Beyond QCI: Understanding Performance Competition Between an MVNO and Its Host MNO

Zeyu Li
Northeastern University
Boston, Massachusetts, USA
li.z29@northeastern.edu

Yufei Feng
Northeastern University
Boston, Massachusetts, USA
feng.yuf@northeastern.edu

Dimitrios Koutsonikolas
Northeastern University
Boston, Massachusetts, USA
d.koutsonikolas@northeastern.edu

ABSTRACT

Mobile virtual network operators (MVNOs) provide service over host mobile network operators (MNOs), yet users often assume that host carriers receive better performance. This paper presents, to our knowledge, the first simultaneous measurement study of MVNO–host MNO throughput competition with radio-side visibility. Across three operator families, five competition scenarios, and three U.S. cities, we run simultaneous TCP throughput tests over commercial subscriptions and collect cross-layer indicators including QCI, PCI, RAT mode, SA/NSA mode, resource block allocation, MCS, and MAC layer NR/LTE throughput.

Our results show that host advantage is conditional rather than universal, and that no single factor consistently explains host–MVNO throughput ordering. Across operator families, observed performance differences are associated with different combinations of QoS class, RAT attachment, serving-cell context, radio-resource allocation, frequency band, and SA/NSA architecture. Under identical radio contexts, QoS-class ordering often coincides with downlink performance ordering. In contrast, when the radio contexts are different, RAT mismatch, cell attachment, or NSA LTE aggregation dominates the observed outcome. City-level comparisons further show that operator ordering can change across locations as RAT attachment, serving-cell selection, bandwidth, and radio conditions vary. Thus, host–MVNO performance must be interpreted through the joint radio and access context rather than from operator identity or any single QoS indicator.

CCS CONCEPTS

• **Networks** → **Network measurement**.

KEYWORDS

5G, MVNO, cellular measurement, cross-layer measurement, quality of service, radio access networks

ACM Reference Format:
Zeyu Li, Yufei Feng, and Dimitrios Koutsonikolas. 2026. Beyond QCI: Understanding Performance Competition Between an MVNO and Its Host MNO. In *Proceedings of the 2026 ACM Internet Measurement Conference (IMC '26), October 12–16, 2026, Karlsruhe, Germany*. ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3777912.3839808

IMC '26, October 12–16, 2026, Karlsruhe, Germany

ACM ISBN 979-8-4007-2327-8/26/10
https://doi.org/10.1145/3777912.3839808

1 INTRODUCTION

Mobile users increasingly encounter service plans that advertise differentiated treatment through terms such as "premium data", "premium network access", or "high-priority data" [7, 46, 48]. These offerings raise a broader question for mobile virtual network operators (MVNOs), which rely on a host mobile network operator (MNO): when host and MVNO subscribers use the same network service at the same place and time, do they experience systematically different performance, and what network factors are associated with those differences?

Answering this question requires more than application-layer speed tests. End-to-end throughput can be shaped by subscriber QoS configuration, radio access technology (RAT), serving-cell attachment, spectrum band and bandwidth, channel conditions, resource allocation, link adaptation, and 5G architecture (Standalone/SA vs. Non-Standalone/NSA). These factors may vary even between co-located devices using the same host network. One observable component of this configuration is the Quality of Service (QoS) class, represented by the QoS Class Identifier (QCI) in 4G and 5G QoS Identifier (5QI) in 5G [11]. Online discussions often use these values as a proxy for subscriber priority [42, 43], but QoS class is only one input to resource treatment and does not directly determine delivered throughput.

Despite extensive work on commercial cellular performance, there is limited measurement evidence showing how the performance of MVNO subscribers compares against the performance of host-MNO subscribers when the two are measured simultaneously, and even less evidence linking those end-to-end outcomes to radio-side context. This paper fills this gap by examining how MVNOs compete with their host MNOs in commercial deployments through four research questions: **RQ1:** How does an MVNO compare with its host MNO in terms of downlink and uplink throughput under matched, simultaneous measurements? **RQ2:** How are observed QoS class, radio-resource allocation, and throughput related when operators use comparable serving contexts? **RQ3:** How do access differences such as RAT attachment, serving-cell selection, frequency band, and 5G architecture relate to host–MVNO performance ordering? **RQ4:** How do these performance and access patterns vary across operator families and measurement locations? Answering these questions requires concurrent measurements that expose both application-layer throughput and lower-layer radio indicators.

We present, to our knowledge, the first cross-layer measurement study of simultaneous throughput competition between an MVNO and its host MNO. Across three operator families, five competition scenarios, and three U.S. cities, we run identical workloads in parallel on commercial host-MNO and MVNO subscriptions. We collect application-layer throughput and radio indicators, including 5QI, channel quality indicator (CQI), modulation and coding scheme (MCS), resource block (RB) allocation, Physical Cell Identifier (PCI), RAT mode, and MAC layer 5G/LTE throughput. This design enables per-run comparison of competitive outcomes and attribution of performance differences to their access context.

Our measurements reveal that host–MVNO performance ordering is conditional and varies across operator families, traffic directions, and access contexts. With a 10% margin threshold, MVNOs achieve higher win rates than their host in a substantial fraction of the measured downlink and uplink scenarios. Additionally, no single observed factor consistently explains these outcomes. Depending on the measurement context, throughput differences are associated with QoS class, RAT attachment, serving-cell selection, spectrum bandwidth, channel conditions, resource allocation, and SA/NSA architecture. In particular, QoS-class ordering is more closely aligned with throughput ordering in some matched serving-cell cases, whereas other access differences become more prominent when operators use different cells or radio configurations.

In summary, the paper makes the following contributions:

- We conduct a cross-layer measurement study of simultaneous throughput competition between MVNOs and host MNOs using identical workloads.
- We provide a cross-layer MVNO measurement dataset that combines application-layer throughput with radio-side indicators, including 5QI, PCI, RB allocation, MCS, and MAC-layer NR/LTE throughput. To enable reproducibility, we make our dataset publicly available [40].
- We characterize how different cross-layer factors are associated with host–MVNO performance outcomes. Depending on the operator family and traffic direction, the dominant observed correlate can be QoS class, resource allocation, RAT attachment, serving-cell/bandwidth differences, or 5G architecture mode.
- We show that host-MNO advantage is conditional rather than universal and that the dominant performance correlate differs across operator families, traffic directions, and locations.

2 BACKGROUND

MVNO–Host MNO Service Path. Mobile service follows an end-to-end path from user equipment (UE), through the radio access network (RAN) and core network, to external data networks. An MVNO uses some or all of a host MNO's infrastructure, so it may share the same RAN while differing in subscriber provisioning, policy configuration, Quality of Service (QoS), and traffic handling. The RAN segment includes radio-related factors such as channel quality, scheduling, RAT mode, and serving-cell association, while the network/service segment includes provisioning, QoS policy, and traffic handling. Therefore, sharing radio access infrastructure does not guarantee identical end-user performance.

QCI and 5QI. QCI is standardized in LTE/Evolved Packet System (EPS) as a scalar QoS identifier, with each value mapped to a resource type, priority level, packet delay budget, and packet error loss rate [1]. In the 5G Core Network (5GC), the corresponding identifier is 5G QoS Identifier (5QI). For standardized values 1–9, 3GPP TS 23.501 maps 5QI directly to the same QCI values and QoS characteristics [2]. Each 5QI/QCI value has a unique priority set by 3GPP. For the 5QI/QCI values 6–9 that are involved in this paper, a smaller value indicates a higher priority. In the remainder of the paper, we use QCI for NSA/EPS, 5QI for SA/5GC, and QoS class when referring to both collectively.

Radio Access Modes and Indicators. Cellular throughput depends on access configuration and radio conditions. In SA 5G, the UE connects to the 5G core via 5G New Radio (NR). In NSA 5G, NR is deployed with LTE anchoring, and user traffic can use both NR and LTE.

PCI identifies the serving physical cell. RSRP measures reference-signal strength, SINR reflects signal quality relative to interference and noise, CQI indicates channel quality for link adaptation, and Carrier Aggregation (CA) indicates how many carriers are serving for the running UE. At the radio layer, the scheduler determines when and how many radio resources are allocated to a UE. The resulting throughput further depends on factors such as channel quality, MCS, spatial transmission configuration, and available spectrum bandwidth. Together, these indicators describe the radio context that shapes observed throughput.

3 METHODOLOGY

Measurement setup and matched comparison design. We used commercial subscriptions from the three major U.S. host MNOs, anonymized as HA, HB, and HC, and their MVNOs. For each host, we selected two MVNOs, denoted VA1/VA2, VB1/VB2, and VC1/VC2, for a total of nine operator subscriptions.

We selected representative MVNOs for each host network using a public U.S. MVNO ranking by estimated customer size [15]. For all operators, we used the highest-tier prepaid plan available at the time of measurement. We focus on prepaid plans because some MVNOs lack postpaid options; using a single subscription type avoids confounding from prepaid–postpaid differences. Table 9 in Appendix B summarizes the subscriptions and plan details. The plan descriptions are taken from operator-advertised material and are not related to the exact RAN scheduling priority.

Our testbed consisted of three Samsung Galaxy S24 smartphones and three XCAL Solo devices [5]. Each phone was paired with one XCAL Solo to collect radio-side measurements. During each test, the three phones were placed side by side and kept stationary.

We designed five test scenarios to compare operators under different contention settings. **Scenario 1:** the host MNO and its two MVNOs run sequentially, each for one minute. **Scenario 2:** the host MNO and virtual operator 1 run simultaneously for one minute. **Scenario 3:** the host MNO and virtual operator 2 run simultaneously for one minute. **Scenario 4:** virtual operator 1 and virtual operator 2 run simultaneously for one minute. **Scenario 5:** the host MNO, virtual operator 1, and virtual operator 2 run simultaneously for one minute. These scenarios are applied to each operator family, i.e., HA/VA1/VA2, HB/VB1/VB2, and HC/VC1/VC2.

We conducted the primary campaign for all three operator families in Boston and collected additional measurements for Families A and B in Atlanta and Family C in Philadelphia. Within each city, we selected 2–5 stationary measurement sites in busy public areas with available service for the participating operators. The choice of the second city was constrained by our travel schedule and subscription/data-plan availability rather than by a controlled matching of radio conditions across cities. We therefore used the multi-city measurements to evaluate geographic variation rather than as strictly controlled city-to-city experiments. At each location, we repeated multiple rounds of downlink and uplink throughput tests. Each round executed scenarios S1 to S5 in order, and for each scenario, we conducted one downlink run and one uplink run back to back. In a whole test round for one operator, there are 8 test runs in total. The measurement locations and the time range of the measurements are shown in Fig. 22 and Table 10 in Appendices C and D.

We used `iperf3` [17] in TCP mode with a Google Cloud VM as the server to generate end-to-end sustained bulk data traffic. We installed `Termux` on each phone and ran synchronization scripts to align the start time of simultaneous tests. This design keeps location, device model, test timing, traffic workload, and measurement pipeline aligned whenever possible, while acknowledging that complete control is infeasible in commercial networks.

Our analysis combines application-layer throughput with radio-side measurements. At the application layer, we collected downlink and uplink throughput samples at a granularity of 500 ms using `iperf3`. On the radio side, XCAL logs QoS class, CQI, MCS, allocated RB/PRB counts, PCI and serving-cell information, RAT mode, and other scheduling and resource-usage indicators. For NSA measurements, the PCI used in our cell-attachment analysis refers to the NR serving-cell PCI, not the LTE anchor PCI. We use NR PCI together with channel number (NR-ARFCN) to identify matched serving contexts for resource-allocation analysis.

Data filtering. Some host and MVNO prepaid plans have monthly data limits (see Table 9). Exceeding these limits triggers a hard cap, reducing both downlink and uplink throughput to below 200 kbps. This is nearly three orders of magnitude lower than normal throughput (tens to hundreds of Mbps) and would dominate the comparison. We therefore exclude rounds affected by this plan-level cap from the analysis. After filtering, each operator has 20–50 valid runs per scenario, each run with about 120 throughput samples.

Cross-layer analysis strategy. We analyze each operator family separately and first characterize application-layer throughput outcomes and RAT attachment. We then distinguish two types of comparison. For measurements with matched serving contexts, we jointly examine observed QoS class, radio-resource allocation, channel indicators, MCS, and throughput to understand how resource allocation relates to delivered performance. A direct RB-sharing interpretation is restricted to cases with matching NR PCI and NR-ARFCN. For measurements with different RATs, serving cells, frequency bands, or architecture modes, we treat the results as access differentiation rather than direct competition for the same radio-resource pool. These comparisons are used to characterize how different access configurations coincide with different user-perceived throughput. Because our measurements are collected from UEs in a commercial network, we report these factors as observed associations and do not attempt to infer the operator's internal scheduler logic or intentional steering policy.

4 CROSS-FAMILY OVERVIEW

We present a high-level comparison of throughput performance across three operator families. Each family consists of one host MNO (H-class) and two MVNOs (V-class) that share the network

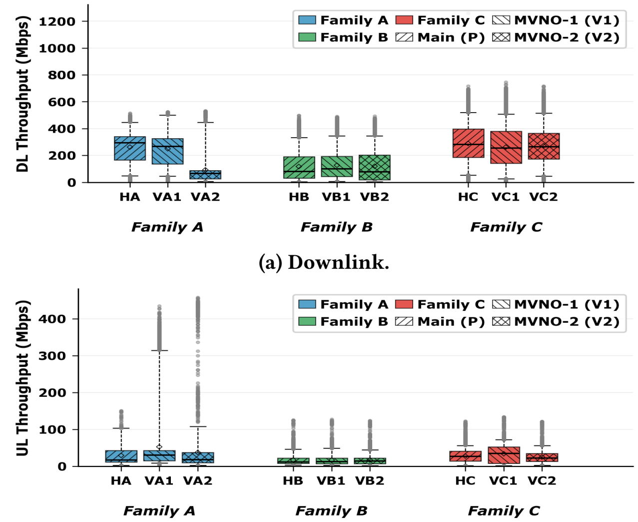

(a) Downlink.

(b) Uplink.

Figure 1: Cross-family throughput distributions aggregating all valid measurement scenarios and locations. Whiskers extend to the 5th and 95th percentile; the diamond marker (⋄) denotes the mean.

infrastructure: Family A (HA, VA1, VA2), Family B (HB, VB1, VB2), and Family C (HC, VC1, VC2).

4.1 Aggregate Throughput Distributions

Fig. 1 shows the aggregated throughput distributions for all operators, combining measurements across every scenario and location, and including all application-layer throughput samples. The distributions reveal that operator families differ substantially in their overall throughput levels. Family C achieves the highest median downlink throughput, followed by Family A and Family B, reflecting differences in network technology deployment and spectrum allocation across carriers. Within each family, H-class and V-class operators exhibit overlapping distributions, indicating that MVNOs can achieve throughput comparable to their host operator under certain conditions. However, overlap in the aggregate distribution does not imply equivalent competitive outcomes; two distributions can overlap substantially while one operator systematically wins per-run contests. In the uplink direction, distributions are generally tighter and the host–MVNO performance gaps are smaller across all families.

4.2 Per-Round Competitive Outcomes

Aggregate distributions describe the marginal performance of each operator. To capture relative outcomes under simultaneous measurements, we instead examine per-run win statistics. For each (location, round, scenario) tuple, an operator is credited with a *win* only if its mean throughput exceeds every rival's by more than 10 %; rounds in which no operator clears that bar are counted as *ties*. Table 1 summarizes the win rates across all families, directions, and scenario 2 to scenario 5. Three observations can be drawn from Table 1.

Host operator advantage is not universal. Across 18 competitive scenarios between the host operator and their virtual operators (3

Table 1: Per-round win rates (%) for Scenario 2 - Scenario 5 across all three operator families.

		H vs. V1			H vs. V2			V1 vs. V2			H vs. V1 vs. V2			
Family	**Dir.**	**H Win**	**V1 Win**	**Tie**	**H Win**	**V2 Win**	**Tie**	**V1 Win**	**V2 Win**	**Tie**	**H Win**	**V1 Win**	**V2 Win**	**Tie**
A	DL	43.1	29.4	27.5	90.0	10.0	0.0	100.0	0.0	0.0	80.0	10.0	0.0	10.0
	UL	21.6	58.8	19.6	30.0	70.0	0.0	80.0	10.0	10.0	10.0	60.0	0.0	30.0
B	DL	27.5	45.0	27.5	46.4	39.3	14.3	48.1	29.6	22.2	14.8	29.6	25.9	29.6
	UL	60.0	20.0	20.0	39.3	42.9	17.9	33.3	37.0	29.6	25.9	11.1	22.2	40.7
C	DL	50.0	42.9	7.1	76.9	15.4	7.7	13.3	40.0	46.7	63.6	9.1	18.2	9.1
	UL	34.3	51.4	14.3	54.3	22.9	22.9	48.6	37.1	14.3	25.7	40.0	17.1	17.1

families × 2 directions × 3 scenarios), the host operator has higher win rate in 9 scenarios, while an MVNO has the higher win rate in the remaining 9. For example, in Family A uplink, VA1 wins 58.8% of rounds against HA (with 19.6% ties), and in Family B downlink, VB1 records a higher win rate than HB (45.0% vs. 27.5%).

Competitive outcomes differ between downlink and uplink. For each family, the win-rate pattern changes between downlink and uplink. In Family A, HA has a higher downlink win rate than VA1 (43.1% vs. 29.4%) but a lower uplink win rate (21.6% vs. 58.8%). In Family B, VB2 trails HB in downlink (39.3% vs. 46.4%) but slightly exceeds HB in uplink (42.9% vs. 39.3%).

MVNOs within the same family can exhibit different competitive patterns. Within each family, the two MVNOs do not always show similar win rates against the host operator. In Family A downlink, VA1 wins 29.4% of rounds against HA, whereas VA2 wins 10.0%. In Family C uplink, VC1 records a 51.4% win rate against HC, while VC2 records 22.9%.

The following sections examine each family in depth, decomposing performance by location, scenario, and radio indicators to identify the cross-layer contexts associated with these patterns.

5 OPERATOR FAMILY A

Family A exhibits different competitive patterns in downlink and uplink. Across all Family A measurements, XCAL logs show a stable QoS class assignment pattern: HA and VA1 carry QoS class 8, while VA2 carries QoS class 9. In the downlink, HA and VA1 both outperform VA2 across the five competition scenarios: HA wins 90% of rounds against VA2, VA1 wins all VA1-vs-VA2 rounds, and HA leads VA1 in 43.1% of their head-to-head rounds (with 27.5% ties). In the uplink, the ordering inverts: HA wins only 21.6% of rounds against VA1 and 30.0% against VA2, while VA1 takes 58.8% against HA and 80% of VA1-vs-VA2 rounds. Far from conforming to a *host-always-better* expectation, an MVNO in Family A can consistently outperform its host carrier when the direction changes.

To assess whether this asymmetry is a coarse RAT-level artifact, we examine which radio generation each operator attaches to in every round. Across all Family A scenarios, RAT mismatches are confined to only two rounds, both within the HA-vs.-VA1 scenario, in which HA connected to LTE while VA1 connected to 5G NR; VA1 won both rounds in both downlink and uplink. While the outcome is unsurprising given the technology disparity, two rounds in 50 HA-vs.-VA1 rounds cannot account for the systematic direction-dependent pattern. We therefore conclude that RAT heterogeneity is not the primary driver of the observed asymmetry, and restrict all subsequent analysis to the remaining 48 rounds in which every participating operator is simultaneously on 5G NR. Fig. 2 shows

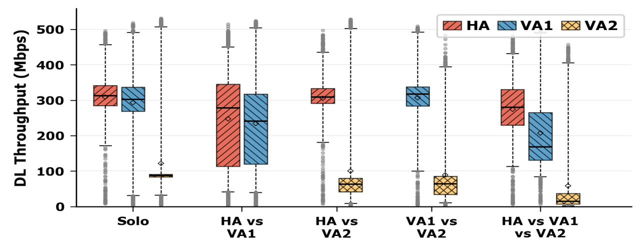

(a) Downlink.

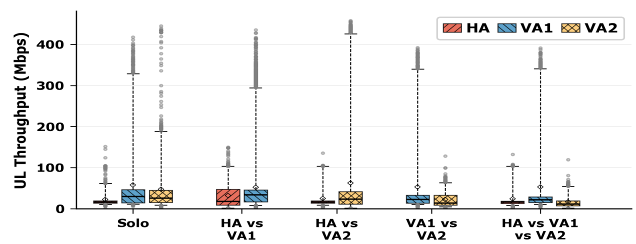

(b) Uplink.

Figure 2: Family A throughput distributions for rounds with 5G NR attachment across five scenarios.

Table 2: Family A per-run win rates (%) for rounds with 5G NR attachment.

	Downlink				*Uplink*			
Scenario	**HA**	**VA1**	**VA2**	**Tie**	**HA**	**VA1**	**VA2**	**Tie**
Solo	30.0	0.0	0.0	70.0	10.0	40.0	40.0	10.0
HA vs. VA1	43.8	27.1	—	29.1	18.4	61.2	—	20.4
HA vs. VA2	90.0	—	10.0	0.0	30.0	—	70.0	0.0
VA1 vs. VA2	—	100.0	0.0	0.0	—	80.0	10.0	10.0
HA vs. VA1 vs. VA2	80.0	10.0	0.0	10.0	10.0	60.0	0.0	30.0

the throughput distributions for rounds with 5G NR attachment; Table 2 summarizes the per-run win rate in both directions.

Downlink baseline. The downlink results under 5G NR attachment are similar to those obtained from the full set of measured rounds, indicating that excluding the two RAT-mismatch rounds does not materially change the aggregate pattern. HA and VA1 both consistently exceed VA2: HA wins 43.8% rounds against VA1 with 29.1% ties, 90% rounds against VA2, and 80% of simultaneous rounds, while VA1 wins all VA1 vs. VA2 rounds.

Uplink baseline. The uplink pattern differs substantially from the downlink pattern. HA wins 18.4% rounds against VA1, 30.0% rounds against VA2, and only 10% of simultaneous rounds. VA1 wins 80% of VA1-vs.-VA2 rounds. This downlink/uplink difference remains after removing the observed RAT mismatch and motivates a closer examination of within-5G factors in §5.1, §5.2.

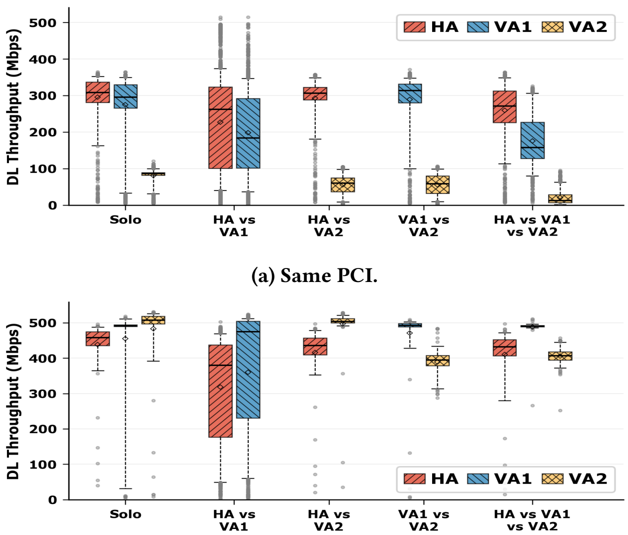

(a) Same PCI.

(b) Different PCI.

Figure 3: Family A downlink throughput for rounds with 5G NR attachment, conditioned on cell-attachment context.

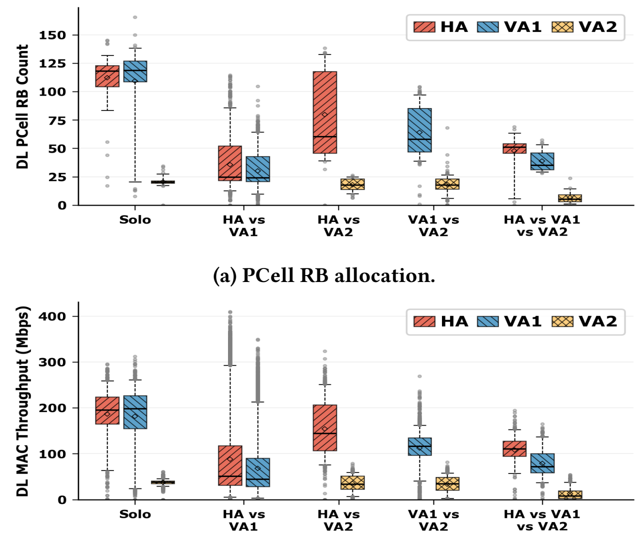

(a) PCell RB allocation.

(b) PCell MAC-layer throughput.

Figure 4: Family A downlink under matched radio context (same PCI, NR-ARFCN, bandwidth).

5.1 QoS class and Radio Context

At the aggregate level, the downlink ordering in Family A is characterized by similar performance between HA and VA1, with both exceeding VA2 in most scenarios. This ordering is consistent with the observed QoS class assignment pattern in Family A, where HA and VA1 carry QoS class 8 and VA2 carries QoS class 9. We therefore examine how this relationship varies with radio context, and, in particular, whether QoS class differences translate into RB allocation differences that in turn explain the observed throughput differences.

Same-PCI vs. Diff-PCI. Fig. 3 separates downlink measurements under 5G NR attachment according to whether competing operators share the same physical cell identifier (PCI). In rounds where competing operators share the same PCI (Fig. 3a), HA and VA1 exhibit higher medians than VA2 across the scenarios shown. In rounds where competing operators attach to different PCIs (Fig. 3b), the separation is smaller, and the ordering is less consistent across scenarios. These comparisons indicate that the association between QoS class and throughput ordering is stronger when competing operators share the same cell context, and weaker when they are attached to different cells.

Matched cell, carrier, and bandwidth. We further examine the case where different operators directly compete for resources from the same radio resource pool when their PCI, NR-ARFCN, and bandwidth are all identical. In real-world deployments, CA makes it relatively uncommon for two UEs to share an identical radio context across all serving cells. We therefore restrict this analysis to the Primary Cell (PCell), where the radio context can be explicitly matched across operators. Fig. 4 shows the distributions of PCell RB allocation and MAC-layer throughput across different competition scenarios for the downlink direction.

This matched-context view gives a mixed answer to whether QoS class explains RB allocation and throughput. When either HA or VA1 (QoS class 8) competes with VA2 (QoS class 9), both RB allocation and MAC-layer throughput consistently favor the QoS-class-8 operator, in both downlink and uplink; the three-operator scenario shows the same pattern, with VA2 receiving substantially fewer RBs and achieving lower throughput than HA and VA1. In these pairings, QoS class is a good predictor of both resource allocation and delivered throughput.

However, Fig. 4 also shows that HA obtains a modestly but consistently higher RB allocation and higher PCell throughput than VA1 across scenarios, even though the two operators carry the same observed QoS class. This asymmetry indicates that equal QoS class does not guarantee equal treatment at the scheduler: some additional, non-QCI factor is separating HA and VA1 even under matched serving conditions.

Illustrative case: competing under matched context. Table 3 reports one HA-vs.-VA1 downlink run (Boston Location 1, Round 2, Scenario 2) that illustrates this same-QCI asymmetry. Both UEs are served by the same physical cell and both carry QoS class 8. Their primary carriers are fully matched: the same PCI 371, the same NR-ARFCN 648672, and the same 100 MHz bandwidth, with closely comparable channel quality (RSRP within 2.7 dB and SINR within 1.7 dB) and link adaptation (MCS within 0.17). The run therefore isolates resource allocation from cell-, carrier-, and bandwidth-selection differences. On this matched primary carrier, HA is nevertheless granted a 55% larger RB allocation than VA1 (97.97 vs. 63.22 RBs on average, i.e., 35.9% vs. 23.2% RB utilization) and achieves 55% higher PCell MAC throughput (163.66 vs. 105.40 Mbps).

Normalizing throughput by allocated RBs, however, shows that HA and VA1 deliver essentially the same throughput per RB on this carrier (163.66/97.97 ≈ 1.67 Mbps/RB for HA vs. 105.40/63.22 ≈ 1.67 Mbps/RB for VA1), as also indicated by the same MCS. Since channel conditions were closely matched and per-RB efficiency is nearly identical, link adaptation and channel quality cannot account for the throughput gap; instead, the gap is attributable to the difference in the number of RBs granted, i.e., to a scheduler

Table 3: HA vs. VA1 downlink run at Boston Location 1, Round 2, Scenario 2, broken down by aggregated carrier. Values are per-run averages with standard deviation.

Operator	Carrier	PCI	NR-ARFCN	BW (MHz)	QoS class	RSRP (dBm)	SINR (dB)	RBs	RB Util. (%)	MCS	MAC TP (Mbps)
HA	PCell	371	648672	100	8	-104.42 ± 1.64	13.82 ± 2.03	97.97 ± 15.84	35.89 ± 5.80	14.39 ± 2.5	163.66 ± 41.53
	SCell[1]	371	176910	10	8	-92.96 ± 1.91	18.44 ± 1.82	27.62 ± 6.67	53.12 ± 12.82	18.13 ± 3.2	17.36 ± 7.62
	SCell[2]	371	653952	60	8	-110.62 ± 1.91	10.10 ± 1.17	72.63 ± 6.85	44.83 ± 4.23	13.99 ± 4.6	91.42 ± 23.58
	Aggregate	—	—	170	8	—	—	198.22	—	—	**272.44**
VA1	PCell	371	648672	100	8	-107.15 ± 1.70	12.11 ± 1.27	63.22 ± 13.49	23.16 ± 4.94	14.56 ± 2.3	105.40 ± 30.96
	SCell[1]	371	653952	60	8	-110.71 ± 1.57	11.03 ± 1.38	39.25 ± 7.84	14.38 ± 2.87	12.50 ± 2.7	54.89 ± 18.01
	Aggregate	—	—	160	8	—	—	102.47	—	—	**160.29**

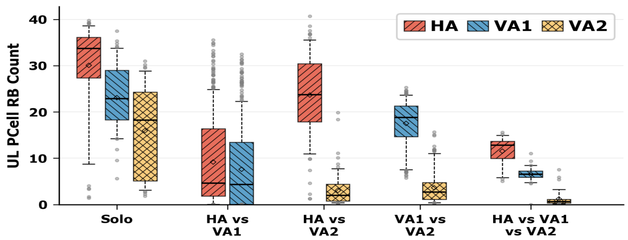

(a) PCell RB allocation.

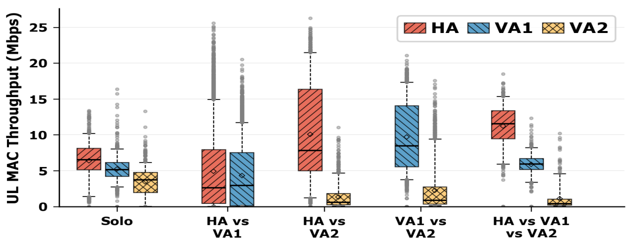

(b) PCell MAC-layer throughput.

Figure 5: Family A uplink under matched radio context (same PCI, NR-ARFCN, bandwidth).

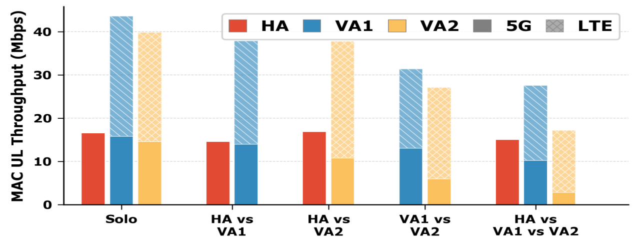

Figure 6: Family A MAC-layer uplink throughput at Boston Location 2. Solid fill: 5G NR component; hatched: LTE component.

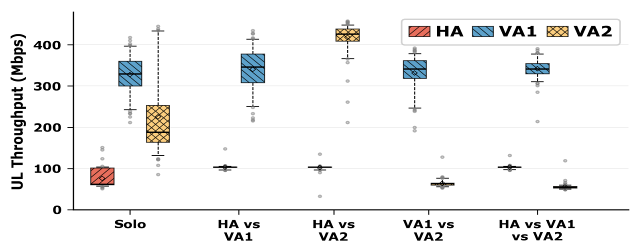

Figure 7: Family A application-layer uplink throughput at Atlanta Location 1. HA: FR1; VA1 and VA2: FR2.

allocation decision rather than to differences in spectral efficiency. Because our measurements are UE-side only, we cannot determine why HA is granted more RBs than VA1 despite the shared QCI value; candidate explanations include operator- or subscriber-level scheduling weights that are not exposed through the standardized QoS class (e.g., Allocation and Retention Priority [1] or other subscriber-profile parameters), or vendor-specific scheduler policies that differentiate host and MVNO traffic above the QCI abstraction. We report this as an observed association rather than evidence of a specific mechanism.

The gap further widens once carrier aggregation is taken into account. In this run, HA aggregates three carriers (100, 10, and 60 MHz) while VA1 aggregates two (100 and 60 MHz), so HA's aggregate MAC throughput reaches 272.44 Mbps against VA1's 160.29 Mbps, a 70% advantage rather than the 55% seen on the primary carrier alone. Differences in the aggregated carrier set can therefore amplify a gap that already exists on the shared carrier. We note that our UE-side traces record only which carriers each UE ends up aggregating, not the network policy or the UE capability signaling behind that choice; we therefore do not attempt to attribute the aggregation difference to a specific mechanism.

Overall, these results show that QoS class explains both RB allocation and throughput ordering when it differs between competing operators (HA/VA1 vs. VA2), but does not fully explain outcomes when QoS class is identical (HA vs. VA1): a residual RB-allocation asymmetry persists at the PCell level, and this residual can be further amplified once carrier aggregation is taken into account.

5.2 Architecture and Frequency Band

The uplink performance under 5G NR attachment (Fig. 2b) shows HA attains lower overall uplink throughput than VA1 and VA2 in multiple scenarios, despite having a comparable or even more advanced QoS class assignment. Inspection of the 5G architecture mode shows that HA is primarily attached in SA mode, whereas VA1 and VA2 operate predominantly in NSA mode. Under NSA operation, total uplink throughput includes both NR and LTE uplink components [3, 4], whereas SA operation contributes only the NR uplink component. Accordingly, comparisons of total uplink throughput between SA and NSA operators reflect both NR-side behavior and the presence or absence of LTE contribution.

Fig. 6 presents a representative example of uplink results for Family A, depicting the MAC-layer throughput breakdown at Boston Location 2. For VA1 and VA2, the LTE component accounts for a large share of total uplink throughput in most scenarios. This indicates that the application-layer uplink advantage of the MVNOs at this location is associated in large part with NSA dual connectivity rather than with NR-side scheduling behavior alone. When only the NR component is examined, the throughputs of HA and VA1 are closer than in the total uplink view, while VA2 throughput remains lower in several scenarios.

Table 4: HA vs. VA1 uplink run at Atlanta Location 5, Round 2, broken down by RAT. Values are per-run averages with standard deviation.

Op.	RAT	PCI	Freq.	BW (MHz)	QoS class	RSRP (dBm)	SINR (dB)	RBs	Util. (%)	MCS	TP (Mbps)
HA	NR	887	396030	10	8	−92.49 ± 1.88	17.82 ± 1.57	21.17 ± 3.26	40.71 ± 6.28	11.33 ± 1.15	5.03 ± 0.76
	LTE	284	132072	20	8	−101.18 ± 0.88	18.51 ± 1.80	65.60 ± 3.24	65.60 ± 3.24	10.62 ± 1.87	20.86 ± 4.35
	Agg.	—	—	30	8	—	—	—	—	—	**25.89**
VA1	NR	887	396030	10	8	−95.32 ± 1.09	9.86 ± 1.27	15.32 ± 2.52	29.46 ± 4.85	18.03 ± 1.38	9.09 ± 1.68
	LTE	284	132522	10	8	−97.06 ± 0.90	21.40 ± 0.79	39.39 ± 1.20	78.78 ± 2.40	20.97 ± 1.33	27.06 ± 1.88
	Agg.	—	—	20	8	—	—	—	—	—	**36.15**

Frequency band selection introduces an additional source of variation at Atlanta Location 1. Fig. 7 shows the application-layer uplink throughput at that location, where VA1 and VA2 attach to millimeter-wave (mmWave) (FR2) while HA attaches to midband (FR1) [4]. In the host-vs.-MVNO scenarios (scenario 2, scenario 3, and scenario 5), both MVNOs achieve substantially higher uplink throughput than HA. This difference is consistent with the distinct frequency band selection observed at this location.

When the two MVNOs are compared under a more similar mmWave context (VA1 vs. VA2), VA1 records the higher throughput across rounds. Under this more comparable radio context, the throughput ordering is consistent with the corresponding operator ordering observed elsewhere in Family A. These results indicate that aggregate uplink ordering in Family A is associated with both 5G architecture mode and frequency band selection.

Matched cell, carrier, and scheduling context The preceding analysis relates aggregate uplink differences in Family A to architecture mode (SA vs. NSA) and frequency band selection. A separate question, parallel to §5.1, is whether QoS class continues to explain RB allocation and throughput once cell, carrier, and bandwidth are held fixed, independent of architecture mode, which can still vary across operators in this comparison (HA is predominantly SA, VA1 predominantly NSA). As in the downlink, we examine PCell RB allocation and MAC-layer throughput under matched PCI, NR-ARFCN, and bandwidth. Fig. 5 shows that the same overall pattern observed for downlink holds for uplink: QoS class is a good predictor of RB allocation and throughput when it differs between the competing operators, but not when it is the same.

To examine this same-QoS-class case at finer granularity, and with architecture mode also controlled, we isolate one uplink run in which HA and VA1 are attached to the same physical cell (matched PCI), carry the same QoS class value, and both happen to operate in NSA mode, shown in Table 4. Their NR carriers are fully matched (PCI 887, NR-ARFCN 396030, 10 MHz), so the NR rows isolate scheduling and link adaptation from any cell- or spectrum-selection difference; their LTE carriers, in contrast, differ in both frequency and bandwidth.

On the matched NR carrier, the outcome is the reverse of the downlink example in Table 3, even though RB allocation follows the same pattern. HA is again granted the larger share (21.17 vs. 15.32 RBs, 40.71% vs. 29.46% utilization) and reports the better channel (RSRP 2.8 dB higher, SINR 8.0 dB higher). Yet, its NR throughput is lower – 5.03 Mbps against VA1's 9.09 Mbps. The gap is carried entirely by MCS: VA1 is scheduled at 18.03 against HA's 11.33, 59% higher, despite the worse reported channel. The LTE carriers show a different split: HA holds twice the bandwidth (20 vs. 10 MHz) but a lower SINR (18.51 vs. 21.40 dB), and VA1 again has higher MCS (20.97 vs. 10.62) and achieves higher throughput (27.06 vs. 20.86 Mbps) from fewer PRBs. Aggregated across both RATs, VA1 reaches 36.15 Mbps against HA's 25.89 Mbps, a 40% advantage. Because HA holds the larger resource share on both carriers yet still trails in throughput, and because the MCS gap appears both where HA's channel is worse (LTE) and where it is better (NR), this is unlikely to be a simple channel-quality artifact. Our UE-side traces expose only the resulting grants, not the scheduler or power-control logic behind them, so we do not attribute it to a specific mechanism. One plausible contributor is that RSRP and SINR, as reported by the UE, characterize the downlink channel, while MCS on the uplink is selected by the gNB based on the uplink channel it observes (e.g., via SRS), together with uplink-side interference from other UEs that neither phone can observe. Downlink channel symmetry does not imply uplink channel symmetry, and the gNB's uplink scheduling decision need not track the downlink SINR reported by XCAL.

Accordingly, per-run throughput outcomes are not determined by a single factor. Even after accounting for RAT, architecture mode, frequency band selection, cell attachment, and RB allocation, link adaptation may still contribute to observed throughput differences in individual rounds.

5.3 City-Level Differences

The same three subscriptions behave very differently in the two cities, both in absolute rate and in the ordering among the operators (Figs. 8, 9). In the downlink, pooling over all scenarios, the Atlanta medians are 262.5, 276.0, and 446.0 Mbps for HA, VA1, and VA2, against 296.0, 261.0, and 58.2 Mbps in Boston. HA and VA1 are broadly comparable across the two cities, but VA2 moves from a distant third in Boston to clearly catching up with VA1 and even surpassing HA in Atlanta, a shift of nearly an order of magnitude. The uplink is faster in Atlanta for all three operators (48.5, 35.9, and 113.5 Mbps against 14.6, 23.6, and 14.8), and the ordering changes again: in Boston VA1 leads its own host, whereas in Atlanta the host has a higher median, while VA1 exhibits a long upper tail. Meanwhile, VA1 clearly outperforms in Atlanta, both in the direct VA1–VA2 comparison and in the three-operator scenario, whereas in Boston their performance is very close. VA2 also significantly outperforms the host in Atlanta, whereas in Boston it is only slightly higher. The Atlanta sites give the MVNOs access to spectrum not available in Boston (the mmWave carriers discussed in §5.2), which accounts for much of both the level shift and the reordering.

Summary. Overall, Family A shows a direction-dependent competitive pattern. Downlink results broadly follow the QoS class assignment pattern, especially under shared-cell attachment. In

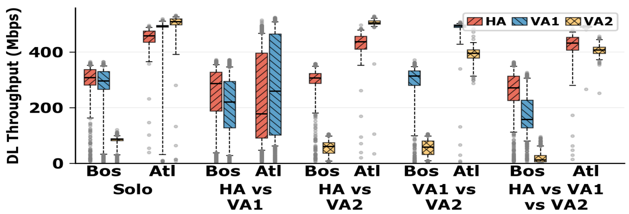

Figure 8: Family A downlink throughput distributions by city (all valid data).

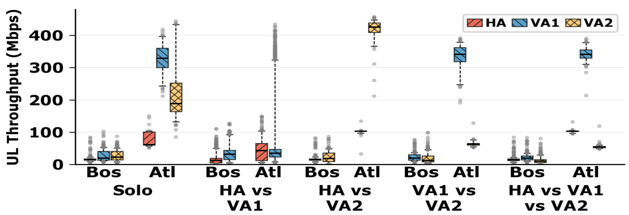

Figure 9: Family A uplink throughput distributions by city (all valid data).

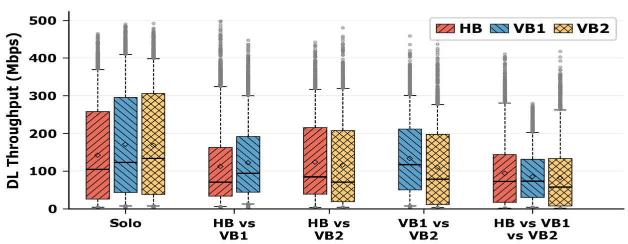

(a) Downlink.

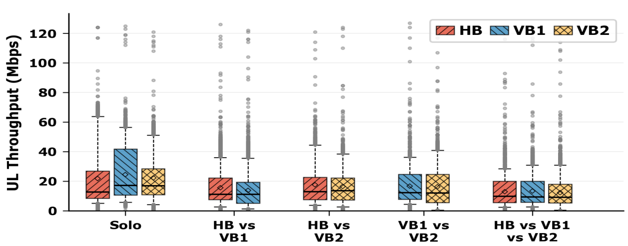

(b) Uplink.

Figure 10: Family B throughput distributions across five scenarios and all measured rounds.

uplink, however, VA1 and VA2 often outperform HA due to NSA LTE contribution and favorable frequency band selection.

6 OPERATOR FAMILY B

Family B differs from Family A in two aspects. First, all three operators consistently employ QCI 8, i.e., the host MNO does not use a more favorable QCI value. Second, all three operate in Non-Standalone (NSA) mode.

Fig. 10 shows the throughput distributions aggregated over all measured rounds, and the Family B rows of Table 1 report the corresponding per-run win rates. We observe that Family B shows substantial overlap among operators and no consistent host advantage. In the downlink, the leading operator varies by pairing: VB1 leads HB in the HB-vs.-VB1 scenario, HB leads VB2 in the HB-vs.-VB2 scenario, and VB1 leads VB2 in the MVNO-only scenario. In the three-operator simultaneous scenario, the win rates are split across all three operators with a large tie fraction. The uplink shows a similarly mixed pattern. HB leads VB1 in their head-to-head comparison, but VB2 slightly exceeds HB in the HB-vs.-VB2 scenario, and no operator dominates the three-operator scenario.

Since QCI and architecture mode are constant across the three operators, we therefore examine two remaining sources of variation: RAT attachment and cell-attachment context.

Table 5: Family B RAT mismatch statistics.

Operator / Metric	Downlink	Uplink
HB on LTE	29/150 (19.3%)	11/150 (7.3%)
VB1 on LTE	12/150 (8.0%)	5/150 (3.3%)
VB2 on LTE	21/150 (14.0%)	18/150 (12.0%)
RAT-mismatch rounds	51/150 (34.0%)	29/150 (19.3%)
Overall mismatch	80/300 (26.7%)	

Table 6: Family B per-run win rates (%) restricted to rounds in which every participating operator is attached to 5G NR.

	Downlink				*Uplink*			
Scenario	**HB**	**VB1**	**VB2**	**Tie**	**HB**	**VB1**	**VB2**	**Tie**
Solo	15.8	21.1	21.1	42.1	24.0	16.0	24.0	36.0
HB vs. VB1	29.6	33.3	—	37.0	59.5	18.9	—	21.6
HB vs. VB2	52.9	—	23.5	23.5	35.0	—	45.0	20.0
VB1 vs. VB2	—	36.8	36.8	26.3	—	10.0	50.0	40.0
HB vs. VB1 vs. VB2	17.6	11.8	23.5	47.1	15.8	10.5	31.6	42.1

6.1 RAT Mismatch

RAT mismatch is much more common in Family B than in Family A. Among 300 valid rounds, 80 (26.7%) include at least one operator attached to LTE rather than 5G NR. The mismatch rate is higher in downlink (51/150, 34.0%) than in uplink (29/150, 19.3%). Table 5 further shows how often each operator is attached to LTE out of all valid rounds.

Within mismatch rounds, win outcomes largely follow RAT attachment. The operator attached to 5G NR wins in 79 of the 80 mismatch rounds. The only exception is one three-operator round at Atlanta Location 2, where HB on LTE slightly exceeds the two 5G-attached operators at unusually low throughput levels.

Because RAT mismatch dominates the outcome whenever it occurs, the remainder of the analysis focuses on the subset of rounds in which all participating operators are attached to 5G NR. Fig. 11 shows the resulting throughput distributions, and Table 6 recomputes the per-run win rates over this subset.

Removing the RAT-mismatch rounds changes the picture partially; the tie fraction rises in most pairings, reaching 37.0% in HB-vs.-VB1 downlink and above 40% in the three-operator scenario in both directions, confirming that a large part of the apparent separation in Table 1 came from LTE-vs.-NR competition rather than from competition within 5G NR. However, the ordering among operators remains the same: HB still leads VB2 in the downlink (52.9% vs. 23.5%) but trails it in the uplink (35.0% vs. 45.0%), and neither MVNO holds a stable advantage over the other. With QoS class and architecture mode fixed across the three operators, this leaves cell-attachment context as the remaining source of variation, which we examine next.

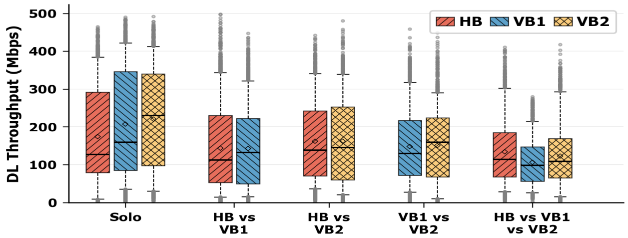

(a) Downlink.

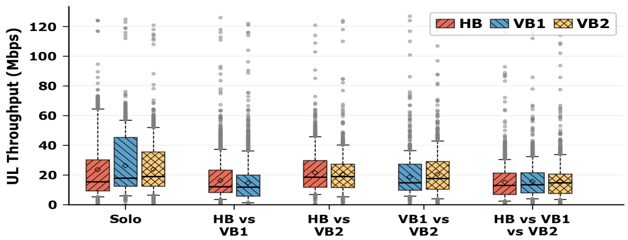

(b) Uplink.

Figure 11: Family B throughput distributions for rounds with 5G NR attachment.

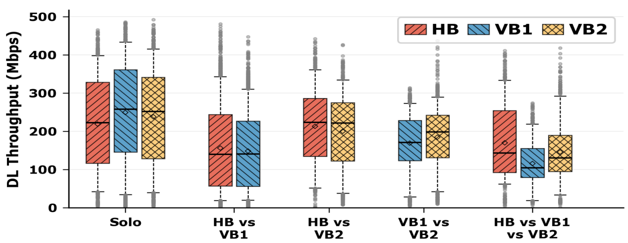

(a) Same PCI.

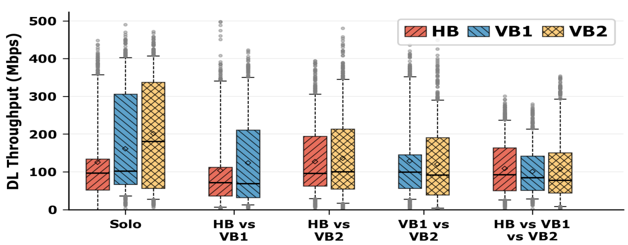

(b) Different PCI.

Figure 12: Family B downlink throughput for rounds with 5G NR attachment, conditioned on cell-attachment context.

6.2 Cell-Attachment Under 5G NR

After restricting the analysis to rounds with 5G NR attachment, QCI and architecture mode remain constant across Family B operators: all three carry QCI 8 and operate in NSA mode. This makes Family B useful for examining variation across cell-attachment contexts.

Figs. 12, 13 compare throughput across Same-PCI and Diff-PCI contexts in Family B, which is asymmetric between the host and MVNOs.

Same-PCI attachment. The three operators share the same physical cell, and their throughput distributions overlap substantially in both downlink and uplink (Figs. 12a, 13a). This is consistent

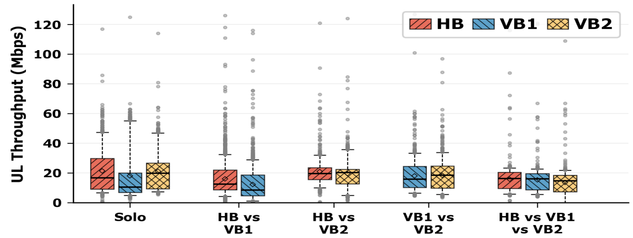

(a) Same PCI.

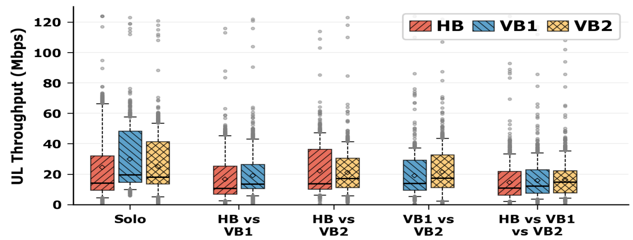

(b) Different PCI.

Figure 13: Family B uplink throughput for rounds with 5G NR attachment, conditioned on cell-attachment context.

with Family B having the same QCI assignment and the same NSA architecture mode across all three operators.

However, when we further tighten the Same-PCI condition as in §5.1 and look at PCell RB allocation and throughput under matched PCI, NR-ARFCN, and bandwidth, we reach a similar conclusion: although all three operators share QCI 8, RB allocation and throughput can still differ under fully matched serving conditions (Figs. 23, 24 in Appendix E). In the downlink competing scenarios, allocation is close across operators, consistent with the shared QCI. In the uplink, however, HB is granted more RBs than VB1 in their head-to-head round but fewer RBs than VB2 in its own head-to-head round (18.2 vs. 14.4, and 10.9 vs. 16.0, respectively) – an inconsistent, direction-dependent split among operators that share the same QCI, cell, carrier, and bandwidth. As in Family A, this residual asymmetry is not attributable to QoS class, and our UE-side measurements cannot determine its cause.

Diff-PCI attachment. Unlike under Same-PCI attachment, here (Figs. 12b, 13b), in the downlink HB trails both VB1 and VB2, and its clear lead in the three-operator scenario shrinks substantially. The same is true for the uplink direction, although the gaps are smaller compared to the downlink case. These Diff-PCI cases expose HB and the MVNOs to different serving configurations, including differences in band, bandwidth, and available radio resources.

This is illustrated in Table 7. Both UEs operate in NSA mode and have QCI 8. The table lists every carrier they hold, separated by RAT. The two RATs behave very differently. On the LTE anchor the two operators are indistinguishable: they are served by the same cell (PCI 437, Band 2, 20 MHz) and receive nearly the same share of it, 44.68 vs. 39.51 RBs, yielding almost identical throughput (25.84 vs. 24.96 Mbps). All of the difference between the two subscriptions therefore originates on the NR side, where HB is attached to a 10 MHz n5 carrier and VB1 to an 80 MHz n77 carrier. What the NR side shows is that the frequency and bandwidth of the carrier a UE lands on set the achievable rate. HB is in fact the more heavily served of the two in relative terms: it occupies 73.97% of its NR carrier

Table 7: HB vs. VB1 downlink run at Boston Location 2, Round 1, broken down by RAT. Both operators carry QCI 8. Values are per-run averages with standard deviation.

Op.	RAT	PCI	Band	BW (MHz)	RBs	Util. (%)	TP (Mbps)
HB	NR	271	n5	10	38.47 ± 10.08	73.97 ± 19.38	51.29 ± 17.26
	LTE	437	Band 2	20	44.68 ± 15.35	44.68 ± 15.35	25.84 ± 9.08
	Agg.	—	—	30	—	—	77.12
VB1	NR	424	n77	80	123.61 ± 15.33	56.96 ± 7.07	308.72 ± 90.02
	LTE	437	Band 2	20	39.51 ± 10.89	39.51 ± 10.89	24.96 ± 6.80
	Agg.	—	—	100	—	—	333.67

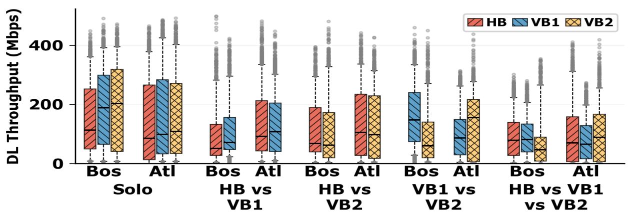

Figure 14: Family B downlink throughput distributions by city.

against VB1's 56.96%. Yet because its carrier is eight times narrower and on a low band, those 38.47 RBs deliver only 51.29 Mbps, while VB1's 123.61 RBs on the wide mid-band carrier deliver 308.72 Mbps, six times more. Aggregated over both RATs the gap is 77.12 vs. 333.67 Mbps, a factor of 4.3. This run therefore reveals that carrier selection, band, and bandwidth act on the outcome independently of how the scheduler divides an individual carrier.

Overall, the Same-PCI/Diff-PCI split shows that Family B's host–MVNO ordering is not determined by QCI or architecture mode, which are fixed across operators. Instead, the ordering changes with cell attachment: when HB and the MVNOs are compared across different cells, differences in band, bandwidth, and available radio resources can reduce or reverse the host advantage.

6.3 City-Level Comparison

Figs. 14, 15 compare Family B throughput distributions between Boston and Atlanta using all valid rounds. These plots include both RAT-matched and RAT-mismatched cases, and therefore capture the combined effect of local RAT availability, serving-cell attachment, band selection, bandwidth, and radio conditions.

In the downlink, both cities show substantial overlap among HB, VB1, and VB2, but the leading operator changes by scenario. Boston shows stronger VB1/VB2 performance in the solo case and stronger VB1 performance in the MVNO-only scenario, whereas Atlanta is more balanced in the solo scenarios and shows stronger VB2 performance in the MVNO-only and three-operator scenarios. This also reveals a location-dependent reversal between the two MVNOs: VB1 leads VB2 in Boston, while VB2 becomes stronger than VB1 in Atlanta. In the uplink, the performance gaps among the three operators become smaller in Atlanta than those observed in Boston, resulting in a more balanced performance across operators but with longer upper tails.

Thus, the city-level results reinforce the Family B finding that host–MVNO ordering is location-dependent.

Summary. Family B shows that even when QCI and architecture mode are fixed across operators, host–MVNO ordering remains variable. RAT mismatch explains many outcomes: when LTE and 5G NR

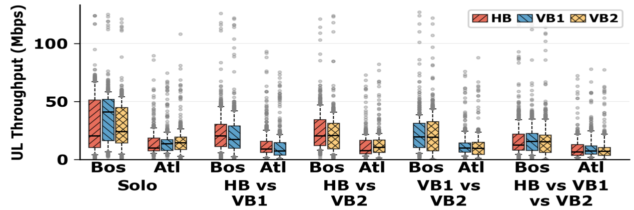

Figure 15: Family B uplink throughput distributions by city.

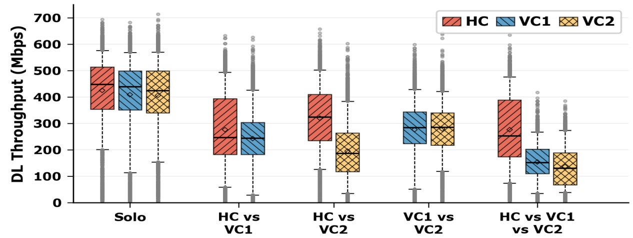

(a) Downlink.

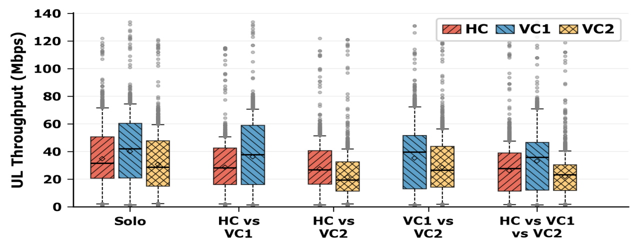

(b) Uplink.

Figure 16: Family C throughput distributions for rounds with 5G NR attachment.

compete, the 5G-attached operator almost always wins. Within 5G NR, the ordering changes with PCI context and city, indicating that serving-cell selection, bandwidth, and local radio conditions can reshape the comparison even without QCI or SA/NSA differences.

7 OPERATOR FAMILY C

Family C maintains a consistent QoS class configuration in every valid round: HC is assigned QoS class 6, and both VC1 and VC2 are assigned QoS class 7. Unlike Family B, we observe no RAT mismatch in either direction; all participating operators are attached to 5G NR in the analyzed rounds.

Fig. 16 shows the aggregate throughput distributions. In the downlink, HC generally has the strongest performance in host-vs.-MVNO scenarios, especially against VC2 and in the three-operator scenario. However, the HC-vs.-VC1 case is more balanced, with VC1 winning a substantial fraction of rounds. In the uplink, the ordering changes: VC1 records the highest win rate in every related scenario. Thus, Family C shows a clear direction asymmetry: the downlink is broadly consistent with HC's more favorable QoS class, whereas the uplink often favors VC1 despite its higher QoS class value.

7.1 QoS class and Radio Context

Fig. 17 separates Family C downlink throughput by cell-attachment context.

Under Same-PCI attachment. HC's advantage is clear in the host-vs.-MVNO scenarios and in the three-operator scenario. This pattern is consistent with HC carrying QoS class 6 while VC1 and VC2 carry QoS class 7. The MVNO-only case is balanced, as both VC1 and VC2 carry the same QoS class.

As in §5.1, §6.2, we further look at PCell RB allocation and throughput under matched PCI, NR-ARFCN, and bandwidth; the Solo and three-operator downlink scenarios retain no round under this restriction. Figs. 25, 26 in Appendix E show the resulting PCell RB allocation and MAC-layer throughput. In the downlink, allocation follows the QoS-class ordering closely: HC receives roughly twice the RBs and throughput of VC1 or VC2 in both host-vs.-MVNO pairings, while VC1 and VC2, which share QoS class 7, receive similar shares with heavily overlapping distributions. The uplink shows the same class-based ordering in the head-to-head pairings, but the MVNO-only pairing breaks it: VC2 receives more than twice VC1's median allocation despite sharing the same QoS class, with VC1's allocation also varying considerably rather than settling at a stable level. Family C therefore reproduces the same pattern observed in Families A and B: QoS class differences map onto RB-allocation and throughput differences, but equal QoS class does not guarantee equal allocation or throughput. As before, this within-class residual is not attributable from UE-side measurements alone.

Under Diff-PCI attachment. The QoS class-aligned ordering becomes less stable. In the HC-vs.-VC1 scenario, VC1 leads HC even though HC carries the higher-priority QoS class. In HC-vs.-VC2, HC remains comparable to or stronger than VC2, but the separation is smaller than in the Same-PCI case. In the three-operator scenario, HC and VC1 are closer, while VC2 remains lower. These results show that the downlink association between QoS class and throughput is strongest under shared-cell attachment and weakens when operators attach to different cells.

This is illustrated in Table 8. The two subscriptions are not short of spectrum relative to each other: HC aggregates four NR carriers totaling 200 MHz, while VC1 aggregates two NR and two LTE carriers totaling 195 MHz. Their two principal NR carriers are configured almost identically – both hold an 80 MHz and a 90 MHz n41 carrier, with comparable RB counts and utilization (51–55% in all four cases) – but the two operators land on entirely different cells. This difference in serving cell shows up directly in the channel: on the matched 80 MHz carrier, VC1 records 15.74 dB SINR against HC's 8.32 dB, and on the matched 90 MHz carrier, 11.50 dB against 1.80 dB. MCS follows the same pattern (8.77 vs. 4.98 and 6.53 vs. 4.98), so equal resources on equal spectrum convert into very unequal rates: 129.91 and 106.13 Mbps for VC1 against 36.29 and 43.30 Mbps for HC. HC's two remaining carriers make the same point more starkly — both are n25 cells with negative SINR, and the 20 MHz one returns only 2.82 Mbps despite HC occupying 62% of it. Aggregated across all carriers, VC1 reaches 298.10 Mbps against HC's 114.78 Mbps, a factor of 2.6, even though HC holds slightly more spectrum and the more favorable QoS class. Hence, serving-cell selection, rather than QoS class or the quantity of resources granted, determines the outcome of this run.

7.2 Architecture Mode

The uplink results do not follow the QoS class ordering. As shown in Fig. 16b, VC1 often achieves the highest uplink throughput, even though HC carries QoS class 6 and VC1 carries QoS class 7.

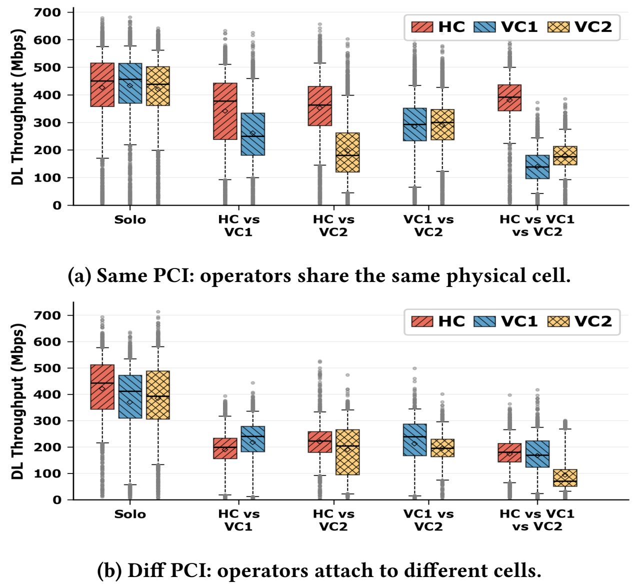

Figure 17: Family C downlink throughput, conditioned on cell-attachment context.

Fig. 18 shows that this VC1 advantage appears under both Same-PCI and Diff-PCI attachment. Under Same-PCI attachment, VC1 is stronger than HC in the HC-vs.-VC1 scenario and remains competitive in the three-operator scenario. Under Diff-PCI attachment, VC1 shows even wider distributions and often higher medians, especially in HC-vs.-VC1 and VC1-vs.-VC2. Thus, unlike the downlink, the uplink advantage of VC1 is not primarily explained by QoS class or PCI context.

Fig. 19 explains this asymmetry. Across every location and round, VC1 includes both NR and LTE uplink components, while HC and VC2 include only NR. This pattern is consistent with VC1 operating in NSA mode and HC/VC2 operating in SA mode. VC1's NR uplink component is comparable to or lower than HC's NR component, but the additional LTE component raises its combined uplink throughput above HC in many scenarios. Thus, the main source of Family C uplink asymmetry is architecture mode rather than QoS class.

7.3 City-Level Comparison

Figs. 20, 21 compare Family C throughput distributions across Boston and Philadelphia under the 677-type QoS class pattern. In the downlink, both cities show evidence of HC advantage in host-vs.-MVNO comparisons, but the strength and ordering vary by city. In Boston, HC leads clearly in HC-vs.-VC2 and in the three-operator scenario, while VC1 and VC2 are closer in the MVNO-only comparison. In Philadelphia, the overall throughput range is lower and the distributions are more variable; HC remains strong in host-vs.-MVNO scenarios, but VC2 becomes stronger than VC1 in the MVNO-only and three-operator scenarios. This city-level difference shows that the 677-type QoS class pattern does not produce a fixed ordering across locations. The host advantage in downlink is visible in both cities, but local access context changes the relative ordering between the two MVNOs and the magnitude of the gap. In the uplink, the city-level comparison follows the aggregate observation

Table 8: HC vs. VC1 downlink run at Philadelphia Location 1, Round 2, broken down by aggregated carrier. Values are per-run averages with standard deviation.

Op.	Carrier	PCI	Band	BW (MHz)	QoS class	RSRP (dBm)	SINR (dB)	RBs	Util. (%)	MCS	TP (Mbps)
HC	NR PCell	171	n41	80	6	−104.29 ± 0.64	8.32 ± 0.99	111.08 ± 22.53	51.19 ± 10.38	4.98 ± 2.40	36.29 ± 13.07
	NR SCell[1]	522	n41	90	6	−110.98 ± 1.24	1.80 ± 0.84	128.02 ± 21.27	52.25 ± 8.68	4.98 ± 2.92	43.30 ± 17.48
	NR SCell[2]	640	n25	20	6	−104.94 ± 2.40	−9.53 ± 1.83	66.14 ± 35.85	62.40 ± 33.82	4.48 ± 3.05	2.82 ± 7.41
	NR SCell[3]	640	n25	10	6	−98.92 ± 0.98	−5.15 ± 1.21	42.60 ± 8.65	81.93 ± 16.63	11.34 ± 3.65	32.37 ± 11.41
	Agg.	—	—	200	6	—	—	—	—	—	**114.78**
VC1	NR PCell	629	n41	80	7	−104.30 ± 0.67	15.74 ± 1.27	118.34 ± 17.03	54.53 ± 7.85	8.77 ± 3.05	129.91 ± 39.35
	NR SCell[1]	67	n41	90	7	−109.38 ± 1.08	11.50 ± 0.87	126.58 ± 14.50	51.67 ± 5.92	6.53 ± 3.31	106.13 ± 29.08
	LTE PCell	333	Band 66	20	7	−99.61 ± 2.84	12.65 ± 1.21	67.99 ± 20.32	67.99 ± 20.32	14.55 ± 2.89	50.55 ± 15.83
	LTE SCell[1]	18	Band 2	5	7	−97.54 ± 1.68	9.47 ± 1.62	13.82 ± 3.59	55.30 ± 14.36	14.29 ± 4.17	11.51 ± 3.38
	Agg.	—	—	195	7	—	—	—	—	—	**298.10**

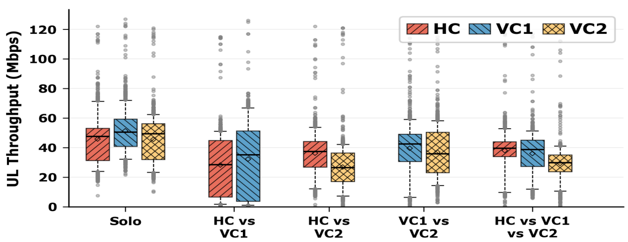

(a) Same PCI.

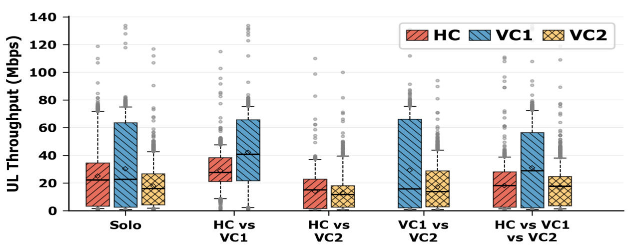

(b) Different PCI.

Figure 18: Family C uplink throughput, conditioned on cell-attachment context.

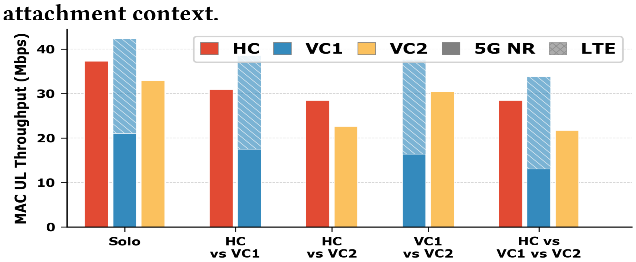

Figure 19: Family C MAC-layer uplink throughput by scenario (Across every location and round). Solid fill: 5G NR; hatched fill: LTE.

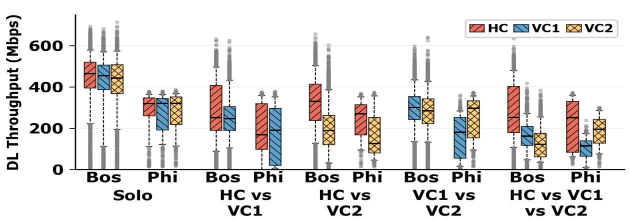

Figure 20: Family C downlink distributions by city.

that VC1 is frequently competitive or stronger, consistent with the NSA aggregation behavior discussed above.

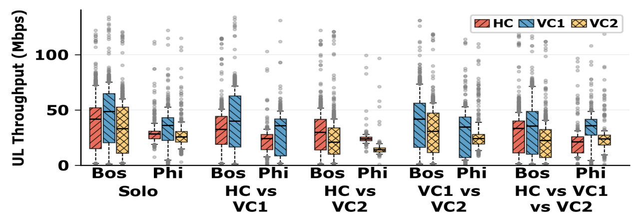

Figure 21: Family C uplink distributions by city.

Summary. Family C shows a stable 677-type QoS class pattern with no RAT mismatch. HC often leads in downlink, especially under Same-PCI attachment, but Diff-PCI cases and NSA aggregation create exceptions. In uplink, VC1 often outperforms HC because its NSA operation adds an LTE component to NR uplink throughput. Thus, Family C outcomes depend largely on PCI context and SA/NSA architecture.

8 DISCUSSION

QoS class explains less than it promises, even in the best case. Across all three families, QoS class predicts throughput ordering most reliably when the competing operators also share a cell, carrier, and bandwidth: different QoS classes then map cleanly onto different RB shares (Family A HA/VA1 vs. VA2; Family C HC vs. VC1/VC2). But even this best case is incomplete in two ways. First, equal QoS class does not imply equal treatment: HA and VA1 (Family A), HB and the two MVNOs (Family B), and VC1/VC2 in one direction (Family C) all show residual RB-allocation or throughput asymmetries under fully matched serving conditions that QoS class cannot explain. Second, even where RB allocation is symmetric, throughput need not be: our matched-cell uplink example (Table 4) shows a case where the operator with more RBs and an equal-or-better reported downlink channel is nonetheless scheduled at a much lower MCS and loses on throughput. QoS class, resource allocation, and link adaptation are therefore three distinct and only loosely coupled layers, and no single one is a reliable proxy for the others.

Access differentiation dominates when contexts differ. When competing operators are not confined to the same cell, carrier, and bandwidth, the performance gap is usually decided before scheduling ever comes into play. RAT mismatch, serving-cell selection, frequency band, bandwidth, and architecture mode can each dominate the outcome on their own, and city-level comparisons show these effects compounding rather than canceling: the same subscriptions can invert their relative ranking, and shift by nearly an order of magnitude, simply because a different location exposes the

MVNOs to different spectrum. These are differences in what network the UE is actually connected to and should not be interpreted as evidence of scheduling priority.

Implications. For users, our results show that neither MNO status (host vs. MVNO) nor a favorable QoS class guarantees better performance: a higher-priority host can still lose to its own MVNO once RAT, cell attachment, architecture mode, or location diverge, so online discussions that treat QoS class alone as a proxy for subscriber priority are often missing many factors that actually determine throughput. For operators, our results point to concrete risks in how access is provisioned rather than just scheduled, with actionable fixes. Keeping the host on SA while leaving MVNOs on NSA can backfire in the uplink, where the MVNO's LTE anchor supplies a throughput component the SA host lacks entirely (Family A, Family C); a host optimized for downlink priority can still lose the uplink comparison for a reason unrelated to QoS class. Likewise, if host and MVNO traffic are steered toward different serving cells, cell selection rather than the QoS class configured for either subscriber determines the outcome whenever one operator lands on a narrower, lower band, or worse-SINR cell (Family B, Family C); a host operator has no throughput advantage to show for its QoS priority if its own subscribers are placed on inferior cells. For the measurement community, our results argue that application-layer throughput comparisons of host and MVNO subscribers are uninterpretable without simultaneous, cross-layer radio-side visibility; the same aggregate gap can arise from RAT mismatch, cell attachment, spectrum differences, RB allocation, or link adaptation, each requiring different data to detect and each carrying a different implication for user experience.

Scope and limitations. Our results come from nine commercial subscriptions across three U.S. cities and a limited set of stationary sites, and are not nationally representative. Measurements spanned multiple days and times, but were not a controlled busy-hour experiment, so temporal and background-load effects cannot be separated from the reported differences. UE-side RB-utilization figures indicate observed usage, not cell-wide load or congestion. The study covers only sustained non-GBR TCP traffic; other workloads (interactive, real-time, GBR) may be treated differently. We also do not measure latency or reconstruct scheduler logic directly; our traces expose resulting grants, not the decisions or policies behind them, which is why we report cross-layer associations rather than causal mechanisms throughout the paper. Closing these gaps is a natural direction for future work.

9 RELATED WORK

5G measurement studies. A large body of work has measured commercial 5G networks from different perspectives. Early studies characterize operational 5G coverage, throughput, latency, power consumption, and application Quality of Experience (QoE) [36, 38, 39, 44, 45, 51]. A second line of work isolates individual mechanisms and deployment choices: SA versus NSA operation and the evolution of NSA deployments [12, 13, 23, 33–35, 53]; spectrum band, bandwidth, carrier aggregation, and MIMO [19, 21, 54]; scheduling, beam management, and PHY latency [16, 18, 20, 22]; mobility, handover, and roaming [8, 24, 28, 30, 31]; and throughput prediction [9, 37]. A third line scales measurements up in space or time, through nationwide or operator-side datasets, longitudinal campaigns, and driving- or transit-based collection [10, 14, 25–27, 29, 32, 41, 49, 52].

These studies show that 5G performance depends on deployment choices such as spectrum band, SA/NSA mode, carrier aggregation, bandwidth, mobility, and local radio conditions. However, they primarily characterize MNO performance, 5G features, mobility, or application QoE. Our work studies a different question: how MVNOs compete with their host MNOs under co-located and synchronized measurements.

MVNO performance and mobile service architecture. Prior work studies MVNOs from application, architecture, and service-ecosystem perspectives. Zarinni et al. [55] provide the closest measurement study to our work, comparing base carriers and MVNOs across web, video, and voice applications and showing that MVNOs sharing the same base carrier can still expose different QoE. Their study predates commercial 5G, however, and does not examine radio-side factors. Our work focuses on TCP throughput alone but uses XCAL traces for cross-layer analysis in the 5G context; consistent with their finding, we show that sharing a host network does not imply identical user-perceived performance. Unlike their application-layer comparison, however, our simultaneous, cross-layer measurements further show that the direction and magnitude of host–MVNO differences vary across operator families and can coincide with RAT attachment, serving-cell configuration, spectrum bandwidth, resource allocation, and SA/NSA operation. Xiao et al. [50] study a large commercial MVNO from architecture, billing, data-usage prediction, reselling, and churn perspectives, while Vallina-Rodriguez et al. [47] and Alcalá-Marín et al. [6] examine higher-layer mobile-network configuration, business relationships, aggregators, and roaming behavior. These works show that virtualized mobile services can differ from traditional MNO service, but they do not study matched 5G host–MVNO competition with radio-side visibility. Our work fills this gap by relating end-to-end throughput differences to the observed cross-layer access context.

10 CONCLUSION

In this paper, we presented a cross-layer measurement study of host-MNO and MVNO performance across three operator families, five competition scenarios, and three U.S. cities. Our results show that host-MNO advantage is conditional rather than universal, and that the factors associated with performance differences vary substantially across operator deployments. Family A exhibits strong downlink–uplink asymmetry associated with architecture and frequency-band differences; Family B is strongly affected by RAT mismatch and serving-cell heterogeneity; and Family C shows a distinct uplink advantage for one MVNO associated with its additional LTE contribution under NSA. Matched serving-context measurements further show that even UEs sharing the same NR cell and carrier can receive different radio-resource allocations and achieve different throughput. Overall, host–MVNO performance cannot be characterized by operator identity or any single network indicator; understanding these differences requires jointly considering radio-resource allocation, RAT and serving-cell attachment, spectrum configuration, link adaptation, and architecture mode.

REFERENCES

[1] 3GPP. 2020. *Policy and Charging Control Architecture.* Technical Specification TS 23.203. 3rd Generation Partnership Project (3GPP). https://www.3gpp.org/ftp/Specs/archive/23_series/23.203/

[2] 3GPP. 2024. *TS 23.501: System Architecture for the 5G System (5GS).* Technical Specification 23.501. 3rd Generation Partnership Project.

[3] 3GPP. 2024. *TS 37.340: NR; Multi-connectivity; Overall Description; Stage 2.* Technical Specification 37.340. 3rd Generation Partnership Project.

[4] 3GPP. 2024. *TS 38.300: NR; NR and NG-RAN Overall Description; Stage 2.* Technical Specification 38.300. 3rd Generation Partnership Project.

[5] Accuver. 2026. Smartphone based network QoS and QoE logging solution including Voice MOS. https://www.accuver.com/products/network-optimization/XCAL-Solo3.

[6] Sergi Alcalá-Marín, Aravindh Raman, Weili Wu, Andra Lutu, Marcelo Bagnulo, Ozgu Alay, and Fabián Bustamante. 2022. Global Mobile Network Aggregators: Taxonomy, Roaming Performance and Optimization. In *Proc. of ACM MobiSys.*

[7] AT&T. 2026. Turbo Live by AT&T: Premium Data Boost for Live Events. https://www.att.com/plans/att-turbo-live/. Accessed 2026.

[8] Eduardo Baena, Zihan Xu, Moinak Ghoshal, Imran Khan, Phuc Dinh, Z. Jonny Kong, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2025. Root Cause Analysis of Cellular Network Throughput Degradations under Vehicular Mobility. In *Proc. of IEEE MASS.*

[9] Omar Basit, Phuc Dinh, Imran Khan, Z. Jonny Kong, Y. Charlie Hu, Dimitrios Koutsonikolas, Myungjin Lee, and Chaoyue Liu. 2024. On the Predictability of Fine-Grained Cellular Network Throughput Using Machine Learning Models. In *Proc. of IEEE MASS.*

[10] Omar Basit, Imran Khan, Moinak Ghoshal, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2025. 5G Metamorphosis: A Longitudinal Study of 5G Performance from the Beginning. In *Proc. of ACM IMC.*

[11] Broadband Map. 2026. Priority On Cell Networks – QCI & 5QI. https://broadbandmap.com/priority/. Accessed 2026.

[12] Giuseppe Caso, Mohammad Rajiullah, Konstantinos Kousias, Usman Ali, Nadir Bouzar, Luca De Nardis, Anna Brunstrom, Ozgu Alay, Marco Neri, and Maria-Gabriella Di Benedetto. 2024. The Chronicles of 5G Non-Standalone: An Empirical Analysis of Performance and Service Evolution. *IEEE Open Journal of the Communications Society* 5 (2024), 7380–7399.

[13] Giuseppe Caso, Mohammad Rajiullah, Konstantinos Kousias, Usman Ali, Luca De Nardis, Anna Brunstrom, Ozgu Alay, Marco Neri, and Maria-Gabriella Di Benedetto. None. An Initial Look into the Performance Evolution of 5G Non-Standalone Networks. In *Proc. of IFIP/IEEE TMA.*

[14] Fukun Chen, Moinak Ghoshal, Enfu Nan, Phuc Dinh, Imran Khan, Z. Jonny Kong, Y. Charlie Hu, and Dimitrios Koutsonikolas. None. A Large-Scale Study of the Potential of Multi-carrier Access in the 5G Era. In *Proc. of PAM.*

[15] DeadCellZones. 2024. Top 50 MVNOs Ranked by Customers and Size (2024). https://www.deadcellzones.com/top-50-mvnos.html. Accessed 2026.

[16] Phuc Dinh, Moinak Ghoshal, Yunmeng Han, Yufei Feng, Dimitrios Koutsonikolas, and Joerg Widmer. 2025. Demystifying Resource Allocation Policies in Operational 5G mmWave Networks. *IEEE/ACM Transactions on Networking* 33 (2025), 745–760. Issue 2.

[17] ESnet and Lawrence Berkeley National Laboratory. 2026. iPerf3 - The TCP, UDP and SCTP network bandwidth measurement tool. https://software.es.net/iperf/

[18] Yufei Feng, Phuc Dinh, Moinak Ghoshal, Eduardo Baena, Hamza Bouchebbah, and Dimitrios Koutsonikolas. None. Vivisecting Beam Management in Operational 5G mmWave Networks. In *Proc. of ACM CoNEXT.*

[19] Yufei Feng, Phuc Dinh, Moinak Ghoshal, Omar Basit, Sizhe Wang, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2026. Disentangling the Throughput Contributions of MIMO and Carrier Aggregation in 5G Networks. In *Proc. of PAM.*

[20] Yufei Feng, Jin Wei, Phuc Dinh, Moinak Ghoshal, and Dimitrios Koutsonikolas. 2023. Beam Management in Operational 5G mmWave Networks. In *Proc. of ACM mmNets Workshop.*

[21] Rostand Fezeu, Claudio Fiandrino, Eman Ramadan, Jason Carpenter, Lilian Freitas, Faaiq Bilal, Wei Ye, Joerg Widmer, Feng Qian, and Zhi-Li Zhang. None. Unveiling the 5G Mid-Band Landscape: From Network Deployment to Performance and Application QoE. In *Proc. of ACM SIGCOMM.*

[22] Rostand Fezeu, Eman Ramadan, Wei Ye, Benjamin Minneci, Jack Xie, Arvind Narayanan, Ahmad Hassan, Feng Qian, Zhi-Li Zhang, Jaideep Chandrashekar, and Myungjin Lee. 2023. An In-Depth Measurement Analysis of 5G mmWave PHY Latency and its Impact on End-to-End Delay. In *Proc. of PAM.*

[23] Rostand A. K. Fezeu, Jason Carpenter, Eman Ramadan, Faaiq Bilal, Ziyan Wu, Nanditha Naik, Duncan Joly, Dehkontee Chea Cuppah, and Zhi-Li Zhang. 2025. A Peek into 5G NSA vs. SA Control Plane Performance. In *Proc. of ACM HotMobile.*

[24] Rostand A. K. Fezeu, Claudio Fiandrino, Eman Ramadan, Jason Carpenter, Daqing Chen, Yiling Tan, Feng Qian, Joerg Widmer, and Zhi-Li Zhang. 2024. Roaming across the European Union in the 5G Era: Performance, Challenges, and Opportunities. In *Proc. of IEEE INFOCOM.*

[25] Claudio Fiandrino, David Juárez Martínez-Villanueva, and Joerg Widmer. None. Uncovering 5G Performance on Public Transit Systems with an App-based Measurement Study. In *Proc. of ACM MSWiM.*

[26] Moinak Ghoshal, Omar Basit, Imran Khan, Z. Jonny Kong, Sizhe Wang, Yufei Feng, Phuc Dinh, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2025. Replication: Performance of Cellular Networks on the Wheels. In *Proc. of ACM IMC.*

[27] Moinak Ghoshal, Omar Basit, Sizhe Wang, Phuc Dinh, Imran Khan, Yufei Feng, Zhekun Yu, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2025. A First Large-Scale Study of Operational 5G Standalone Networks. In *Proc. of ACM CoNEXT.*

[28] Moinak Ghoshal, Imran Khan, Phuc Dinh, Z. Jonny Kong, Omar Basit, Sizhe Wang, Yufei Feng, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2025. Handover Configurations in Operational 5G Networks: Diversity, Evolution, and Impact on Performance. arXiv:2511.03116.

[29] Moinak Ghoshal, Imran Khan, Z. Jonny Kong, Phuc Dinh, Jiayi Meng, Y. Charlie Hu, and Dimitrios Koutsonikolas. None. Performance of Cellular Networks on the Wheels. In *Proc. of ACM IMC.*

[30] Ahmad Hassan, Arvind Narayanan, Anlan Zhang, Wei Ye, Ruiyang Zhu, Shuowei Jin, Jason Carpenter, Z. Morley Mao, Feng Qian, and Zhi-Li Zhang. None. Vivisecting Mobility Management in 5G Cellular Networks. In *Proc. of ACM SIGCOMM.*

[31] Michail Kalntis, José Suárez-Varela, Jesús Omaña Iglesias, Anup Kiran Bhattacharjee, George Iosifidis, Fernando A. Kuipers, and Andra Lutu. None. Through the Telco Lens: A Countrywide Empirical Study of Cellular Handovers. In *Proc. of ACM IMC.*

[32] Imran Khan, Moinak Ghoshal, Joana Angjo, Sigrid Dimce, Mushahid Hussain, Paniz Parastar, Yenchia Yu, Claudio Fiandrino, Charalampos Orfanidis, Shivang Aggarwal, Ana C. Aguiar, Ozgu Alay, Carla F. Chiasserini, Falko Dressler, Y. Charlie Hu, Steven Y. Ko, Dimitrios Koutsonikolas, and Joerg Widmer. 2024. How Mature is 5G Deployment? A Cross-Sectional, Year-Long Study of 5G Uplink Performance. In *Proc. of IFIP/IEEE Networking.*

[33] Konstantinos Kousias, Mohammad Rajiullah, Giuseppe Caso, Ozgu Alay, Anna Brunstrom, Usman Ali, Luca De Nardis, Marco Neri, and Maria-Gabriella Di Benedetto. 2024. Empirical Performance Analysis and ML-Based Modeling of 5G Non-Standalone Networks. *Computer Networks* 241 (2024), 110207.

[34] Konstantinos Kousias, Mohammad Rajiullah, Giuseppe Caso, Ozgu Alay, Anna Brunstrom, Luca De Nardis, Marco Neri, Usman Ali, and Maria-Gabriella Di Benedetto. None. Coverage and Performance Analysis of 5G Non-Standalone Deployments. In *Proc. of ACM WiNTECH.*

[35] Joao Moreno, Marcello Contini, and Ana Aguiar. 2024. 5G NSA Performance: A Measurement Study. In *Proc. of IFIP WONS.*

[36] Arvind Narayanan, Eman Ramadan, Jason Carpenter, Qingxu Liu, Yu Liu, Feng Qian, and Zhi-Li Zhang. None. A First Look at Commercial 5G Performance on Smartphones. In *Proc. of ACM WWW.*

[37] Arvind Narayanan, Eman Ramadan, Rishabh Mehta, Xinyue Hu, Qingxu Liu, Rostand A. K. Fezeu, Udhaya Kumar Dayalan, Saurabh Verma, Peiqi Ji, Tao Li, Feng Qian, and Zhi-Li Zhang. None. Lumos5G: Mapping and Predicting Commercial mmWave 5G Throughput. In *Proc. of ACM IMC.*

[38] Arvind Narayanan, Muhammad Iqbal Rochman, Ahmad Hassan, Bariq S. Firmansyah, Vanlin Sathya, Monisha Ghosh, Feng Qian, and Zhi-Li Zhang. 2022. A Comparative Measurement Study of Commercial 5G mmWave Deployments. In *Proc. of IEEE INFOCOM.*

[39] Arvind Narayanan, Xumiao Zhang, Ruiyang Zhu, Ahmad Hassan, Shuowei Jin, Xiao Zhu, Xiaoxuan Zhang, Denis Rybkin, Zhengxuan Yang, Zhuoqing Morley Mao, Feng Qian, and Zhi-Li Zhang. None. A Variegated Look at 5G in the Wild: Performance, Power, and QoE Implications. In *Proc. of ACM SIGCOMM.*

[40] NUWiNS Lab. 2026. MVNO–Host Performance Comparison Dataset. https://github.com/Zack-ZeyuLi/NU-VMNOs. Created in 2025-2026.

[41] Paniz Parastar, Andra Lutu, Ozgu Alay, Giuseppe Caso, and Diego Perino. None. Spotlight on 5G: Performance, Device Evolution and Challenges from a Mobile Operator Perspective. In *Proc. of IEEE INFOCOM.*

[42] Reddit. 2025. MobileX is no longer QCI 8, which leaves US Mobile as the only independent MVNO besides the cable companies to have QCI 8 data. https://www.reddit.com/r/NoContract/comments/1n2me7i/mobilex_is_no_longer_qci_8_which_leaves_us_mobile/. Reddit discussion, r/NoContract.

[43] Reddit. 2025. QCI real talk. https://www.reddit.com/r/NoContract/comments/1hxx2yw/qci_real_talk/. Reddit discussion, r/NoContract.

[44] Muhammad Iqbal Rochman, Vanlin Sathya, Damian Fernandez, Norlen Nunez, Ahmed S. Ibrahim, William Payne, and Monisha Ghosh. 2023. A Comprehensive Analysis of the Coverage and Performance of 4G and 5G Deployments. *Computer Networks* 237 (2023), 110060.

[45] Muhammad Iqbal Rochman, Wei Ye, Zhi-Li Zhang, and Monisha Ghosh. 2025. A Comprehensive Real-World Evaluation of 5G Improvements Over 4G in Low- and Mid-Bands. *IEEE Transactions on Cognitive Communications and Networking* 11 (2025), 1427–1440. Issue 3.

[46] T-Mobile. 2026. Compare Our Best Prepaid Phone Plans. https://prepaid.t-mobile.com/prepaid-plans. Accessed 2026.

[47] Narseo Vallina-Rodriguez, Srikanth Sundaresan, Christian Kreibich, Nicholas Weaver, and Vern Paxson. 2015. Beyond the Radio: Illuminating the Higher Layers of Mobile Networks. In *Proc. of ACM MobiSys.*

[48] Verizon. 2026. Best Prepaid Phone Plans. https://www.verizon.com/prepaid/plans/phone-plans/. Accessed 2026.

[49] Sizhe Wang, Moinak Ghoshal, Yufei Feng, Imran Khan, Phuc Dinh, Omar Basit, Zhekun Yu, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2026. Exploring the 5G Digital Divide in the Non-Contiguous US: LEO Satellites to the Rescue?. In *Proc. of ACM SIGMETRICS*.

[50] Ao Xiao, Yunhao Liu, Yang Li, Feng Qian, Zhenhua Li, Sen Bai, Yao Liu, Tianyin Xu, and Xianlong Xin. 2019. An In-depth Study of Commercial MVNO: Measurement and Optimization. In *Proc. of ACM MobiSys*.

[51] Dongzhu Xu, Anfu Zhou, Xinyu Zhang, Guixian Wang, Xi Liu, Congkai An, Yiming Shi, Liang Liu, and Huadong Ma. None. Understanding Operational 5G: A First Measurement Study on Its Coverage, Performance and Energy Consumption. In *Proc. of ACM SIGCOMM*.

[52] Xinlei Yang, Hao Lin, Zhenhua Li, Feng Qian, Xingyao Li, Zhiming He, Xudong Wu, Xianlong Wang, Yunhao Liu, Zhi Liao, et al. 2022. Mobile Access Bandwidth in Practice: Measurement, Analysis, and Implications. In *Proc. of ACM SIGCOMM*.

[53] Wei Ye, Jason Carpenter, Zejun Zhang, Rostand A. K. Fezeu, Feng Qian, and Zhi-Li Zhang. 2023. A Closer Look at Stand-Alone 5G Deployments from the UE Perspective. In *Proc. of IEEE MeditCom*.

[54] Wei Ye, Xinyue Hu, Steven Sleder, Anlan Zhang, Udhaya Dayalan, Ahmad Hassan, Rostand Fezeu, Akshay Jajoo, Myungjin Lee, Eman Ramadan, Feng Qian, and Zhi-Li Zhang. None. Dissecting Carrier Aggregation in 5G Networks: Measurement, QoE Implications and Prediction. In *Proc. of ACM SIGCOMM*.

[55] Fatima Zarinni, Ayon Chakraborty, Vyas Sekar, Samir R. Das, and Phillipa Gill. 2014. A First Look at Performance in Mobile Virtual Network Operators. In *Proc. of ACM IMC*.

APPENDIX

A ETHICS

All measurements were conducted using commercial subscriptions and ordinary user-plane traffic from our own devices, without modifying or interfering with network infrastructure. Each throughput test lasted one minute and involved at most three simultaneously active UEs. We use TCP streams for data transmission to ensure that network congestion control is effective. All collected data originated from our own devices and contain no personally identifiable information or traffic from other users.

B OPERATOR SUBSCRIPTIONS

Table 9 lists the nine subscriptions used in the measurements.

Table 9: Operator subscriptions and operator-advertised plan descriptions used in the measurements.

Operator	Type	Plan Description
HA	Host	Unlimited data
VA1	MVNO	Unlimited data
VA2	MVNO	50GB “high-speed” data/month
HB	Host	Unlimited data
VB1	MVNO	Unlimited data
VB2	MVNO	50GB “high-speed” data/month
HC	Host	50GB “premium” data/month
VC1	MVNO	50GB “premium” data/month
VC2	MVNO	50GB “premium” data/month

C MEASUREMENT LOCATIONS

Fig. 22 shows the locations where all measurements were conducted.

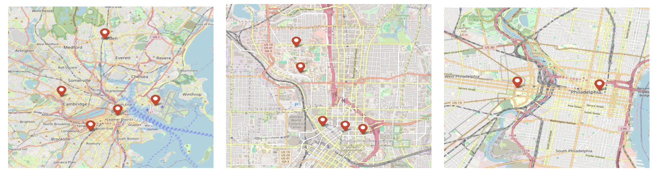

(a) Boston. (b) Atlanta. (c) Philadelphia.

Figure 22: Measurement locations across the three cities. Each pin marks one test site.

D MEASUREMENT TIME RANGE

Table 10 shows the time range of every measurement campaign.

Table 10: Measurement campaign timeline over all valid rounds.

Fam.	City	Site	Rds	Date	Local time
A	Atlanta	1	1	2026-02-26	16:47–17:00
		2	5	2026-02-26	17:35–18:53
		3	5	2026-02-27	09:51–10:37
		4	5	2026-02-27	13:16–14:04
		5	5	2026-02-27	15:04–15:45
	Boston	1	5	2025-12-30	13:43–15:29
		2	5	2025-12-31	15:35–17:16
		3	5	2025-12-31	18:38–20:35
		4	5	2026-01-01	14:50–16:21
		5	5	2026-01-07	17:12–18:47
B	Boston	1	5	2026-04-01	20:37–22:30
		2	5	2026-04-02	10:50–12:56
		3	5	2026-04-02	13:30–15:08
	Atlanta	1	5	2026-02-28	10:47–12:20
		2	5	2026-02-28	13:04–14:38
		3	5	2026-02-28	15:22–16:42
		4	5	2026-02-28	17:50–18:36
		5	5	2026-02-28	20:07–20:42
C	Boston	1	5	2025-09-28	15:06–17:30
		2	5	2025-09-29	14:53–18:08
		3	5	2025-10-23	12:14–14:27
		4	5	2025-10-01	12:12–14:38
		5	5	2025-10-01	16:18–18:52
	Philadelphia	1	4	2026-02-07	12:23–13:31
		2	5	2026-02-07	15:32–17:14

E MATCHED RADIO CONTEXT

Figs. 23, 24, 25, 26 report the PCell RB allocation and MAC-layer throughput for the rounds in which every participating operator is served by the same PCI, NR-ARFCN, and bandwidth. Figs. 23, 24 cover Family B with an 8/8/8 QoS-class pattern, and support the discussion in §6.2; Figs. 25, 26 cover Family C with a 6/7/7 QoS-class pattern, and support the discussion in §7.1.

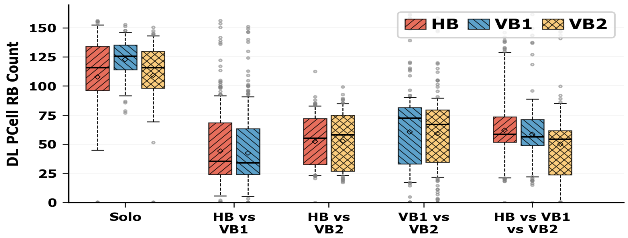

(a) PCell RB allocation.

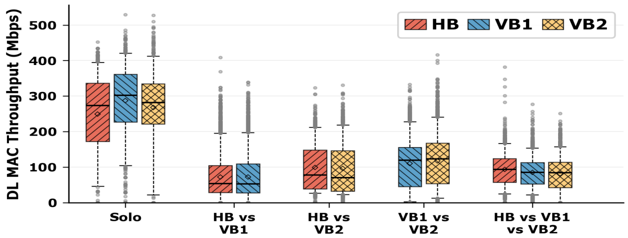

(b) PCell MAC-layer throughput.

Figure 23: Family B downlink under matched radio context (same PCI, NR-ARFCN, bandwidth).

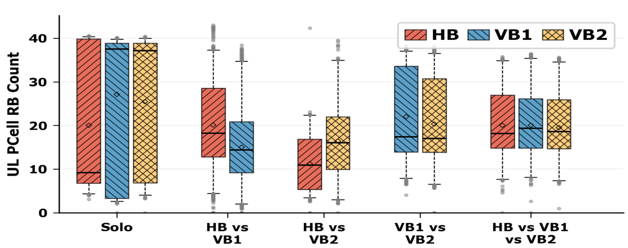

(a) PCell RB allocation.

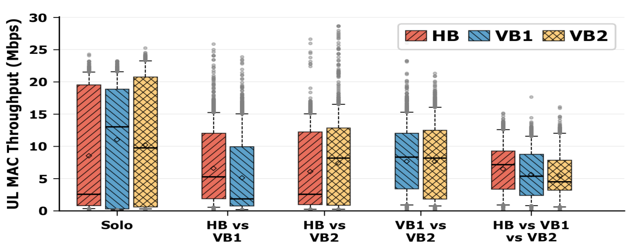

(b) PCell MAC-layer throughput.

Figure 24: Family B uplink under matched radio context (same PCI, NR-ARFCN, bandwidth).

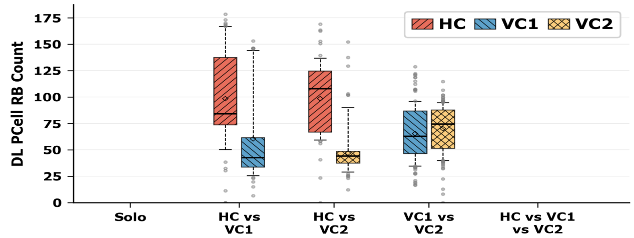

(a) PCell RB allocation.

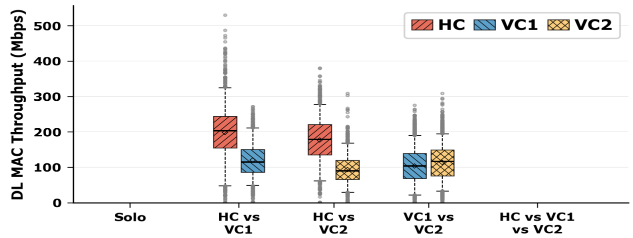

(b) PCell MAC-layer throughput.

Figure 25: Family C downlink under matched radio context (same PCI, NR-ARFCN, bandwidth).

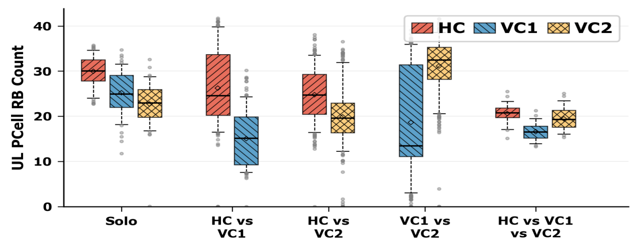

(a) PCell RB allocation.

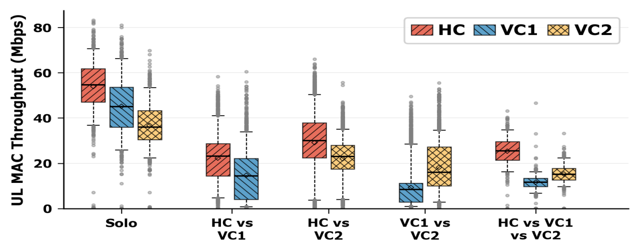

(b) PCell MAC-layer throughput.

Figure 26: Family C uplink under matched radio context (same PCI, NR-ARFCN, bandwidth).