}} = T_0 \tag{2}$$

• **CA throughput** ($\mathbf{T}_{\text{CA}}$): the throughput when CA is active, representing the combined throughput across all active CCs. The contribution of each CC

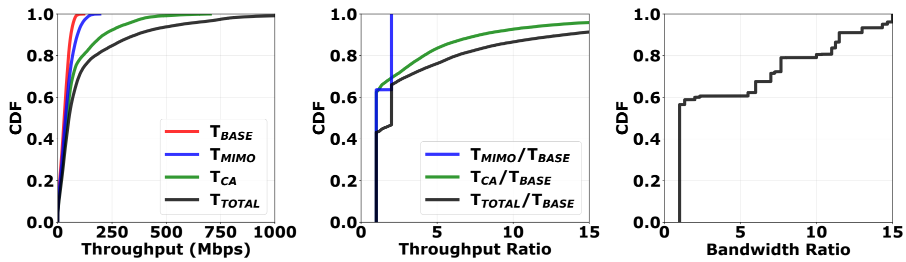

(a) Throughput. (b) Throughput ratio. (c) Bandwidth ratio.

Fig. 6: Low band throughput, throughput ratio, and bandwidth ratio (T-Mobile).

is calculated as the normalized per-layer throughput to remove the impact of MIMO.

$$T_{\text{CA}} = \sum_{i=0}^{N} \frac{T_i}{L_i} \tag{3}$$

• **Total throughput** ($\mathbf{T}_{\text{TOTAL}}$): the measured MAC-layer throughput, reflecting the combined effect of MIMO and CA.

$$T_{\text{TOTAL}} = \sum_{i=0}^{N} T_i \tag{4}$$

By comparing these quantities across frequency ranges (low, mid, and mmWave) and across operators, we assess how each mechanism contributes to throughput performance in real deployments.

For a more comprehensive view, we show in the following three types of cumulative distribution functions (CDFs) when analyzing each band. First, we show the CDFs of $\mathbf{T}_{\text{BASE}}$, $\mathbf{T}_{\text{MIMO}}$, $\mathbf{T}_{\text{CA}}$, and $\mathbf{T}_{\text{TOTAL}}$. Second, we include the CDFs of throughput ratios, $\mathbf{T}_{\text{MIMO}}/\mathbf{T}_{\text{BASE}}$, $\mathbf{T}_{\text{CA}}/\mathbf{T}_{\text{BASE}}$, and $\mathbf{T}_{\text{TOTAL}}/\mathbf{T}_{\text{BASE}}$, to better visualize the direct throughput enhancement effects contributed by MIMO, CA, and combination of the two technologies, respectively. Finally, we show the CDFs of the aggregate-to-primary-carrier bandwidth ratio to reveal how much spectrum is actually used under CA operation.

Since both MIMO and CA are used sparsely or not at all in the uplink direction (Figs. 3, 5), in the rest of this section and the next section, we limit our analysis to the downlink direction only.

5.1 Low Bands

Fig. 6 shows the CDFs of throughput, throughput ratio, and bandwidth ratio for T-Mobile in the low bands. Recall that T-Mobile is the only operator that employs CA in the low bands.

Fig. 6a shows that the throughput enhancement from MIMO and CA are limited in the lower percentile range. For example, the 25-th percentiles for

$\mathbf{T}_{\mathrm{MIMO}}$ and $\mathbf{T}_{\mathrm{CA}}$ are 15.7 Mbps and 19.6 Mbps, respectively, vs. 12.8 Mbps for $\mathbf{T}_{\mathrm{BASE}}$. Even in the median case, the improvement remains modest – 29.4 Mbps for $\mathbf{T}_{\mathrm{BASE}}$, 35.4 Mbps for $\mathbf{T}_{\mathrm{MIMO}}$, and 42.8 Mbps for $\mathbf{T}_{\mathrm{CA}}$. The gap widens only at higher percentiles; e.g., the 90-th percentiles for $\mathbf{T}_{\mathrm{BASE}}$, $\mathbf{T}_{\mathrm{MIMO}}$, and $\mathbf{T}_{\mathrm{CA}}$ are 60.4 Mbps, 90.6 Mbps, and 210.9 Mbps, respectively.

These trends reflect the sparse use of high-order MIMO and CA configurations at low frequencies: all low-band cells operate in 2×2 MIMO mode (Fig. 2c) with only a single layer being active about half the time (Fig. 2f), and about 60% of samples show no CA (Fig. 4c). The ratio plots in Fig. 6(b) mirror this behavior: $\mathbf{T}_{\mathrm{MIMO}}/\mathbf{T}_{\mathrm{BASE}}$ is limited to 1–2, reflecting the capped number of spatial layers in deployment, while $\mathbf{T}_{\mathrm{CA}}/\mathbf{T}_{\mathrm{BASE}}$ exhibits a much broader spread.

The broader spread of $\mathbf{T}_{\mathrm{CA}}/\mathbf{T}_{\mathrm{BASE}}$ in Fig. 6(b) reflects the flexibility of CA in real-world deployments. Unlike MIMO, whose gain is capped by the configured number of spatial layers, CA can exploit additional spectrum with varying bandwidth per added cell. T-Mobile often aggregates low band PCells with mid band SCells, leveraging wider bandwidth from higher-frequency bands while maintaining robust coverage on the low band. Fig. 6c shows the CDF of the total-to-primary bandwidth ratio: when aggregation is active, the effective bandwidth ratio ranges from 2× to 15×, with roughly 40% of samples exceeding 5×. This variability demonstrates the higher flexibility of CA in improving throughput whenever there is available spectrum, while MIMO focuses on maximizing spectral efficiency within a given bandwidth.

5.2 Mid Bands

Figs. 7a-7c, 7d-7f, 7g-7i show the CDFs of throughput, throughput ratio, and bandwidth ratio for AT&T, Verizon, and T-Mobile, respectively, in the mid band deployments. The three operators exhibit noticeably different preferences for enhancing throughput via MIMO and CA.

For AT&T, throughput enhancement is primarily driven by MIMO rather than CA (Fig. 7a). This trend aligns with the limited CA deployment previously shown in Fig. 4a. In the median case, throughput increases from 68.0 Mbps for $\mathbf{T}_{\mathrm{BASE}}$ to 131.7 Mbps for $\mathbf{T}_{\mathrm{MIMO}}$, while $\mathbf{T}_{\mathrm{CA}}$ only reaches 99.6 Mbps. Similarly, at the 90th percentile, $\mathbf{T}_{\mathrm{MIMO}}$ (402.8 Mbps) remains higher than $\mathbf{T}_{\mathrm{CA}}$ (245.8 Mbps). The corresponding median throughput ratios (Fig. 7b), $\mathbf{T}_{\mathrm{MIMO}}/\mathbf{T}_{\mathrm{BASE}} = 2.0$ and $\mathbf{T}_{\mathrm{CA}}/\mathbf{T}_{\mathrm{BASE}} = 1.0$, further highlight MIMO's dominant role in AT&T's mid-band performance. This imbalance stems from the restricted CA scaling flexibility seen in Fig. 7c: the bandwidth ratio takes only three values (1x, 2x, 3x), its median value is 1x, and it reaches its maximum value (3x) only at the 80-th percentile.

In contrast, the throughput enhancement contributions from MIMO and CA are much more balanced for Verizon and T-Mobile, as shown in Figs. 7d, 7g, respectively. For Verizon, the median throughput rises from 91.2 Mbps to 168.7 Mbps with MIMO and to 187.3 Mbps with CA, while for T-Mobile, the corresponding values are 87.2 Mbps, 150.3 Mbps, and 178.6 Mbps. The median throughput ratios (Figs 7e, 7h) reinforce this pattern: $\mathbf{T}_{\mathrm{CA}}/\mathbf{T}_{\mathrm{BASE}}$ equals 1.7

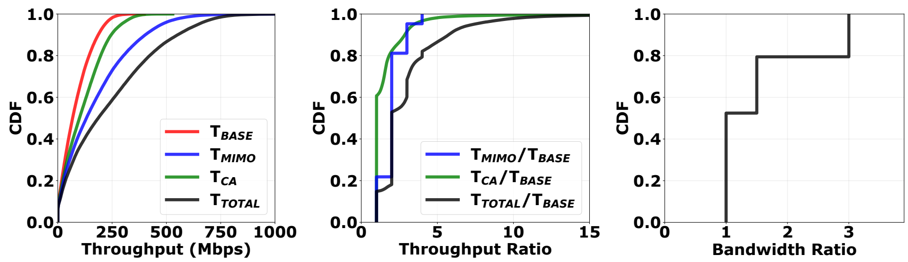

(a) Throughput. (b) Throughput ratio. (c) Bandwidth ratio.

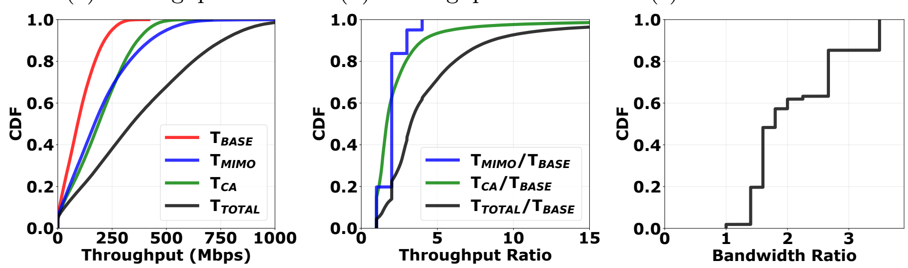

(d) Throughput. (e) Throughput ratio. (f) Bandwidth ratio.

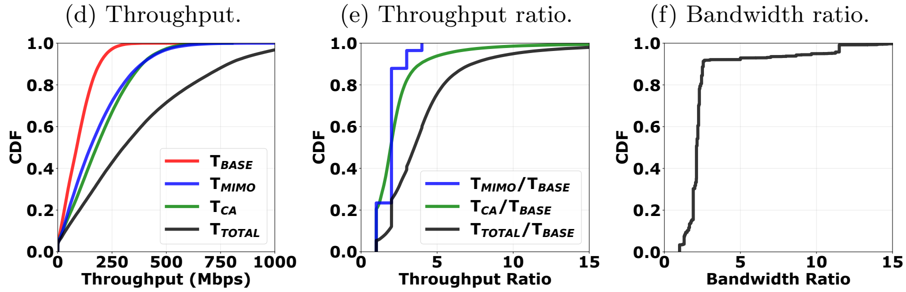

(g) Throughput. (h) Throughput ratio. (i) Bandwidth ratio.

Fig. 7: Mid band throughput, throughput ratio, and bandwidth ratio for AT&T (a)-(c), Verizon (d)-(f), T-Mobile (g)-(i).

for Verizon and 1.9 for T-Mobile, compared to only 1.0 for AT&T. These gains align with their more flexible CA configurations— Fig. 7i and Fig. 7f show significantly more fine-grained bandwidth ratio distributions than AT&T's, which has only three discrete levels (1×, 2×, 3×). The smoother CDF slopes for T-Mobile and Verizon reflect a larger number of unique aggregation configurations with varying total aggregated bandwidths, enabling more flexible spectrum use. This is particularly true for T-Mobile, due to its inter-band CA. Interestingly, the bandwidth ratio for T-Mobile can exceed 10x, thanks to the operator's 3CC and 4CC configurations, which are not supported by AT&T and Verizon. However, the 90-th percentile of the bandwidth ratio for T-Mobile is lower than for Verizon (2.5x vs. 3.5x), as most 3CC and 4CC configuration include low-

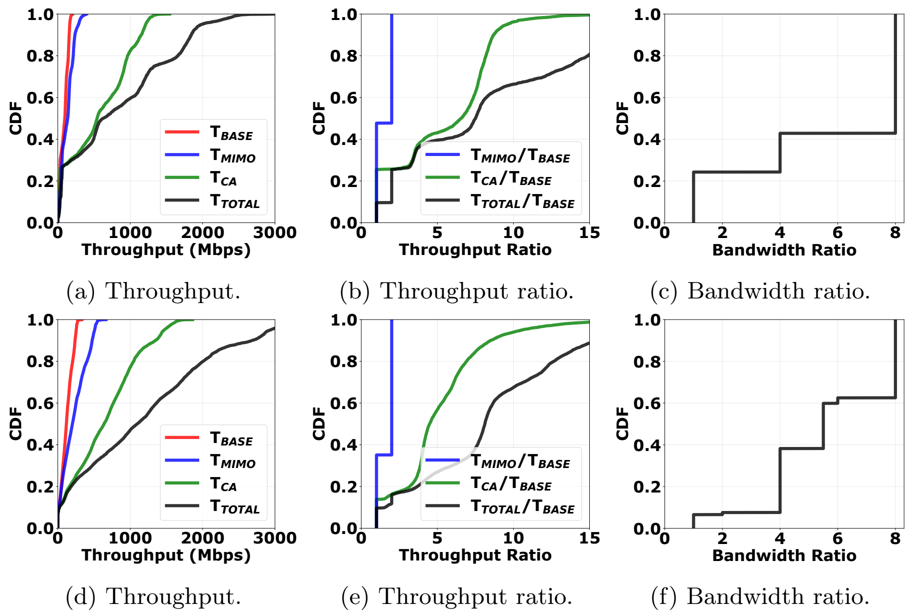

(a) Throughput. (b) Throughput ratio. (c) Bandwidth ratio.

(d) Throughput. (e) Throughput ratio. (f) Bandwidth ratio.

Fig. 8: mmWave throughput, throughput ratio, and bandwidth ratio for AT&T (a)-(c) and Verizon (d)-(f).

bandwidth SCells for robustness instead of high-bandwidth cells for throughput enhancement.

5.3 mmWave Bands

Figs. 8a-8c, 8d-8f show the CDFs of throughput, throughput ratio, and bandwidth ratio for AT&T and Verizon, respectively, in mmWave deployments. Unlike in the lower bands, here both operators rely almost entirely on CA for throughput enhancement, while MIMO plays a secondary role.

For AT&T (Fig. 8a), the median throughput increases from 96.7 Mbps for $\mathbf{T}_{\text{BASE}}$ to 137.6 Mbps with MIMO and to 542.2 Mbps with CA. Similarly, at the 90th percentile, $\mathbf{T}_{\text{CA}}$ (1151.3 Mbps) far exceeds $\mathbf{T}_{\text{MIMO}}$ (236.3 Mbps). The corresponding median throughput ratios (Fig. 8b) can reach up to 15x for CA but are limited to 2× for MIMO. Verizon follows a similar trend, but with slightly higher overall performance and more heterogeneity across CA deployments (compare Figs. 8f vs. 8c). The median throughput for Verizon ((Fig. 8d)) increases from 118.5 Mbps to 201.8 Mbps with MIMO and to 639.0 Mbps with CA, reaching 1348.9 Mbps at the 90th percentile. The throughput ratio (median 4.5×, 90-th percentile 8.5x) for CA (vs. 2× for MIMO) aligns with its bandwidth ratio distribution (median 5.5×, 90th percentile 8×; Fig. 8f), indicating a wider range of aggregated bandwidths across locations thanks to the higher diversity of CA configurations compared to AT&T (see Figs. 4b vs. 4a). Overall, the dominant role of CA in throughput enhancement in mmWave bands is expected, given the abundant spectrum available in these bands and the propagation characteristics of mmWave frequencies – highly directional links, strong path loss, and limited multipath richness, which restrict MIMO spatial multiplexing gains.

Summary Overall, we observe clear operator-specific differences in how throughput scales with MIMO and CA across frequency bands. (1) In the low bands (for T-Mobile only), throughput gains are modest because advanced configurations are rarely used: MIMO operation is typically limited to 2×2 and single-layer operation is often preferred. CA contributes mainly in the upper tail, where inter-band aggregation provides additional bandwidth. (2) In the mid bands, operator strategies diverge: AT&T's improvements are largely driven by MIMO, with CA constrained to a few discrete bandwidth ratios, whereas T-Mobile and Verizon achieve greater improvements from CA, reflecting more diverse aggregation settings and larger aggregate bandwidths. (3) In the mmWave bands, CA dominates as a throughput enhancing factor for both operators, combining multiple carriers to reach multi-Gbps rates, while MIMO gains remain bounded (at most 2×) due to limited multipath and highly directional propagation, which prevent deployments of higher-order MIMO modes.

6 Scaling Analysis

In this section, we evaluate the impact of higher-order configurations on throughput distributions in commercial 5G deployments. We observe how throughput evolves with increasing numbers of MIMO layers or CCs across frequency bands and operators to understand how spatial and spectral parallelism manifest in practice. This analysis provides an empirical view of how effectively 5G networks realize the theoretical scaling promised by high-rank MIMO and multi-carrier aggregation, and where those gains begin to saturate.

6.1 MIMO Scaling Analysis

Low Bands Figs. 9a, 9e, 9i show that two-layer MIMO substantially improves throughput across all operators. The median throughput rises from 9.4 to 30.2 Mbps for AT&T, 14.3 to 20.2 Mbps for Verizon, and 29.0 to 63.7 Mbps for T-Mobile, with corresponding maximum values of 59.9 to 115.5 Mbps, 68.5 to 131.5 Mbps, and 119.9 to 224.2 Mbps, respectively. Interestingly, while the gains at the maximum throughput values are close to the theoretical value of 2x for all three operators (1.87 for T-Mobile, 1.92 for AT&T and Verizon), the gains in the median case vary significantly: 1.4x for Verizon (much lower than the theoretical gain), 2.2x for T-Mobile (slightly higher than the theoretical gain), and 3.2x for AT&T (much higher than the theoretical gain).

Figs. 9d, 9h, 9l show that the cell bandwidth distributions are very similar for 1- and 2-layer configurations for all three operators, suggesting that they do not contribute to the differences in the observed throughput gains. However, Figs. 9c, 9g, 9k show that the median channel quality indicator (CQI) is higher

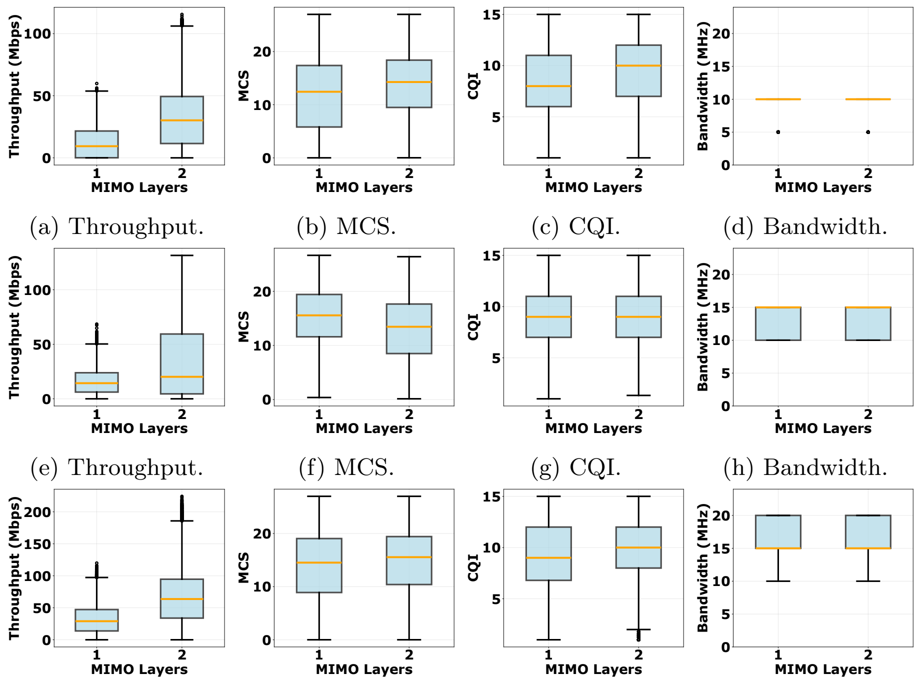

(a) Throughput. (b) MCS. (c) CQI. (d) Bandwidth.

(e) Throughput. (f) MCS. (g) CQI. (h) Bandwidth.

(i) Throughput. (j) MCS. (k) CQI. (l) Bandwidth.

Fig. 9: MIMO scaling analysis in the low bands for AT&T (a)-(d), Verizon (e)-(h), T-Mobile (i)-(l).

for 2-layer configurations compared to 1-layer configurations for AT&T and T-Mobile, but slightly lower for Verizon. In other words, AT&T and T-Mobile conservatively activate spatial multiplexing for high-quality links, while Verizon is more aggressive activating spatial multiplexing even for links of lower channel quality (see also Figs. 2d, 2f vs. Fig. 2e). As a result, the median MCS (modulation and coding schemes) is lower for 2-layer configurations compared to 1-layer configurations in the case of Verizon (Fig. 9f vs. Figs. 9b, 9j), canceling out part of the multiplexing gains, which explains the lower gains from MIMO in the case of Verizon.

Mid Bands Figs. 10a, 10e, 10i show that, in the mid bands, the throughput increases with the number of layers for up to 3 layers with AT&T and up to 2 layers with the other two operators, but decreases with higher layer configurations. The median throughput changes from 41 to 167 to 209 to 155 Mbps for AT&T as the number of layers changes from 1 to 4, from and 85 to 211 to 188 to 145 Mbps for Verizon, and from 63 to 151 to 103 to 74 Mbps for T-Mobile. The largest gains occur between one and two layers (4.1x, 2.48x, 2.4x in the median case for AT&T, Verizon, and T-Mobile, respectively, higher or much higher – for

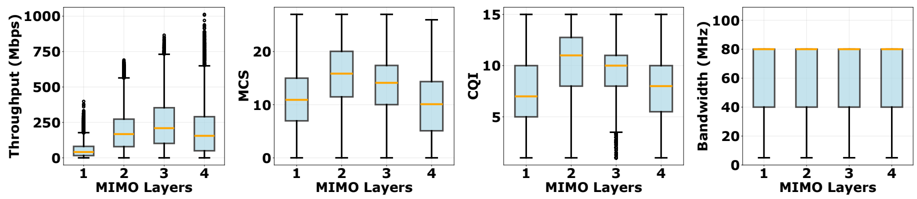

(a) Throughput. (b) MCS. (c) CQI. (d) Bandwidth.

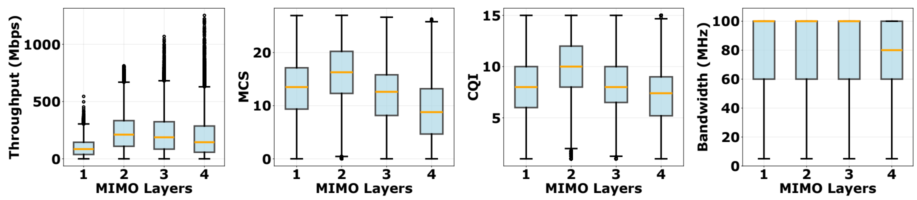

(e) Throughput. (f) MCS. (g) CQI. (h) Bandwidth.

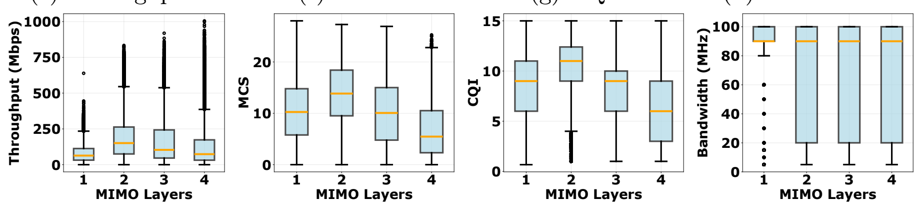

(i) Throughput. (j) MCS. (k) CQI. (l) Bandwidth.

Fig. 10: MIMO scaling analysis in the mid bands for AT&T (a)-(d), Verizon (e)-(h), T-Mobile (i)-(l).

AT&T – than the theoretical value of 2x), with AT&T showing slight additional improvement at three layers (1.25x over 2 layers in the median case, lower than the theoretical value of 1.5x).

Figs. 10d, 10h, 10l show that the cell bandwidth distributions remain essentially similar across ranks (except for T-Mobile's one-layer cases that generally use higher bandwidth), confirming that the gains still primarily arise from spatial multiplexing rather than additional spectrum. Figs. 10c, 10g, 10k and Figs. 10b, 10f, 10j show that CQI and MCS exhibit a similar peak at two layers across all operators, but start dropping at higher-order configurations, explaining the reduced gains at those configurations. Together the low band and mid band results suggest that two-layer MIMO represents a practical sweet spot for today's FR1 deployments; higher ranks tend to yield diminishing or negative returns, likely due to limited spatial separation and inter-layer interference under real-world conditions.

mmWave Bands Figs. 11a, 11e show that mmWave throughput scales strongly with the number of MIMO layers for both operators. The median throughput increases from 116 to 213 Mbps (1.8x) for AT&T and from 60 to 264 Mbps (4.4x)

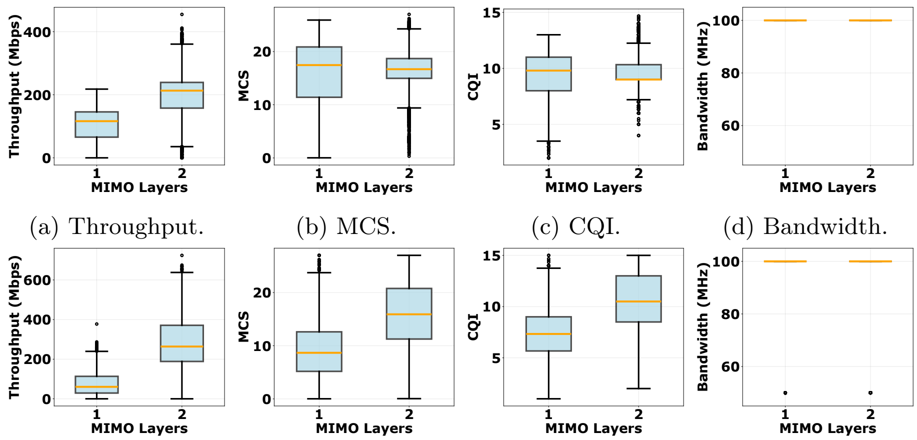

(a) Throughput. (b) MCS. (c) CQI. (d) Bandwidth.

(e) Throughput. (f) MCS. (g) CQI. (h) Bandwidth.

Fig. 11: MIMO scaling analysis in the mmWave bands for AT&T (a)-(d), Verizon (e)-(h).

for Verizon, while the maximum throughput roughly doubles (218 $\rightarrow$ 455 Mbps, 2.1x for AT&T and 377 $\rightarrow$ 724 Mbps, 1.9x for Verizon). The cell bandwidth remains constant at 100 MHz (Figs. 11d, 11h), isolating the observations from bandwidth scaling effects and indicating that the observed gain primarily arises from spatial multiplexing. However, AT&T's median CQI and MCS are similar under 1-layer and 2-layer configurations (Figs. 11c, 11b), indicating that its scheduler enables 2-layer transmission broadly, even without particularly favorable channel conditions. Verizon, in contrast, activates two layers mainly under high-CQI states (Fig. 11g), combining multiplexing gains with favorable channel conditions that support high MCS values (Fig. 11b) and reinforce throughput growth far beyond 3x. Overall, even without CQI improvement, AT&T's near-2x scaling demonstrates that spatial multiplexing at mmWave remains effective, though slightly weaker than in FR1, while Verizon's stronger scaling reflects the combined effects of beam alignment and selective rank scheduling.

6.2 CA Scaling Analysis

Low Bands Fig. 12 summarizes low-band CA scaling for T-Mobile. We focus on 1CC, 2CC, and 4CC configurations, as the 3CC case (only about 500 samples observed compared to more than 5000 for other configurations) appears to be a transient rather than a steady network state. Fig. 12a shows that the median throughput increases from 33.1 to 111.1 (3.6x gain) to 333.2 (10x gain) Mbps as CA expands from 1 to 2 to 4 CCs, i.e, it scales *superlinearly* with the number of carriers, driven by a corresponding aggregate bandwidth increase from 15 MHz to 115 MHz and 225 MHz (Fig. 12e) and by an increase to the fraction of higher-layer MIMO configurations (Fig. 12e). This is because secondary cells are often

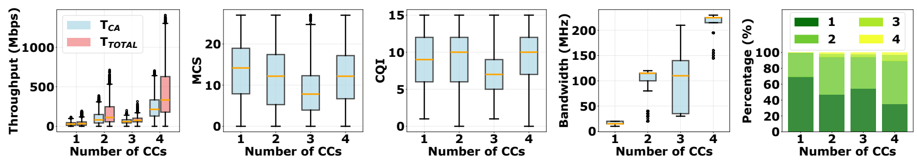

(a) Throughput. (b) MCS. (c) CQI. (d) Bandwidth. (e) MIMO.

Fig. 12: CA scaling analysis in the low bands for T-Mobile.

mid-band cells with much higher bandwidth and higher MIMO support than the low-band PCell. However, the throughput gain does not scale linearly with the aggregated bandwidth. When the CA level increases from 1CC to 2CC, the bandwidth gain is 7.7x, but the throughput gain is only 3.6x in the median case. One possible reason for this is the lower MCS in 2CC configurations (Fig. 12b) despite the higher CQI (Fig. 12c). On the other hand, when the CA level increases from 2CC to 4CC, the bandwidth gain is 1.96x, but the throughput gain is 3x in the median case, possibly due to higher MCS (at the 25-th percentile) and a higher fraction of 2 MIMO layers at 4CC configurations. To remove the effect of MIMO, we also plot $\mathbf{T}_{\text{CA}}$ (see §5 in Fig. 12a). The $\mathbf{T}_{\text{CA}}$ gains at 2CC vs. 1CC and at 4CC vs. 2CC are 2.8x and 2.6x, respectively, lower than the corresponding throughput gains.

Mid Bands Figs. 13a-13e, 13f-13j, 13k-13o summarize mid band CA scaling for AT&T, Verizon, and T-Mobile, respectively. Similar to the low band, throughput scales again superlinearly with the number of CCs for all three operators. For AT&T and Verizon, the median throughput increases from 105 Mbps (1CC) to 329 Mbps (2CC) and from 21 Mbps (1CA) to 401 Mbps (2CA), respectively, when the bandwidth increases for 80 MHz to 120 MHz and from 100 MHz to 160 MHz, respectively. Similarly, for T-Mobile, the throughput rises from 61 Mbps (1CC) to 257 Mbps (2CC), 364 Mbps (3CC), and 501 Mbps (4CC), while the aggregate bandwidth changes from 100 MHz to 190 MHz, 140 MHz, and 220 MHz, respectively. We observe that the throughput scales much faster than the bandwidth for all three operators; in fact, notice that, in the case of T-Mobile, the median aggregate bandwidth is lower at 3CC configurations compared to 2CC configurations and yet, the throughput is higher. The $\mathbf{T}_{\text{CA}}$ gains as the CA level increases from 1CC to 2CC for AT&T and Verizon and from 1CC to 4CC for T-Mobile are very similar to the corresponding throughput gains (e.g., 4.25x vs. 4.21x, 1.4x vs. 1.3x, and 1.37x vs. 1.4x for T-Mobile), suggesting that MIMO is not a contributing factor to the throughput scaling. For Verizon and T-Mobile, Figs. 13g, 13l show much higher MCS values at higher CA levels, which explain the superlinear throughput increase. In contrast, for AT&T, the MCS is already high at 1CC and only slightly increases at 2CC (Fig. 13b); this explains the lower throughput gains for AT&T compared to the other two operators as the CA levels rises from 1CC to 2CC.

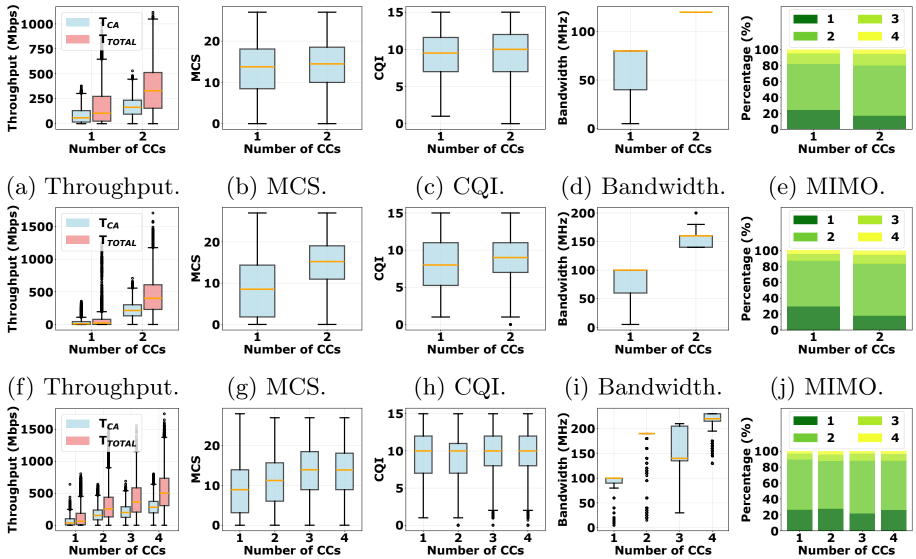

(a) Throughput. (b) MCS. (c) CQI. (d) Bandwidth. (e) MIMO.

(f) Throughput. (g) MCS. (h) CQI. (i) Bandwidth. (j) MIMO.

(k) Throughput. (l) MCS. (m) CQI. (n) Bandwidth. (o) MIMO.

Fig. 13: CA scaling analysis in the mid bands for AT&T (a)-(e), Verizon (f)-(j), T-Mobile (k)-(o).

mmWave Bands Figs. 14a–14e, 14f–14j summarize mmWave CA scaling for AT&T and Verizon, respectively. Here, almost all carriers have the same bandwidth (100 MHz, Figs. 14d, 14i), and hence, the aggregate bandwidth scales linearly with the CA level. Figs. 14a, 14f show that throughput scales superlinearly with the number of CCs for both operators, similar to the low and mid bands. For AT&T, the median throughput increases from 46.6 Mbps at 1CC to 499.1 Mbps at 4CC and 1252.5 Mbps at 8CC; for Verizon, the median total throughput increases from 4.2 Mbps at 1CC to 920.1 Mbps at 4CC, 1315.4 Mbps at 6CC, and 1713.5 Mbps at 8CC. For ATT, we observe a higher CQI (Fig. 14c) and a much higher MCS (Fig. 14b) for configurations that employ CA (4CC and 8CC) compared to single-carrier configurations (1CC). This explains the large throughput improvements at 4CC and 8CC despite the fact that the fraction of 2-layer configurations drops drastically at 4CC (Fig. 14e). For Verizon, CQI, MCS, and the fraction of 2-layer configurations are all much higher at 4CC-8CC compared to 1CC (Figs. 14h, 14g, 14j), which again explain the superlinear throughput increase with the number of CCs. However, in contrast to AT&T, the MCS for Verizon plateaus at 4CC and drops at 8CC; this explains the lower scaling rate for Verizon with higher CA levels (the throughput gain at 8CC over 4CC is 2.5x for AT&T but 1.86x for Verizon).

Summary Throughput scaling varies substantially across operators and between MIMO and carrier aggregation. For MIMO, FR1 deployments show consistent gains up to two layers and, in some cases, moderate improvement at

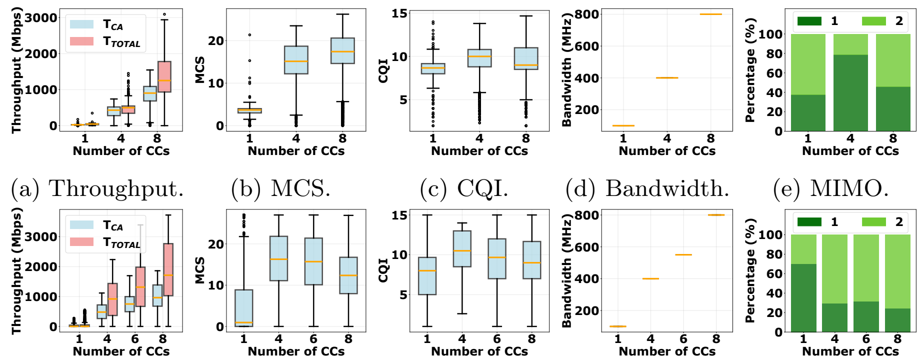

(a) Throughput. (b) MCS. (c) CQI. (d) Bandwidth. (e) MIMO.

(f) Throughput. (g) MCS. (h) CQI. (i) Bandwidth. (j) MIMO.

Fig. 14: CA scaling analysis in the mmWave bands for AT&T (a)-(e) and Verizon (f)-(j).

three layers, after which CQI and MCS degradation degradation lead to diminishing returns. In contrast, FR2 (mmWave) deployments exhibit strong but more variable scaling: AT&T enables higher-layer transmission broadly, yielding near-2× gains even without CQI improvement, whereas Verizon schedules additional layers selectively under high-CQI conditions, where multiplexing gains are combined with favorable channels to produce 4× throughput increases. For CA, throughput increases monotonically across all bands, typically superlinearly when higher-order CCs coincide with improved CQI and MCS, but rarely in direct proportion to aggregate bandwidth.

7 Related Work

5G measurements. Since its initial roll-out in 2019, numerous studies have evaluated various aspects of 5G networks including performance [4,6,7,12,18–21, 24,26,27,31,32,35,38,39,41,43,44,46], latency [13], handovers [15,22,23], resource allocation [9], power consumption [35], impact on application QoE [6, 25, 35], beamforming [10, 34], roaming [6, 14], multi-carrier access [8], and throughput prediction [5, 33, 45]. Among them, the works in [4, 11, 18–20, 22, 41] followed a methodology similar to ours, collecting datasets through cross-country drives, using multiple phones connected to different cellular networks and XCAL to log lower layer data. While some of these works briefly review CA [10, 18, 19, 34] and MIMO [19] deployments in operational 5G networks, none of them has attempted to systematically disentangle the relative throughput contributions of these two technologies in commercial 5G deployments.

CA and MIMO studies. The two works closest to our study are those in [40, 45]. Rochman et al. [40] use a dataset collected in 2022/2023 (when only T-Mobile supported CA in the low-/mid-bands) in Chicago and Minneapolis to study the factors that contribute to 5G performance improvements over 4G in low- and mid-bands, including channel bandwidths, CA, MIMO, and higher-order MCS. Their analysis shows that much of 5G's observed throughput gains over 4G stems from wider channel bandwidths, both over single and aggregated channels, rather than from PHY techniques, such as MIMO or higher-order MCS, alone. Ye et al. [45] focus exclusively on CA, characterizing its deployment in operational networks in two cities and developing a CA-aware predictor that improves throughput estimation. Both these studies provide valuable insights into CA behavior and MIMO deployments, although in a limited geographic area, but none of them quantifies the relative contributions of MIMO and CA to the overall 5G throughput or studies how throughput scales with each of these two factors. *Our work addresses precisely this gap by systematically comparing their throughput gains and scaling behavior across all three frequency bands via a large dataset obtained from a cross-country drive covering large urban centers as well as suburban and rural areas.*

Complementary to measurement studies, foundational theoretical and simulation based research has modeled the expected performance of MIMO and CA under idealized conditions. For MIMO, representative analytical works have characterized capacity scaling, spatial diversity, and antenna selection tradeoffs in multi-antenna systems [29,30,36]. For CA, system-level models have examined load balancing, packet scheduling, and capacity or admission control across aggregated carriers [42, 47]. These studies establish theoretical bounds and design principles under simplified assumptions such as perfect channel knowledge and stationary propagation conditions, but they do not capture the variability and constraints of commercial deployments that our measurement-driven analysis quantifies at scale. Different from these works, the work in [28] identifies experimentally limitations of the sequential, cell-by-cell CA operations employed by today's 5G networks and proposes PHY layer enhancements to enable inferring the channel quality of multiple carriers by measuring only one carrier. However, it neither considers MIMO nor does it explore the contribution of CA to the overall end user throughput.

8 Conclusion

We presented, to our best knowledge, the first quantitative comparison of the throughput gains attributable to MIMO and carrier aggregation (CA) in operational 5G networks. Our large-scale measurements show that real-world performance diverges largely from theoretical expectations. While both technologies enhance end user throughput, their scaling behaviors and deployment priorities differ. MIMO's spatial multiplexing gains often plateau and can even degrade beyond two to three layers as channel correlation and CQI loss offset the benefits of additional spatial streams. In contrast, CA yields steadier, often superlinear throughput growth across frequency bands, though the gains typically scale nonlinearly with aggregated bandwidth. Together, these findings expose two persistent gaps in today's 5G landscape: a theory–operation gap, where empirical scaling falls short of analytical models, and a deployment-standard gap, where commercial networks implement only a subset of 3GPP capabilities. These gaps, along with distinct operator strategies in balancing spectral expansion against hardware and signaling complexity, underscore the importance of understanding these trade-offs to guide spectrum policy, base station design, and the transition toward 6G.

Acknowledgement. We thank the anonymous reviewers for their helpful comments. This work is supported in part by NSF grants 2312834 and 2340283.

References

1. 3GPP: Release 15 Description; Summary of Rel-15 Work Items (TR 21.915) (2019), https://www.3gpp.org/DynaReport/21915.htm, 3rd Generation Partnership Project (3GPP)
2. XCAL Solo, https://accuver.com/sub/products/view.php?idx=11;
3. Dataset: Disentangling the Throughput Contributions of MIMO and Carrier Aggregation in 5G Networks (2025), https://github.com/NUWiNS/pam25-mimo-ca-dataset
4. Baena, E., Xu, Z., Ghoshal, M., Khan, I., Dinh, P., Kong, Z.J., Hu, Y.C., Koutsonikolas, D.: Root Cause Analysis of Cellular Network Throughput Degradations under Vehicular Mobility. In: Proc. of IEEE MASS (2025)
5. Basit, O., Dinh, P., Khan, I., Kong, Z.J., Hu, Y.C., Koutsonikolas, D., Lee, M., Liu, C.: On the Predictability of Fine-grained Cellular Network Throughput using Machine Learning Models. In: Proc. of IEEE MASS (2024)
6. Caso, G., Rajiullah, M., Kousias, K., Ali, U., Bouzar, N., Nardis, L.D., Brunstrom, A., Alay, O., Neri, M., di Benedetto, M.G.: The Chronicles of 5G Non-Standalone: An Empirical Analysis of Performance and Service Evolution. IEEE Open Journal of the Communications Society **5**, 7380–7399 (2024)
7. Caso, G., Rajiullah, M., Kousias, K., Ali, U., De Nardis, L., Brunstrom, A., Alay, O., Neri, M., Di Benedetto, M.G.: An Initial Look into the Performance Evolution of 5G Non-Standalone Networks. In: Proc. of IFIP/IEEE TMA (2023)
8. Chen, F., Ghoshal, M., Nan, E., Dinh, P., Khan, I., Kong, Z.J., Hu, Y.C., Koutsonikolas, D.: A Large-Scale Study of the Potential of Multi-carrier Access in the 5G Era. In: Proc. of PAM (2025)
9. Dinh, P., Ghoshal, M., Han, Y., Feng, Y., Koutsonikolas, D., Widmer, J.: Demystifying resource allocation policies in operational 5g mmwave networks. IEEE Transactions on Networking **33**(2), 1–16 (2024)
10. Feng, Y., Dinh, P., Ghoshal, M., Baena, E., Bouchebbah, H., Koutsonikolas, D.: Vivisecting Beam Management in Operational 5G mmWave Networks. In: Proc. of ACM CoNEXT (2025)
11. Fezeu, R., Carpenter, J., Fiandrino, C., Ramadan, E., Ye, W., Widmer, J., Qian, F., Zhang, Z.L.: Unveiling the 5G Mid-Band Landscape: From Network Deployment to Performance and Application QoE. In: Proc. of ACM SIGCOMM (2024)
12. Fezeu, R., Fiandrino, C., Ramadan, E., Carpenter, J., Freitas, L., Bilal, F., Ye, W., Widmer, J., Qian, F., Zhang, Z.L.: Unveiling the 5G Mid-Band Landscape: From Network Deployment to Performance and Application QoE. In: Proc. of ACM SIGCOMM 2024 (2024)
13. Fezeu, R., Ramadan, E., Ye, W., Minneci, B., Xie, J., Narayanan, A., Hassan, A., Qian, F., Zhang, Z.L., Chandrashekar, J., Lee, M.: An In-Depth Measurement

Analysis of 5G mmWave PHY Latency and its Impact on End-to-End Delay. In: Proc. of PAM (2023)
14. Fezeu, R.A.K., Fiandrino, C., Ramadan, E., Carpenter, J., Chen, D., Tan, Y., Qian, F., Widmer, J., Zhang, Z.L.: Roaming across the European Union in the 5G Era: Performance, Challenges, and Opportunities. In: Proc. of IEEE INFOCOM (2024)
15. Fiandrino, C., Juárez Martínez-Villanueva, D., Widmer, J.: Uncovering 5G Performance on Public Transit Systems with an App-based Measurement Study. In: Proc. of ACM MSWiM (2022)
16. Ghasempour, Y., Haider, M.K., Cordeiro, C., Koutsonikolas, D., Knightly, E.: Multi-Stream Beam-Training for mmWave MIMO Networks. In: Proc. of ACM MobiCom (2018)
17. Ghasempour, Y., Knightly, E.: Decoupling Beam Steering and User Selection for Scaling Multi-User 60 GHz WLANs. In: Proc. of ACM MobiHoc (2018)
18. Ghoshal, M., Basit, O., Wang, S., Dinh, P., Khan, I., Feng, Y., Yu, Z., Hu, Y.C., Koutsonikolas, D.: A first large-scale study of operational 5g standalone networks. In: Proc. of ACM CoNEXT (2025)
19. Ghoshal, M., Basit, O., Khan, I., Kong, Z.J., Wang, S., Feng, Y., Dinh, P., Hu, Y.C., Koutsonikolas, D.: Replication: Performance of cellular networks on the wheels. In: Proc. of ACM IMC (2025)
20. Ghoshal, M., Khan, I., Kong, Z.J., Dinh, P., Meng, J., Hu, Y.C., Koutsonikolas, D.: Performance of cellular networks on the wheels. In: Proc. of ACM IMC (2023)
21. Ghoshal, M., Kong, Z.J., Xu, Q., Lu, Z., Aggarwal, S., Khan, I., Li, Y., Hu, Y.C., Koutsonikolas, D.: An In-Depth Study of Uplink Performance of 5G mmWave Networks. In: Proc. of the ACM SIGCOMM 5G-MeMU Workshop (2022)
22. Hassan, A., Narayanan, A., Zhang, A., Ye, W., Zhu, R., Jin, S., Carpenter, J., Mao, Z.M., Qian, F., Zhang, Z.L.: Vivisecting mobility management in 5G cellular networks. In: Proc. of ACM SIGCOMM (2022)
23. Kalntis, M., Suárez-Varela, J., Iglesias, J.O., Bhattacharjee, A.K., Iosifidis, G., Kuipers, F.A., Lutu, A.: Through the Telco Lens: A Countrywide Empirical Study of Cellular Handovers. In: Proc. of ACM IMC (2024)
24. Khan, I., Ghoshal, M., Angjo, J., Dimce, S., Hussain, M., Parastar, P., Yu, Y., Fiandrino, C., Orfanidis, C., Aggarwal, S., Aguiar, A., Alay, O., Chiasserini, C., Dressler, F., Hu, Y., Ko, S., Koutsonikolas, D., Widmer, J.: How mature is 5g deployment? a cross-sectional, year-long study of 5g uplink performance. In: Proceedings of the IFIP/IEEE Networking 2024 Conference (2024)
25. Khan, I., Tran, T.X., Hiltunen, M., Karagioules, T., Koutsonikolas, D.: An Experimental Study of Low-Latency Video Streaming over 5G. In: Proc. of IEEE MeditCom (2024)
26. Kousias, K., Rajiullah, M., Caso, G., Alay, O., Brunstorm, A., Nardis, L.D., Neri, M., Ali, U., Benedetto, M.G.D.: Coverage and Performance Analysis of 5G Non-Standalone Deployments. In: Proc. of ACM WiNTECH (2022)
27. Kousias, K., Rajiullah, M., Caso, G., Alay, O., Brunstrom, A., Ali, U., Nardis, L.D., Neri, M., Benedetto, M.G.D.: Empirical performance analysis and ML-based modeling of 5G non-standalone networks. Elsevier Computer Networks **241**(C) (2024)
28. Li, Q., Zhang, Z., Liu, Y., Tan, Z., Peng, C., Lu, S.: CA++: Enhancing Carrier Aggregation Beyond 5G. In: Proc. of ACM MobiCom (2023)
29. Liu, Z., Dai, L.: A comparative study of downlink mimo cellular networks with co-located and distributed base-station antennas. IEEE Transactions on Wireless Communications **13**(11), 6259–6274 (2014)

30. Molisch, A.F., Win, M.Z., Choi, Y.S., Winters, J.H.: Capacity of mimo systems with antenna selection. IEEE Transactions on Wireless Communications **4**(4), 1759–1772 (2005)
31. Moreno, J., Contini, M., Aguiar, A.: 5G NSA Performance: A Measurement Study. In: Proc. of IFIP WONS (2024)
32. Narayanan, A., Ramadan, E., Carpenter, J., Liu, Q., Liu, Y., Qian, F., Zhang, Z.L.: A first look at commercial 5G performance on smartphones. In: Proc. of ACM WWW (2020)
33. Narayanan, A., Ramadan, E., Mehta, R., Hu, X., Liu, Q., Fezeu, R.A., Dayalan, U.K., Verma, S., Ji, P., Li, T., et al.: Lumos5G: Mapping and predicting commercial mmWave 5G throughput. In: Proc. of ACM IMC (2020)
34. Narayanan, A., Rochman, M.I., Hassan, A., Firmansyah, B.S., Sathya, V., Ghosh, M., Qian, F., Zhang, Z.L.: A Comparative Measurement Study of Commercial 5G mmWave Deployments. In: In Proc. of IEEE INFOCOM (2022)
35. Narayanan, A., Zhang, X., Zhu, R., Hassan, A., Jin, S., Zhu, X., Zhang, X., Rybkin, D., Yang, Z., Mao, Z.M., Qian, F., Zhang, Z.L.: A Variegated Look at 5G in the Wild: Performance, Power, and QoE Implications. In: Proc. of ACM SIGCOMM (2021)
36. Ngo, H.Q., Larsson, E.G., Marzetta, T.L.: Energy and spectral efficiency of very large multiuser mimo systems. IEEE Transactions on Communications **61**(4), 1436–1449 (2013)
37. nuttcp - Network Performance Measurement Tool, https://www.nuttcp.net
38. Parastar, P., Alay, A.L.O.O., Caso, G., Perino, D.: Spotlight on 5G: Performance, Device Evolution and Challenges from a Mobile Operator Perspective. In: Proc. of IEEE INFOCOM (2023)
39. Rochman, M.I., Sathya, V., Fernandez, D., Nunez, N., Ibrahim, A.S., Payne, W., Ghosh, M.: A comprehensive analysis of the coverage and performance of 4G and 5G deployments. Computer Networks **237** (2023)
40. Rochman, M.I., Ye, W., Zhang, Z.L., Ghosh, M.: A Comprehensive Real-World Evaluation of 5G Improvements Over 4G in Low- and Mid-Bands. IEEE Transactions on Cognitive Communications and Networking **11**(3), 1427–1440 (2025)
41. Wang, S., Ghoshal, M., Feng, Y., Khan, I., Dinh, P., Basit, O., Yu, Z., Hu, Y.C., Koutsonikolas, D.: A first large-scale study of operational 5g standalone networks. In: Proc. of ACM SIGMETRICS (2026)
42. Wang, Y., Pedersen, K.I., Sørensen, T.B., Mogensen, P.E.: Carrier load balancing and packet scheduling for multi-carrier systems. IEEE Transactions on Wireless Communications **9**(5), 1780–1789 (2010)
43. Xu, D., Zhou, A., Zhang, X., Wang, G., Liu, X., An, C., Shi, Y., Liu, L., Ma, H.: Understanding operational 5G: A first measurement study on its coverage, performance and energy consumption. In: Proc. of ACM SIGCOMM (2020)
44. Ye, W., Carpenter, J., Zhang, Z., Fezeu, R.A.K., Qian, F., Zhang, Z.L.: A closer look at stand-alone 5g deployments from the ue perspective. In: Proc. of IEEE MeditCom (2023)
45. Ye, W., Hu, X., Sleder, S., Zhang, A., Dayalan, U.K., Hassan, A., Fezeu, R.A.K., Jajoo, A., Lee, M., Ramadan, E., Qian, F., Zhang, Z.L.: Dissecting Carrier Aggregation in 5G Networks: Measurement, QoE Implications and Prediction. In: Proc. of ACM SIGCOMM (2024)
46. Yuan, X., Wu, M., Wang, Z., Zhu, Y., Ma, M., Guo, J., Zhang, Z.L., Zhu, W.: Understanding 5G performance for real-world services: a content provider's perspective. In: Proc. of ACM SIGCOMM (2022)

47. Zhang, R., Zheng, Z., Wang, M., Shen, X., Xie, L.L.: Equivalent Capacity in Carrier Aggregation-Based LTE-A Systems: A Probabilistic Analysis. IEEE Transactions on Wireless Communications **13**(11), 6444–6460 (2014)