Too Many Reports, Too Few Handovers: A Large-Scale Study of Configuration Inefficiencies in Operational 5G Networks

MOINAK GHOSHAL, IMRAN KHAN, and PHUC DINH, Northeastern University, USA
Z. JONNY KONG, Meta Platforms, USA
OMAR BASIT, Purdue University, USA
SIZHE WANG and YUFEI FENG, Northeastern University, USA
Y. CHARLIE HU, Purdue University, USA
DIMITRIOS KOUTSONIKOLAS, Northeastern University, USA

Mobility management in cellular networks, and in particular the handover (HO) process, is central to providing seamless Internet access. The wide-scale deployment of 5G has greatly increased its complexity, yet how operators configure HOs, and whether those configurations are efficient, remains largely unknown. Through four cross-country US driving trips over a 27-month period, we conduct the first large-scale measurement study of HO configurations across all three major US operators. We find that current configurations are often overly aggressive: devices report far more measurement events than necessary, with only 13–31% actually triggering HOs, and we show this overhead is a direct consequence of configuration choices rather than an inherent property of the events. Despite the large configuration spaces operators maintain, actual HO behavior concentrates in a tiny dominant subset. Finally, optimizing configurations is fundamentally challenging: the metrics operators directly control are only weakly correlated with user-perceived performance, and a deployment-density confound makes it hard to tell whether a configuration is well-calibrated. Our findings have concrete implications for operators seeking to reduce signaling overhead and improve HO quality in 5G networks.

CCS Concepts: • **Networks** → **Mobile networks**; **Network measurement**; **Network performance analysis**.

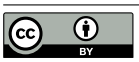
Additional Key Words and Phrases: 5G, Measurement, Network Performance, Dataset

ACM Reference Format:
Moinak Ghoshal, Imran Khan, Phuc Dinh, Z. Jonny Kong, Omar Basit, Sizhe Wang, Yufei Feng, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2026. Too Many Reports, Too Few Handovers: A Large-Scale Study of Configuration Inefficiencies in Operational 5G Networks. *Proc. ACM Netw.* 4, CoNEXT4, Article 64 (December 2026), 22 pages. https://doi.org/10.1145/3856978

1 Introduction

Mobility management in cellular networks, and in particular the handover (HO) process, plays a critical role in providing seamless and ubiquitous Internet access. In operational networks, HO decisions are governed by a policy-based approach: each HO is triggered based on signal

Authors' Contact Information: Moinak Ghoshal, ghoshal.m@northeastern.edu; Imran Khan, khan.i@northeastern.edu; Phuc Dinh, Northeastern University, Boston, MA, USA, dinh.p@northeastern.edu; Z. Jonny Kong, Meta Platforms, Menlo Park, CA, USA, kong102@purdue.edu; Omar Basit, Purdue University, West Lafayette, IN, USA, obasit@purdue.edu; Sizhe Wang, wang.sizh@northeastern.edu; Yufei Feng, Northeastern University, Boston, MA, USA, feng.yuf@northeastern.edu; Y. Charlie Hu, Purdue University, West Lafayette, IN, USA, ychu@purdue.edu; Dimitrios Koutsonikolas, Northeastern University, Boston, MA, USA, d.koutsonikolas@northeastern.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License.
© 2026 Copyright held by the owner/author(s).
ACM 2834-5509/2026/12-ART64
https://doi.org/10.1145/3856978

strength conditions, cell priorities, and a set of measurement events, all controlled by configuration parameters that operators set freely within the 3GPP framework [2].

The wide-scale deployment of 5G over the past six years has significantly increased the complexity of mobility management. 5G's support for diverse radio bands (low, mid, high), different architectures (non-standalone (NSA) vs. standalone (SA)), and the coexistence of LTE and 5G radio access technologies (RATs) make HO management in the 5G era far more complex than before. Prior measurement studies have shown that HOs can significantly degrade throughput and increase latency [8, 11–14, 18, 19, 29, 30, 34]. However, these studies focus on the *outcome* of the HO process — its impact on user performance — without examining its *root cause*: the configuration policies that govern when and how HOs are triggered. Whether current HO configurations are efficient, or whether they introduce unnecessary overhead and suboptimal decisions, remains largely unknown.

The answer matters because HOs are not free. Before an HO can occur, the network configures the UE with a set of reporting events, criteria that specify what to measure and when to report. The UE then measures the serving and neighboring cells and sends a measurement report (MR) whenever a criterion is met. Every report consumes uplink radio resources and signaling capacity, whether or not it leads to an HO. Configurations therefore determine not only when HOs occur, but also how much signaling the network spends to make them happen. Yet 3GPP only defines the configuration parameters, without prescribing how to set them efficiently. Assessing how well operators make this choice requires examining their configuration policies directly, across operators and at scale, which, to the best of our knowledge, no prior measurement study has done in the 5G era. We therefore revisit mobility management from this configuration-centric perspective and ask a fundamental question: *Are current HO configurations in operational 5G networks efficient and well-calibrated?*

To answer this question, we conduct the first large-scale measurement study of HO configurations in the 5G era. Through four cross-country driving trips (Los Angeles to Boston in Aug. '22 and Nov. '24, Boston to Chicago in May '23, and Boston to Atlanta in Aug. '23) covering 15,000+ km over a 27-month period, we collect a dataset of 48K HOs and 239K HO-related reporting events spanning all three major US operators (Verizon, AT&T, T-Mobile), both RATs (LTE and 5G), both 5G architectures (NSA and SA), and all 5G frequency bands (low, mid, high). We analyze this dataset along four axes: the HO-triggering event landscape, the diversity and evolution of configuration parameters, the aggressiveness of those configurations and its role in generating overhead, and the impact of configurations on signal strength and ping-pong performance. Our key findings are:

• **HO-triggering events (§4).** The 5G HO landscape differs fundamentally from LTE: new inter-RAT events (B1, B2) have appeared while periodic events have been practically eliminated. Despite the large overall event volume, actual HO behavior converges on a single dominant event per HO type. The overhead is strongly bimodal — events are either largely decisive (A3, B1, A5) or almost entirely wasted (A1, A6, periodic) — and the conversion rate of the same event varies up to 10x across operators, confirming that overhead is a configuration choice rather than an inherent property of the event.

• **Diversity and evolution of configurations (§5.1).** LTE configurations have grown substantially more complex since prior work [7], with some parameters now taking 43–56 distinct values. 5G configurations have lower richness (fewer than 12 distinct values for most parameters) but still exhibit large dispersion, indicating that operators tune configurations differently across bands, area types, and deployment regions. LTE configurations have largely stabilized, while 5G configurations continue to diversify rapidly, with some parameters gaining more than 10 new values in two years.

• **Aggressiveness, overhead, and operational concentration (§5.2,§5.3).** We introduce a framework to classify each configuration as aggressive, moderate, or conservative. The decisive events are configured more tightly in 5G than LTE, showing that operators have learned to apply discipline

where it matters most, but have not yet extended it to other events. Finally, despite the large and diverse configuration spaces operators provision, actual HO traffic concentrates in a tiny dominant subset. Most provisioned configurations are operationally irrelevant and this concentration is tighter in 5G than LTE.

• **Impact on performance (§6).** Optimizing HO configurations is fundamentally challenging: no single configuration simultaneously optimizes post-HO signal improvement, pre-HO signal degradation, and ping-pong rate, and a pervasive coverage confound means counter-intuitive outcomes across operators and configurations are frequently explained by deployment density rather than configuration quality. TTT is the most reliable lever for ping-pong control across all event types. Most strikingly, Verizon's 5G→LTE demotion strategy, while rational from a signal strength perspective, results in throughput degradation 80% of the time, a cost that RSRP-based optimization cannot capture.

Contributions. In summary, this paper makes the following contributions: (1) We collect a large multi-carrier, multi-technology, multi-band dataset of cellular mobility management during four cross-country driving campaigns spanning 15,000+ km over a 27-month period. (2) We develop a revised methodology to accurately identify HO-triggering events in 5G NSA deployments, correcting a flawed assumption in prior work that affects 12% of HOs. (3) Using the collected dataset and revised methodology, we analyze the HO-triggering event landscape, the diversity and evolution of configuration parameters, the aggressiveness of those configurations and its role in generating overhead, and the impact of configurations on signal strength and ping-pong performance. (4) Our dataset and processing scripts are publicly released [1].

2 HO Primer

A HO is used to transfer a user's network session from the cell it is attached to (serving cell) to another cell (neighboring cell), because either the user moves out of the serving cell's coverage area or there is another cell operating at a different RAT. HOs can be classified into idle-state HOs and active-state HOs, based on whether the UE is idle or sending/receiving traffic to/from the network. Idle-state HOs are performed by the device, which selects an appropriate cell for future access. Active-state HOs are initiated by the network; the serving cell migrates the UE to another neighboring cell. In this work, we focus on active-state HOs.

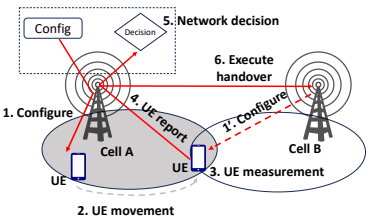


Fig. 1. Basic HO procedure.

Fig. 1 shows a basic cellular HO procedure. When a UE connects to a new cell, the cell advertises an HO configuration (1) to the UE. Based on the configuration, the UE learns the criteria to trigger an HO, including whether to invoke measurement, what and when to measure, and when/what to report. As the UE starts moving from the coverage area of Cell A to a new coverage area under Cell B, it starts measuring (3) different signal metrics (RSRP, RSRQ). 3GPP defines a set of measurement report (MR) mechanisms based on *events*, which are triggered when certain signal-related criteria are met. When an event is triggered, the UE sends an MR (4), which the network uses to decide (5) whether to perform an HO or not. Based on the network's decision, the HO is executed (6) from Cell A to B, followed by a new set of configurations (1') advertised from Cell B to the UE.

Reporting events. Table 1 lists the HO-related reporting events defined by 3GPP and observed in our dataset, together with their triggering conditions. In the following, we briefly describe each one and its role in mobility management.

• *Intra-RAT events (A1–A6)* are triggered by conditions on the signal quality of the serving cell (M_S) or a neighboring cell (M_N) *within* the same RAT. **A1** is triggered when the serving cell becomes *better* than a threshold. Since this indicates that the serving link has recovered, A1 is typically

Table 1. LTE/5G-NR HO-related reported events
(M = measurement, C = serving cell, N = neighboring cell)

Event Type	Event Description	trigger Condition
A1	Serving cell becomes better than a threshold	$M_S > \Theta_{A_1} + H_{A_1}$
A2	Serving cell becomes worse than a threshold	$M_S < \Theta_{A_2} - H_{A_2}$
A3	Neighboring cell becomes offset better than the serving cell	$M_N > M_S + \Delta_{A_3} + H_{A_3}$
A4	Neighboring cell becomes better than threshold	$M_N > \Theta_{A_4} + H_{A_4}$
A5	Serving cell becomes worse than a threshold and neighboring cell becomes better than another threshold	$M_S < \Theta^S_{A_5} - H_{A_5}$ $M_N > \Theta^N_{A_5} + H_{A_5}$
A6	Neighboring cell becomes offset better than the serving cell	$M_N > M_S + \Delta_{A_6} + H_{A_6}$
B1	Inter-RAT neighbor becomes better than threshold	$M_N > \Theta_{B_1} + H_{B_1}$
B2 (5G only)	Primary cell becomes worse than a threshold and inter-RAT neighbor becomes better than a threshold	$M_S < \Theta^S_{B_2} - H_{B_2}$ $M_N > \Theta^N_{B_2} + H_{B_2}$
P	Periodic reporting of serving and neighboring cell's radio signal quality.	N/A

used to cancel a pending HO or stop ongoing neighbor measurements rather than to trigger an HO. **A2** is its counterpart: it is triggered when the serving cell becomes *worse* than a threshold and is used to prepare for an HO, initiate neighbor measurements, or facilitate fallback to another RAT. **A3** is triggered when a neighbor becomes better than the serving cell by an *offset*; because it directly compares the serving and neighboring cells, it is the primary event used for intra-RAT HOs. **A4** is triggered when a neighbor becomes better than an *absolute* threshold, irrespective of the serving-cell quality; it is commonly used for mobility and load-balancing decisions, and in our dataset appears only in 5G. **A5** combines two conditions: the serving cell must fall below one threshold *while* a neighbor rises above another. Requiring both conditions makes A5 more selective than A4, since an HO is considered only after the serving link has degraded. **A6** is structurally similar to A3, but compares a neighbor against a *secondary* cell (SCell) rather than the primary serving cell. It is therefore primarily used for secondary-cell management under carrier aggregation (CA), rather than for primary-cell HOs.

• *Inter-RAT events (B1, B2)* handle transitions between LTE and 5G and mirror the A4/A5 intra-RAT events. **B1** is triggered when an inter- RAT neighbor becomes better than an absolute threshold and is commonly used to drive LTE→5G promotions. **B2** additionally requires the serving cell to fall below a threshold and is typically used for 5G SA→LTE demotions.

• *Periodic reporting (P)* is not condition-based. Instead, the UE periodically reports the signal quality of the serving and neighboring cells at a fixed interval, regardless of whether a particular triggering condition is met. It can therefore generate reports continuously and, as we show in §4, only rarely precedes an HO in our dataset.

Configuration parameters. Each event is governed by the parameters listed in Table 4 in Appendix in §B. The three most relevant to our analysis are (a) the *threshold* Θ or *offset* Δ, which defines the required signal condition, (b) the *hysteresis* margin H, which prevents rapid back-and-forth triggering around the decision boundary, and (c) the *time-to-trigger* (TTT), which specifies how long the condition must remain satisfied before a report is sent. A fourth parameter, the *trigger quantity*, determines whether the condition is evaluated using RSRP or RSRQ. Together, the threshold/offset and TTT determine how *easily* and how *quickly* an event is triggered — the two dimensions along which we characterize configuration aggressiveness in §5.2.

While events and parameters are standardized by 3GPP, operators are free to select different configurations (i.e., values for each parameter), implementing different HO policies. This is the central subject of this paper.

3 Methodology

UEs and measurement tools. We used Samsung S21 Ultra (in 2022, 2023) and S24 (in 2024) phones as UEs, the state-of-the-art devices at the time of our measurement campaigns. Since we only focus

on *active-state* HOs, which are controlled by the network, *our findings are independent of the UE model and can be generalized for any smartphone model.* Each UE was connected to an Accuver XCAL Solo device [6], which extracts and logs PHY layer KPIs and RRC layer signaling messages such as HOs, configurations, and MRs. We performed measurements *concurrently* over the three operators, using 3 UEs, each connected to a different operator.

Driving Tests. We drove 15,000+ km during the four road trips covering all major cities on each route (Fig. 17a in §C). In addition to driving on highways, we conducted additional downtown driving tests in select major cities on our routes to obtain data from urban environments.

5G Carriers. We purchased multiple unlimited data plans from all three major US service providers, Verizon, T-Mobile, and AT&T, to perform the experiments. Fig. 17b in Appendix §C shows the overall cellular coverage landscape for all three operators in our dataset. The 5G coverage varies from 43.3% (for Verizon) up to 89.5% (T-Mobile). Further comparison of the coverage over the same route in 2022 and 2024 shows an increase in 5G coverage for all 3 operators: from 32% to 50% for AT&T, from 73% to 94% for T-Mobile, and from 33% to 54% for Verizon. Notably, all three operators exhibit substantial coverage of *both* LTE and 5G throughout our measurement campaigns, allowing us to observe and study the full range of HO types.

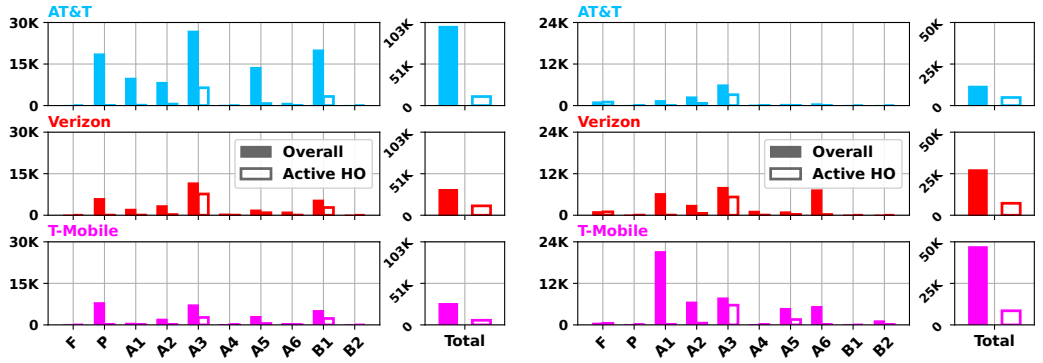
Traffic workload. Our analysis relies on the UE being in *active* (RRC-connected) state throughout the drives, since active-state HOs are only performed while the UE has traffic to send or receive. We therefore ran tests continuously on each UE for the entire duration of every trip, cycling in round-robin fashion through backlogged downlink/uplink throughput tests, ICMP-based ping RTT tests, and application tests — that is, a mix of backlogged (saturated) and non-backlogged traffic. For the performance analysis in §6, where we examine the impact of HO configurations on end-user throughput, we focus on the backlogged tests.

Choice of signal quality metrics. 3GPP permits an event to be evaluated on either RSRP or RSRQ, chosen per configuration via the *trigger quantity* parameter (Table 4), and we use whichever metric is appropriate for the analysis at hand. Our configuration-aggressiveness analysis (§5.2) spans both LTE and 5G, and AT&T and Verizon configure a substantial fraction of their LTE events on RSRQ; we therefore classify both RSRP- and RSRQ-based configurations there, using separate cutoffs for each (Table 3). Our performance-impact analysis (§6), in contrast, focuses on 5G HOs, where every configuration we observe in our dataset is RSRP-based; we consequently use RSRP for the post-HO signal-quality analysis.

Extracting HO-triggering events. Prior work [7] assumed the last reporting event before an HO is the one that triggered it. We found this assumption fails in 12% of HOs in our dataset, particularly in NSA deployments, and developed a revised methodology that inspects all MRs preceding an HO in both LTE and 5G UL DCCH messages to identify the one whose cell ID and frequency match the new cell. This approach uncovered several unusual cases: a B1 event, normally used for LTE→5G HOs, can sometimes trigger an LTE→LTE HO; certain A1/A2 reports include neighboring cell measurements, making them capable of triggering an HO directly; and for a subset of 5G HOs, complete configuration parameters must be traced back through multiple prior HOs due to the network sending only partial RRC reconfiguration messages when a UE reconnects to a previously visited cell.

4 HO Reporting Events: Diversity and Signaling Overhead

Distribution of HO-related reporting events. Figs. 2a and 2b break down the reporting events for each operator, for LTE and 5G respectively. For each event, we show both the total number of events reported by the UE ("Overall") and the number of events that actually led to an HO ("Active HO"). The three operators generate vastly different overall numbers of events: 98K/32K/26K (AT&T/Verizon/T-Mobile) in LTE and 11.5K/27K/47K in 5G.



(a) LTE. (b) 5G.

Fig. 2. Distribution of HO-related reporting events. F: 5G failures, P: periodic events.

Fig. 2a shows that AT&T triggers the highest number of reporting events in LTE, followed by Verizon, and T-Mobile the fewest. The lower event count for T-Mobile is consistent with its significantly lower LTE coverage compared to the other two operators (Fig. 17b). The distribution of events also differs substantially across operators, suggesting the usage of different event parameter configurations. While A3 is the dominant event for all three operators in LTE (7K–26K), AT&T triggers a much larger number of A1, A2, A5, and B1 events (8K–20K) compared to Verizon and T-Mobile (0.5K–5K). Notably, AT&T and Verizon have comparable LTE coverage (Fig. 17b), yet their event distributions differ significantly, highlighting operator-specific diversity in event triggering policies, which we explore in detail in §5.

While a similar disparity in triggered event counts exists for 5G, the ordering among operators reverses: T-Mobile generates the most 5G events, consistent with its ~90% 5G coverage, while AT&T generates the fewest. In terms of event distribution, however, we observe notable contrasts relative to LTE. While A3 remains the dominant event for AT&T and Verizon in 5G, consistent with their LTE behavior, T-Mobile's most frequently triggered event in 5G is A1, an event that was practically never triggered in LTE. A6 forms a significant fraction of total 5G events for both Verizon and T-Mobile compared to LTE. We also observe new event types absent in the pre-5G era: B1 (LTE→5G NSA) and B2 (5G SA→LTE), and various 5G failure events. Notably, periodic events (P) are practically absent in 5G across all operators.

Dominant triggering events. When we look at the events that actually triggered HOs in Figs. 2a and 2b, the picture is strikingly different from the overall distribution. The total number of HO-triggering events is very similar across operators (11K/11K/5K for AT&T/Verizon/T-Mobile in LTE and 4K/7K/8K in 5G), and only three events dominate: A3 (60–70% of all HOs), B1 (15–20%), and A5 (15% for T-Mobile, 5–6% for the other two). Fig. 2 summarizes the dominant triggering event for each HO type. Despite the increased complexity of the 5G HO landscape, operators converge on a single dominant event for each HO type. The majority of 5G→LTE demotions are an exception to this pattern: 60–69% are caused by various 5G failure events rather than by a reporting event, meaning they lie outside the RSRP-based HO framework entirely, and cannot be addressed through configuration tuning alone.

Table 2. Dominant triggering event per HO type.

HO type	Dominant triggering event
LTE→LTE	A3 (80-86%)
5G→5G intra-band	A3 (79-96%)
5G→5G inter-band	A5 (52-72%)
LTE→5G	B1 (98-99%)
5G (NSA)→LTE	5G failures (60-69%)
5G (SA)→LTE	B2 (75%)

Quantifying the event triggering overhead. Figs. 2a and 2b suggest that all three operators use aggressive configurations, forcing the UE to trigger a very large number of events, most of which are ignored by the network. We quantify this overhead as the ratio of reports that actually

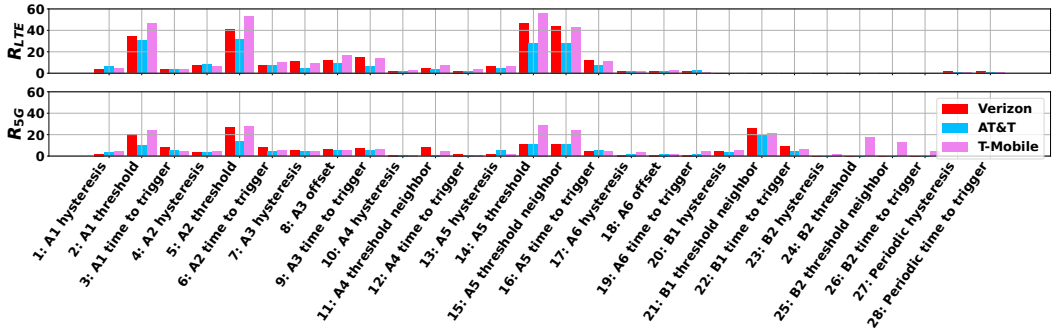


Fig. 4. Richness of HO configuration parameters.

triggered HOs to the total number of reports, for each operator and technology, and find that only 11%/21%/36% (LTE) and 45%/18%/27% (5G) of reported events actually triggered HOs for AT&T/T-Mobile/Verizon. Fig. 3 breaks down the overhead by event type, revealing three important patterns.

First, the conversion rate is strongly bimodal: events are either largely decisive (A3, B1, A5, A2, and A4 in 5G only) or almost entirely wasted. Periodic, A1, and A6 events convert at less than 3.5% across all operators and both technologies despite being reported in large volumes — these are the dominant source of overhead. The most extreme case is T-Mobile, which triggers 21K A1 events in 5G of which less than 1% actually trigger an HO. This is partly by design — A1 signals that the serving cell has recovered and cancels a pending HO, while A6 is responsible for carrier aggregation — but the large volume of these events represents a substantial and avoidable signaling cost. The absence of periodic events in 5G (Fig. 2b) suggests that operators are actually trying to reduce unnecessary signaling on the 5G interface, since periodic events do not typically lead to HOs for LTE (Fig. 2a); the same reasoning should be extended to A1 and A6.

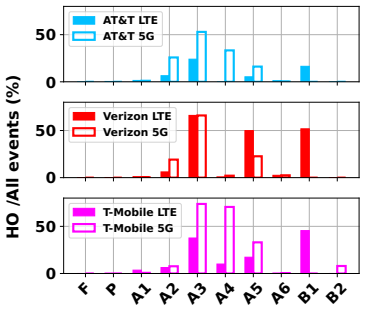


Fig. 3. Event overhead breakdown.

Second, the conversion rate of the same event varies dramatically across operators, confirming that overhead is a configuration choice rather than an inherent property of the event. In LTE, an A5 event triggers an HO 49.9% of the time for Verizon but only 5.3% for AT&T, nearly a 10x difference. Similar gaps exist for B1 (52% vs. 16%) and A3 (66% vs. 24%). AT&T is the least efficient operator overall: even its decisive events convert poorly in LTE, consistent with its lowest aggregate conversion rate (11%), whereas Verizon is the most efficient (36%).

Third, 5G configurations are generally tighter than LTE: for most decisive events, a larger fraction of reports lead to an HO in 5G than in LTE (e.g., A3 rises from 24% to 53% for AT&T and from 38% to 74% for T-Mobile), suggesting more targeted HO decisions in 5G. A notable exception is A5 for Verizon, whose conversion drops from 50% in LTE to 23% in 5G, consistent with the A4-equivalent A5 misconfiguration we identify in §6.1.2.

Summary. All three operators force the UE to report a very large number of events, the vast majority of which the network ignores — an inefficiency that is a direct consequence of configuration choices, which §5.2 examines in detail.

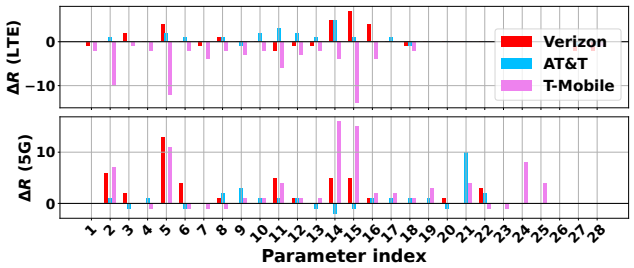


Fig. 5. Richness comparison, R(2024) - R(2022).

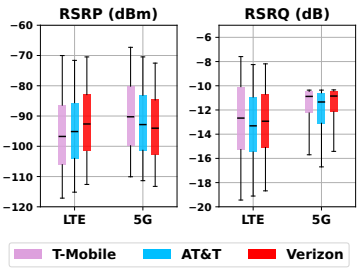


Fig. 6. RSRP/RSRQ distributions.

5 HO Configuration Parameters: Diversity and Aggressiveness

5.1 Diversity of HO configurations

Fig. 4 shows the richness R for 28 parameters related to different reporting events in LTE and 5G, which measures the number of unique values for a given parameter. We observe that most LTE parameters take multiple values. On the extreme end, A3 offset has 9–17 distinct values for each operator, and A1 threshold, A2 threshold, and the two A5 thresholds take 28–31 distinct values in AT&T, 34–44 in Verizon, and 43–56 in T-Mobile – even higher than those reported for AT&T in [7], where the largest richness observed was 21. Even parameters like A3 hysteresis, reported as single-valued in [7], now have 5–11 different values. The only parameters that exhibit low richness are those related to events used infrequently (A4, A6, and periodic). Interestingly, T-Mobile exhibits the highest richness for several parameters despite its lower LTE coverage, indicating that richness reflects operator policy rather than sampling volume. On the other hand, the richness of 5G values is much lower, with most parameters having fewer than 12 distinct values. This is true even for T-Mobile, despite its much higher 5G than LTE coverage (Fig. 17b), reinforcing our conjecture that observed richness is not an artifact of the number of encountered cells.

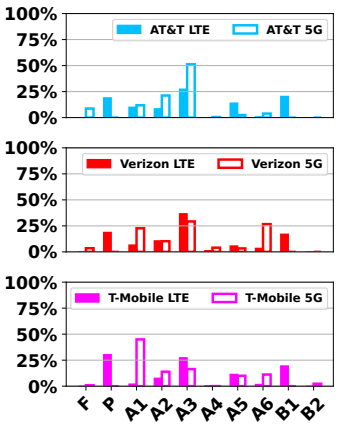


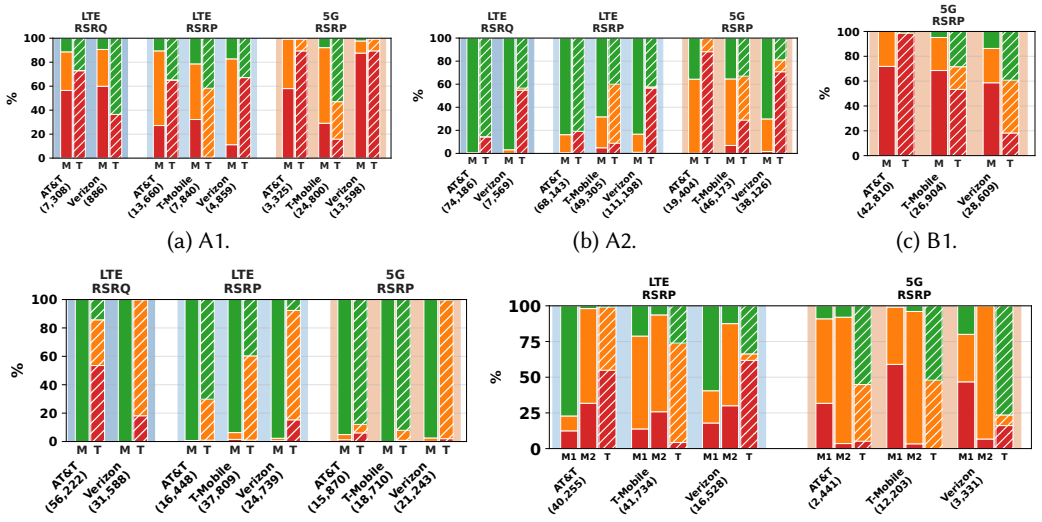
Fig. 7. Normalized triggered events.

Evolution. Fig. 5 shows the difference in richness for each parameter in 2022 vs. 2024: $\Delta R = R(2024) - R(2022)$. For LTE, ΔR is zero or negative for most parameters: operators have largely stopped tweaking LTE configurations as the technology matures, and the negative values for T-Mobile reflect its sharply reduced LTE coverage in 2024. For 5G, by contrast, ΔR is mostly positive, with some parameters gaining more than 10 new values in two years, reflecting ongoing adaptation to expanding coverage and evolving traffic demands. Notably, the parameters of the most popular event A3 (indexes 7, 8, 9) show only minor variation over the same period, suggesting that operators have converged on stable A3 configurations while continuing to fine-tune less dominant events.

5.2 *Aggressiveness* of HO Configurations

Fig. 7 normalizes the reported events in Fig. 2 by each operator's LTE/5G event count. This removes differences due to coverage and highlights how heavily each operator relies on each event type.

Defining aggressiveness. A HO configuration governs how easily an event is triggered through two major components: *metric*, which is the threshold (or offset) that the measured signal must cross, and the *time-to-trigger* (TTT) for which the condition must hold. We classify each observed configuration along both components as *aggressive*, *moderate*, or *conservative*. A configuration is aggressive if it makes the event trigger easily and quickly (an easily-satisfied threshold/offset and a



(a) A1. (b) A2. (c) B1.

(d) A3. (e) A5.

Fig. 8. Aggressiveness of an event to be triggered. M = Threshold/Offset Metric, M1 = Threshold Serving (A5), M2 = Threshold Neighbor (A5), T = TTT. Aggressive, Moderate, Conservative.

short TTT) and conservative if otherwise. We set these cutoffs per event: threshold-based events (A1, A2, B1, and the two A5 thresholds) use absolute RSRP/RSRQ signal-quality bands chosen relative to the operating range we observe across our measurements (Fig. 6), while offset-based events (A3, A6) use offset cutoffs adopted in the spirit of the 3GPP reference parameter sets [3]. For events that trigger when a cell's signal strength rises *above* a threshold (A1, B1, and the neighboring cells threshold of A5), a *low* threshold is aggressive; for events that trigger when the serving cell's signal strength falls *below* a threshold (A2 and the serving cell threshold of A5), a *high* threshold is aggressive; and for offset-based events (A3, A6), a *small* offset is aggressive. Table 3 shows the cutoffs we use to classify the metric component of each HO event configuration as aggressive or conservative. For TTT, we label ≤ 200 ms as aggressive and ≥ 640 ms as conservative, following the discrete 3GPP TTT ladder [4, 5].

Table 3. Aggressiveness classification.

Evt.	RSRP		RSRQ	
	Agg.	Con.	Agg.	Con.
A1	≤-115 dB	≥-90 dB	≤-15 dB	≥-12 dB
A2	≥-105 dB	≤-120 dB	≥-13 dB	≤-16 dB
A3	≤2 dB	≥5 dB	≤1 dB	≥2.5 dB
A5-S	≥-105 dB	≤-120 dB	≥-12 dB	≤-17 dB
A5-N	≤-115 dB	≥-90 dB	≤-16 dB	≥-12 dB
A6	≤3 dB	≥8 dB	≤2 dB	≥4 dB
B1	≤-110 dB	≥-95 dB	≤-18 dB	≥-12 dB

Fig. 8 shows, for each event, the breakdown of its configurations into aggressive, moderate, and conservative categories, separately for the metric M (threshold/offset) and the TTT (T), using the trigger quantities each operator actually deploys (RSRP and RSRQ for LTE, RSRP only for 5G; see §3).

We discuss the aggressiveness of event triggering by grouping the events into two categories: (1) events that do not directly lead to a HO, and (2) otherwise.

Group 1: Events that do not directly lead to a HO.

A1. A1 is triggered when the *serving* cell rises above a threshold, so a *low* threshold and a short TTT make it aggressive: the serving link needs to recover only slightly, and only briefly, before the UE reports. Verizon configures A1 the most aggressively in 5G (88% aggressive thresholds, 89% aggressive TTTs), followed by AT&T (58%, 89%), while T-Mobile is the least aggressive (29%, 16%).

Interestingly, the resulting event volumes follow the *opposite* order. A1 is T-Mobile's single largest 5G event type, accounting for 45% of all its 5G events — more than any other event for any operator in either RAT — compared with 23% for Verizon and 12% for AT&T. Aggressiveness explains why

Verizon triggers more A1 events than AT&T, but it does not explain why the most conservative operator triggers the most events. This is largely due to T-Mobile's considerably stronger 5G RSRP distribution compared with the other two operators (Fig. 6). As a result, even relatively conservative A1 thresholds are satisfied frequently. Because A1 is primarily used to *cancel* pending HOs rather than trigger them, these reports contribute signaling cost without directly leading to HOs, and T-Mobile's 5G A1 conversion rate is correspondingly below 1%.

A6. A6 compares a neighbor against a *secondary* cell (SCell), so like A3, a *small* offset and a short TTT make it more aggressive. For A6, however, the more interesting distinction is how actively an operator uses the event, and here the three operators separate clearly. Verizon uses A6 extensively — with an 8 dB offset and an aggressive 160 ms TTT — and it accounts for 27% of its 5G events. T-Mobile and AT&T progressively suppress A6 triggering instead: 71% and 93% of their configurations, respectively, pin the offset to its maximum value alongside a 5120 ms TTT, and their A6 shares fall correspondingly to 11% and 4%.

Although A6 does not directly drive primary-cell HOs, its relatively large reporting volume does not necessarily represent unnecessary signaling. A6 reports are used for secondary-cell management under CA, allowing the network to identify and reconfigure SCells as radio conditions change. Thus, the differences we observe across operators largely reflect different CA management strategies.

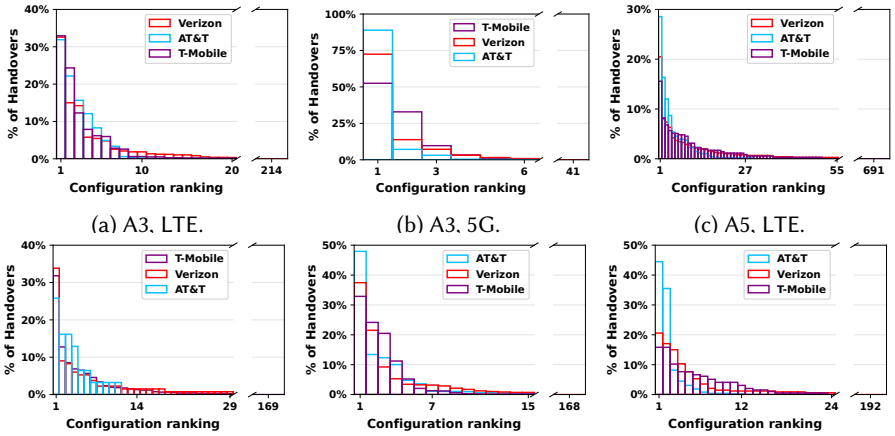
Periodic reporting. Periodic reports are not condition-based, so their aggressiveness is determined mainly by the reporting interval. While they contribute noticeably to LTE reporting, they are practically absent from 5G across all three operators. The near-elimination of periodic reporting in 5G shows that operators can substantially reduce reporting overhead by avoiding events that provide limited value for HO decisions.

Group 2: Events that actually drive HOs.

A3. A3 is triggered when a neighboring cell becomes better than the serving cell by an offset, so a *small* offset makes the configuration more aggressive. In general, A3 is configured conservatively: the majority of observed configurations use offsets of ≥5 dB (Fig. 8d). One notable exception is AT&T in LTE, which relies primarily on RSRQ and sets the RSRP offset and hysteresis to their maximum values 60% of the time, effectively suppressing RSRP-based A3 triggering. In 5G, A3 configurations are more tightly controlled than in LTE and their reports translate into HOs at correspondingly higher rates (Fig. 3).

B1. B1 is triggered when an inter-RAT neighbor exceeds an absolute threshold, so a *low* threshold and a short TTT make the configuration more aggressive. All three operators use very similar thresholds, clustered around -111 to -112 dBm (Fig. 8c), suggesting a common criterion for when a 5G cell is considered viable. The main difference across operators lies in TTT. AT&T is the most aggressive, with a median TTT of 40 ms, while Verizon is more conservative with a median of 256 ms. Verizon's aggregate median, however, masks a clear band-dependent strategy: very short TTTs for mmWave, moderate TTTs for 5G mid-band, and longer TTTs for 5G low-band. We examine this strategy, and how TTT differences across operators affect ping-pong behavior, in §6.1.3.

A2. A2 is triggered when the *serving* cell drops below a threshold, so a *high* threshold makes the configuration more aggressive: the serving link needs to degrade only slightly before the UE reports the A2 event. In LTE, all three operators generally use conservative A2 thresholds, which naturally limits triggering; Verizon is the main exception in practice because it pairs conservative thresholds with aggressive RSRQ-based TTTs. In 5G, all three operators become less conservative — the conservative share drops from 68-99% in LTE to 35-70% — allowing the UE to report A2 earlier as the 5G link degrades. This is particularly relevant for Verizon and AT&T, whose 5G coverage is slightly lower than 50% (Fig. 17). The main difference across operators again lies in TTT: AT&T



(a) A3, LTE. (b) A3, 5G. (c) A5, LTE.

(d) A5, 5G. (e) B1. (f) A2.

Fig. 9. Concentration of HOs across different parameter configurations.

uses aggressive TTTs 88% of the time, consistent with its larger A2 event share, while T-Mobile remains the most conservative, consistent with keeping UEs on its stronger 5G network for longer.

A5. A5 requires the serving cell to fall below one threshold *and* a neighboring cell to rise above another, so it becomes more aggressive when the serving-cell threshold is *high* or the neighbor-cell threshold is *low*. Across operators, the neighbor threshold remains relatively moderate, making the serving-cell threshold the main tuning lever. In 5G, T-Mobile and Verizon frequently set this threshold to -31 dBm (42% of Verizon's configurations and 16% of T-Mobile's), a value that is almost always satisfied. This effectively removes the serving-cell degradation condition and makes A5 behave like an A4 event, eliminating the additional safeguard that distinguishes A5 from A4 (§2).

The deployment-wide configuration shares, however, do not directly reflect how often these settings affect actual HOs. Among configurations that *trigger* HOs (§6.1.2), the fraction of A4-equivalent A5 configurations is 43% for Verizon and 25% for AT&T, but only 9% for T-Mobile. This indicates that many of T-Mobile's -31 dBm configurations are deployed in locations where the corresponding A5 conditions are rarely satisfied.

Summary. Overall, these results highlight three broader observations. First, reporting volume is shaped jointly by configuration choices and deployment conditions, rather than by aggressiveness alone. Second, operators adapt the same event differently across LTE and 5G; for example, more aggressive A2 configurations in 5G are consistent with the need for timely fallback where 5G coverage remains incomplete. Third, some parameter choices, such as overly high A5 serving-cell thresholds, can effectively change the intended behavior of an event. We examine how these configuration choices translate into HO performance in §6.

5.3 Operational Concentration of HO Configurations

We now examine what fraction of this large configuration space actually drives HOs in practice. For each event, we rank an operator's configurations (combinations of threshold/offset, hysteresis, and TTT) by the number of HOs they trigger and show each configuration's share of that operator's HOs (Fig. 9).

Across all events and operators, the distribution is steeply concentrated: a small set of configurations accounts for the majority of HOs, while the remaining ones form a long tail. For A3 in LTE (Fig. 9a), 5–14 configurations cover 90% of an operator's HOs, even though the total number of distinct configurations is 34/64/214 for AT&T/T-Mobile/Verizon. The concentration is even tighter in 5G: the number of distinct A3 configurations falls to 24–41 (Fig. 9b), and only 2–3 of them

account for 90% of the HOs. AT&T's 5G A3 is the extreme case, with a single configuration driving 89% of its A3 HOs. A5 sharpens this contrast further. In LTE, A5 configurations form the longest tail with 267/691/569 distinct combinations for AT&T/T-Mobile/Verizon (Fig. 9c), yet only 11–34 of them contribute 90% of the HOs. As with A3, the A5 5G tail collapses relative to LTE, with T-Mobile having the most at 169 combinations (Fig. 9d) and only 8–18 configurations contributing 90% of the HOs. B1 (Fig. 9e) is similarly steep, with its top five configurations covering 77–94% of LTE→5G HOs. The 5G A2 (Fig. 9f) is the most operator-dependent: AT&T's HOs collapse onto just four configurations (95% within its top five), while T-Mobile and Verizon spread theirs across 122 and 192 combinations, needing 12–15 to reach 90%. This is consistent with AT&T's highly standardized A2 TTT — 88% of its demotions use TTT = 160 ms (§5.2) — leaving little room for configuration diversity to matter in practice.

This concentration has an important implication: most of the configuration diversity is operationally irrelevant. Although operators provision a large parameter space, real HO traffic flows through a tiny dominant subset, so tuning effort can focus on the few combinations that actually drive HOs.

6 Impact on Performance

The 5G HO configurations we analyze are defined in terms of RSRP thresholds and offsets — the levers operators control are signal strength conditions. The central question of this section is therefore whether the configurations operators have chosen actually achieve their intended goal: improving signal strength post-HO, limiting signal degradation pre-HO, and avoiding wasteful ping-pong HOs. We analyze the four major HO types in turn: A3 and A5 for intra-RAT 5G→5G HOs, B1 for LTE→5G promotions, and A2 for 5G→LTE demotions.

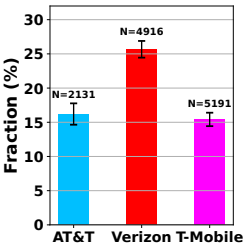
Following §5.3, we analyze only HO-triggering configurations, and among those, the dominant ones that together account for at least 90% of each event's HOs; these may differ from the most frequently provisioned configurations, a distinction that matters, as we show for A5 below.

6.1 Impact on signal strength and ping-pong HOs

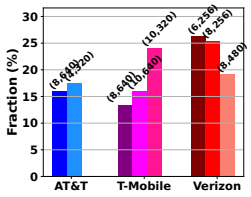
We use three metrics to evaluate configuration effectiveness: $\Delta RSRP = RSRP_{post} - RSRP_{pre}$ quantifies post-HO signal improvement; $RSRP_{pre}$ quantifies how much signal quality has degraded before the HO executes; ping-pong rate (the fraction of HOs where the UE returns to the previous cell within a short amount of time (15 s in our analysis)) quantifies wasteful HO decisions. Together, these metrics capture the fundamental three-way tradeoff operators face: a larger offset or threshold produces better post-HO signal at the cost of worse pre-HO signal, while a longer TTT reduces ping-pong risk but delays the HO decision and worsens pre-HO signal. Optimizing all three simultaneously is extremely difficult, as we show below.

Note that the 15 s ping-pong window is an observation interval chosen for our analysis, not a configured parameter. The TTTs of the dominant configurations we analyze are at most 1.28 s, and even the largest TTT allowed by 3GPP (5.12 s) is well below 15 s, leaving ample time for a return HO to be triggered and executed within the window. At the same time, 15 s is short relative to the time a vehicle typically spends traversing a cell at highway speed.

One important subtlety is worth noting before proceeding: unlike $\Delta RSRP$, which is a relative metric, $RSRP_{pre}$ is an absolute metric that



(a) Across operators.

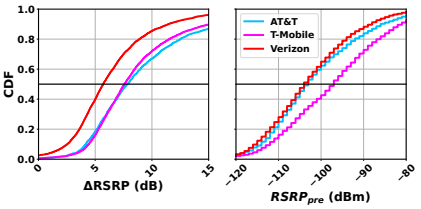


(b) Across A3 configurations.

Fig. 10. Fraction of A3 ping-pong HOs.

AT&T	T-Mobile	Verizon
(8, 640): 89% (4, 320): 7% Count: 2	(8, 640): 53% (10, 640): 33% (10, 320): 10% Count: 3	(6, 256): 72% (8, 256): 14% (8, 480): 7% Count:3

(a) Dominant combinations ($\Delta_{A_3} + H_{A_3}, TTT$).



(b) Performance.

Fig. 11. Dominant parameter combinations and RSRP analysis for the A3 event.

depends not only on HO configuration choices, but also on the operator's overall deployment density and transmit power. An operator with a denser deployment may achieve higher $RSRP_{pre}$ than a competitor using a more aggressive configuration simply because its cells are closer together. This coverage confound means that counter-intuitive $RSRP_{pre}$ results across operators should not be interpreted as configuration failures without first checking whether deployment density explains them.

6.1.1 A3 HOs. An A3 event is triggered when the RSRP of a neighboring cell remains $\Delta_{A_3} + H_{A_3}$ dB higher than the RSRP of the serving cell for TTT ms. Fig. 11a shows the dominant configurations — tuples of $((\Delta_{A_3} + H_{A_3}, TTT))$ — for all three operators. AT&T has 2 dominant configurations responsible for 96% of its HOs, while T-Mobile and Verizon each have 3, responsible for 96% and 93%, respectively. The configurations differ across operators.

RSRP analysis. Regardless of the A3 parameter choices, 98% or more of A3 HOs result in a positive $\Delta RSRP$ (Fig. 11b). This is a direct consequence of the A3 trigger condition, which requires the neighbor's RSRP to exceed the serving cell's by an offset before an HO is triggered. However, the magnitude of improvement and the degree of pre-HO degradation vary significantly across configurations.

AT&T primarily uses one configuration (8, 640), while T-Mobile uses the same configuration 53% of the time and another configuration with higher $\Delta_{A_3} + H_{A_3}$ (10, 640) 33% of the time. In contrast, Verizon uses two configurations with lower TTT and/or offset 86% of the time ((6, 256) and (8, 256)). As a result, AT&T and T-Mobile achieve similar $\Delta RSRP$, outperforming Verizon by 2 dB at the median. The $RSRP_{pre}$ results, however, are counter-intuitive at first glance: T-Mobile has much higher $RSRP_{pre}$ than AT&T despite using the same primary configuration, and Verizon has the lowest $RSRP_{pre}$ despite using the most aggressive configurations, which should in principle trigger HOs earlier and thus, at higher signal levels.

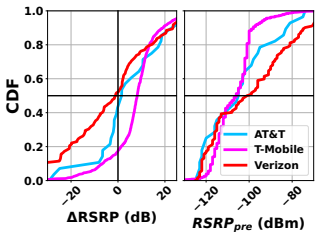
The cross-operator comparison is explained by coverage differences rather than configuration choices. T-Mobile has the highest 5G RSRP and Verizon the lowest among the three operators (Fig. 6), which directly accounts for T-Mobile's higher $RSRP_{pre}$ compared to AT&T in spite of using the same configuration, and Verizon's lower $RSRP_{pre}$ in spite of its more aggressive configurations.

Ping-pong analysis. Verizon experiences the highest A3 ping-pong rate (25%) compared to 15–16% for AT&T and T-Mobile (Fig. 10a), directly attributable to its low TTT values of 256 ms vs. 640 ms for the other two operators. Within each operator, longer TTT consistently produces lower ping-pong rates for T-Mobile and Verizon (Fig. 10b). The exception is AT&T, where the very different configurations (8, 640) and (4, 320) produce similar ping-pong rates, suggesting that local deployment density can override the TTT effect when cell boundaries are sparse and unstable. TTT is thus the dominant, but not universal lever for ping-pong control — a finding we confirm across A5 and B1 below.

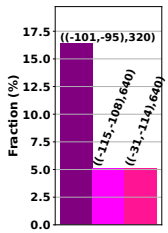
6.1.2 A5 HOs. An A5 event requires two conditions to hold simultaneously for TTT ms: the serving cell's RSRP drops below $\Theta^S_{A_5} - H_{A_5}$ and a neighboring cell's RSRP rises above $\Theta^N_{A_5} + H_{A_5}$. The use of two independent thresholds rather than a single relative offset gives A5 finer-grained control

AT&T	T-Mobile	Verizon
((-31, -102), 640): 25%	((-101, -95), 320): 32%	((-31, -103), 640): 33%
((-104, -100), 320) : 16%	((-115, -108), 640): 13%	((-42, -118), 640): 10%
((-82, -110), 640): 16%	((-31, -114), 640): 9%	((-122, -111), 128): 8%
...	...	...
((-122, -118), 160): 3%	((-108, -108), 320): 2%	((-115, -111), 640): 1%
Count: 8	Count: 14	Count: 18

(a) Dominant combinations $((\Theta^S_{A_5} - H_{A_5}, \Theta^N_{A_5} + H_{A_5}), TTT)$.



(b) Performance.



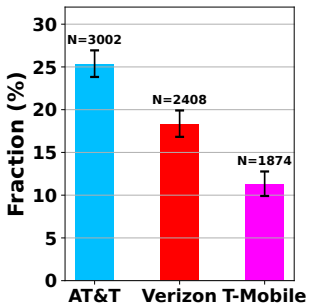
(c) Ping-pong HOs.

Fig. 12. Dominant parameter combinations and performance analysis for the A5 event.

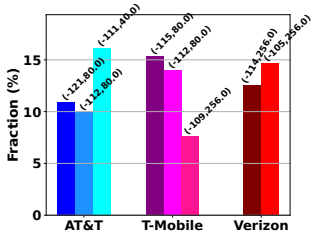
than A3; operators can independently set how bad the serving cell must be and how good the neighbor must be before an HO is triggered. As we show below, this advantage is real, but requires careful configuration to realize. Fig. 12a shows that there are 8/14/18 dominant configurations for AT&T/T-Mobile/Verizon — substantially more than for A3 — with no single configuration exceeding a 33% share. The configurations differ again across operators.

A critical misconfiguration. The primary A5 configurations for both AT&T and Verizon are functionally equivalent to A4 events: the serving cell threshold $\Theta^S_{A_5} - H_{A_5}$ is set to -31 or -42 dBm, a level that is almost always satisfied in practice, meaning the HO is triggered whenever a neighbor satisfies the second condition regardless of the serving cell's signal quality. AT&T uses this A4-equivalent configuration 25% of the time and Verizon 43% of the time, compared to only 9% for T-Mobile. This is a misconfiguration: it negates the key advantage of A5 over A3, which is the ability to condition the HO on actual serving cell degradation rather than just neighbor strength.

RSRP analysis. The performance consequences of the A4-equivalent misconfiguration are clearly visible in Fig. 12a. Both AT&T and Verizon have negative $\Delta RSRP$ approximately 50% of the time, with Verizon achieving the worst outcomes. However, there is a compensating effect on $RSRP_{pre}$: by not waiting for the serving cell to degrade, these configurations ensure the HO is often completed while signal quality is still relatively high. Verizon achieves the highest $RSRP_{pre}$ among the three operators 60% of the time, and AT&T outperforms T-Mobile on this metric 35% of the time. This is the characteristic tradeoff of A4-equivalent A5 configurations; they sacrifice post-HO signal quality in exchange for executing the HO while the serving cell is still strong.



(a) Across operators.



(b) Across B1 configurations.

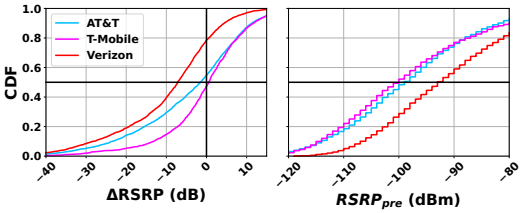
Fig. 13. % of LTE-5G ping-pong HOs.

T-Mobile, which uses properly configured A5 events for the vast majority of the time, achieves the highest $\Delta RSRP$ (positive 80% of the time). It also achieves the highest $RSRP_{pre}$ 40% of the time, primarily due to the aggressive TTT in its primary configuration (320 ms vs. 640 ms for the other operators), but lags behind Verizon 60% of the time and behind AT&T 35% of the time.

Ping-pong analysis. Fig. 12c plots the fraction of A5 ping-pong HOs for the top-3 configurations of T-Mobile. Consistent with the A3 findings, TTT is the dominant lever: ((-101, -95), 320) produces a 16.3% ping-pong rate compared to 5% for configurations with TTT = 640 ms (Fig. 12c), despite having a higher neighbor threshold that should in principle ensure stronger post-HO signal. The higher threshold is insufficient to compensate for the low measurement confidence of a short TTT.

AT&T	T-Mobile	Verizon
(-112, 80): 48%	(-112, 80): 33%	(-114, 256): 37%
(-111, 40): 13%	(-115, 80): 24%	(-116, 256): 22%
(-121, 80): 12%	(-109, 256): 21%	(-105, 256): 9%
...	...	...
(-117, 80): 4%	(-103, 640): 5%	(-111, 1280): 2%
Count: 6	Count: 5	Count: 10

(a) Dominant combinations ($\Theta_{B_1} + H_{B_1}, TTT$)



(b) RSRP analysis for B1 event.

Fig. 14. Dominant parameter combinations and performance analysis for B1 event.

6.1.3 B1 HOs. B1 events are used to promote UEs from LTE to 5G. A report is sent when the RSRP of an inter-RAT 5G neighboring cell exceeds $\Theta_{B_1} + H_{B_1}$ for TTT ms. Notably, B1 does not condition on the serving LTE cell's signal quality; promotion to 5G is desired as long as the 5G cell is sufficiently strong, regardless of how good the LTE connection currently is.

Fig. 14a shows the dominant configurations. All three operators use similar $\Theta_{B_1} + H_{B_1}$ values but very different TTT values, reflecting different promotion strategies. AT&T uses the shortest TTT (40–80 ms), aggressively promoting UEs to 5G at the first opportunity. T-Mobile uses a mix of short and long TTT values (80, 256, and 640 ms). For Verizon, we further break down the TTT configurations based on the carrier frequency of the target 5G cell. Verizon uses a deliberately band-differentiated strategy: very short TTT (128 ms) for mmWave, moderate TTT (160–512 ms) for 5G mid, and long TTT (1024–1280 ms) for 5G low. This reflects a rational policy — mmWave cells have limited coverage and must be seized quickly when available, while 5G low competes directly with Verizon's strong LTE signal, justifying conservatism.

RSRP analysis. Since B1 ignores the serving cell, RSRP improvement is not guaranteed. About 50% of B1-triggered HOs for AT&T and T-Mobile and 80% for Verizon result in a signal strength reduction (Fig. 14b). Verizon's high rate of signal-reducing B1 HOs is explained by its network characteristics: it has the highest LTE RSRP but the lowest 5G RSRP among the three operators (Fig. 6), meaning promotion to 5G almost always means accepting a weaker signal. Its long TTT for 5G low reflects awareness of this tradeoff. T-Mobile has the opposite profile — highest 5G RSRP, lowest LTE RSRP — and achieves the best $\Delta RSRP$ despite using more aggressive TTT values than Verizon, again illustrating that coverage context, not configuration alone, determines outcomes.

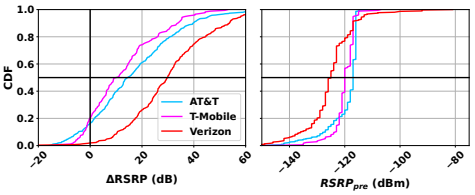
Ping-pong analysis. We found the Verizon configuration (-116, 256) did not appear at all in 2024. Since the 2022 dataset contained a much smaller number of LTE-5G HOs and the 2023 dataset was collected only on the East Coast, our analysis is based on the 2024 dataset only for a fair comparison across configurations. T-Mobile has the lowest B1 ping-pong rate (10%), reflecting its strong 5G coverage (Fig. 13a). AT&T's aggressive short TTT produces the highest rate (18%). Verizon, despite using longer TTT than T-Mobile, experiences nearly as many ping-pong HOs as AT&T, because its weak 5G signal means that even carefully timed promotions frequently fail to hold. Within operators, longer TTT consistently reduces ping-pong rates for both AT&T and T-Mobile (Fig. 13b). Between configurations with the same TTT, a higher $\Theta_{B_1} + H_{B_1}$ reduces ping-pong only for AT&T, confirming again that TTT is the primary lever and threshold effects are secondary.

6.1.4 A2 HOs. A2 events are the primary mechanism for 5G NSA→LTE demotions. A report is sent when the serving 5G cell's RSRP remains below $\Theta_{A_2} - H_{A_2}$ for TTT ms. Since A2 does not consider neighboring cell signal quality, RSRP improvement after the demotion is not guaranteed, but in practice all three operators use conservative enough thresholds (Fig. 15a) ensuring that $\Delta RSRP$ is positive 75–95% of the time (Fig. 15b).

Fig. 15a shows the dominant configurations. All three operators use similar, mostly conservative $\Theta_{A_2} - H_{A_2}$ values but differ substantially in TTT. AT&T uses TTT = 160 ms for 88% of its demotions —

	AT&T	T-Mobile	Verizon
5G → LTE	(-117, 160): 44% (-115, 160): 36% (-125, 160): 8% (-116, 320): 4% Count: 4	(-118, 640): 16% (-121, 160): 16% (-118, 160): 10% . . (-82, 160): 3% Count: 12	(-120, 128): 21% (-124, 128): 17% (-126, 40): 15% . . (-96, 128): 1% Count: 15

(a) Dominant combinations ($\Theta_{A_2} - H_{A_2}, TTT$)



(b) RSRP analysis for A2 event.

Fig. 15. Dominant parameter combinations and performance analysis for A2 event.

a standardized, moderately aggressive approach. T-Mobile distributes TTT broadly across 160–640 ms (86% of the time), with a small fraction using 40 ms (4%) — the most conservative overall, consistent with its goal of keeping UEs on 5G as long as possible given its strong 5G signal. Verizon uses the shortest TTT values: 128 ms alone accounts for 69% of demotions and TTT ≤ 160 ms covers 88% — the most aggressive strategy among the three operators.

These strategies make clear sense given each operator's network characteristics. Verizon has the lowest 5G RSRP and the strongest LTE RSRP, so it does not want to keep UEs on 5G when signal degrades, hence the short TTT and quick demotion. Fig. 15b shows indeed that Verizon has the lowest $RSRP_{pre}$ under 5G but switching to LTE results in the highest $\Delta RSRP$. T-Mobile has the opposite profile (highest 5G RSRP, lowest LTE RSRP) so it holds on to 5G connections as long as possible, accepting lower $\Delta RSRP$ as the cost of maximizing 5G time. AT&T sits between the two. Among the four event types analyzed in this section, A2 configurations are the most clearly well-matched to each operator's deployment context — a useful contrast to the misconfigured A5 events and aggressive B1 TTTs identified above. We omit ping-pong analysis for A2 since rapid return to 5G after a demotion is governed by B1 configurations, analyzed in §6.1.3.

6.1.5 Summary. In summary, the three metrics cannot be optimized simultaneously, and two of them are strongly influenced by deployment density. TTT is consistently the most reliable lever for ping-pong control, though sparse deployments can override it, and counter-intuitive outcomes are frequently explained by deployment density rather than configuration quality, so operators cannot evaluate configurations from aggregate metrics alone.

6.2 RSRP vs. throughput improvement

While the 5G HO configurations we analyze are defined in terms of RSRP, a natural question is whether the signal strength improvements we analyzed in §6.1 translate to throughput improvements for users. We briefly examine this relationship to characterize the limits of RSRP-based optimization. Since the goal here is to isolate the impact of HO configurations on end-user throughput, we restrict this analysis to the backlogged, saturated tests in our dataset (§3); non-backlogged traffic is rate-limited by the application rather than by the radio link, and would therefore mask the very effect we are trying to measure. We calculate the average throughput 1 s before (T_{pre}) and 1 s after (T_{post}) an HO and then define $\Delta T = T_{post} - T_{pre}$; results are robust across intervals from 0.5 s to 5 s (Fig. 18 in §D).

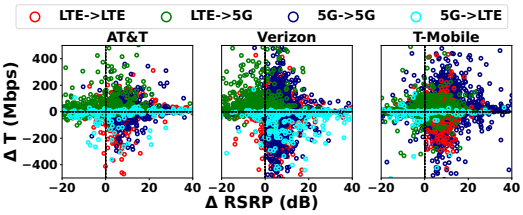


Fig. 16. RSRP and throughput correlation.

Fig. 16 shows scatterplots of ΔT against $\Delta RSRP$ for all three operators and HO types. There is no meaningful correlation between signal strength improvement and throughput improvement across any operator or HO type. RSRP therefore remains the appropriate metric for judging whether the RSRP-based configurations of §6.1 behave as intended, but it is not a proxy for

user-perceived performance, which also depends on factors outside the HO framework, such as bandwidth, scheduling, and congestion. We return to this limitation in §7.

Three additional observations are worth noting. First, for intra-RAT LTE→LTE and 5G→5G HOs (mostly triggered by A3 events), the majority improve RSRP but ΔT is negative for 40–50% of cases across all operators, indicating that signal-improving HOs frequently harm user experience regardless of how well the RSRP outcome is optimized. Second, LTE→5G HOs result in signal strength reduction 45–75% of the time across operators, yet ΔT is negative for only 20–30%; aggressive 5G promotion often benefits users even when it worsens signal strength, due to the much higher bandwidth of 5G mid and high bands. Third, and most strikingly, Verizon's 5G→LTE demotions improve RSRP 85% of the time yet result in throughput degradation 80% of the time, with losses exceeding 150 Mbps in 18% of cases. AT&T and T-Mobile achieve a substantially better balance between RSRP and throughput for the same HO type. This suggests that Verizon's aggressive A2 demotion strategy, while rational from a signal strength perspective, comes at a significant throughput cost that RSRP-based optimization cannot capture.

7 Implications for Operators

For operators. The most immediate opportunity is reducing wasteful signaling from A1 and A6, which together dominate reported events while almost never leading to HOs. Operators already eliminated periodic events from 5G for exactly this reason; the same logic applies to A1 and A6. For A1, conservative configurations alone cannot prevent excessive triggering in networks with strong 5G coverage; operators should reconsider whether the current event triggering policy is necessary at current 5G deployment densities.

Operators should also audit their A5 configurations. Verizon and AT&T configure A5 in a way that makes it functionally equivalent to a simpler event in 43% and 25% of HO-triggering configurations respectively, negating its key design advantage. The performance consequences are clear: misconfigured A5 events produce negative post-HO RSRP approximately half the time, while properly configured ones achieve improvement the vast majority of the time. T-Mobile's HO-triggering A5 configurations provide a concrete reference point.

TTT is the most reliable lever for ping-pong control, but cannot be evaluated from aggregate metrics alone; deployment density confounds the relationship and must be controlled for. Finally, operators relying solely on RSRP to evaluate HO performance will miss throughput consequences entirely (§6.2).

For the research community. The coverage confound we document, where deployment density explains counter-intuitive performance outcomes across operators and configurations, poses a fundamental challenge for automated HO tuning systems (e.g., [10]). Ground-truth characterization of this confound, as we provide here, is a necessary input for any tuning system to define meaningful optimization targets.

The weak RSRP-throughput correlation suggests the 3GPP HO framework is optimizing the wrong objective from the user's perspective, particularly for inter-RAT HOs where bandwidth differences dominate. Designing HO policies that account for technology-specific characteristics is an open and practically important problem. The operational concentration finding further suggests that data-driven optimization need not search the full parameter space; the dominant configurations we identify provide a compact, empirically grounded starting point.

Finally, while quantifying the precise cost of signaling overhead in terms of UE energy and network resources is beyond our scope, the rates we observe represent continuous unnecessary radio activity whose energy cost is well-established [14, 15].

8 Related Work

HOs in the pre-5G era. Prior work studied LTE HOs in operational networks [7, 20–23, 31–33, 35, 36], examining their performance impact [20, 21, 31, 33, 35, 36], mitigation solutions [31, 35, 36], and instability or unreachability issues [22, 23, 32]. Closest to ours is [7], the first large-scale characterization of LTE HO configurations; we provide the first such study in the 5G era.

HOs in the 5G era. Many experimental studies have evaluated 5G HOs [8, 11–14, 16–19, 24–26, 29, 30, 34, 37], most at small scale across a few km or 1–3 cities [8, 18, 19, 24–26, 29, 30, 34, 37]. The only large-scale 5G HO studies [11–14, 17] characterize what happens during HOs: frequency, duration, and impact on transport and application performance from the user's perspective [11–14], or geo-temporal dynamics and failure causes over a 4-week dataset from an operator's perspective [17] — but none examine the configuration parameters that cause HOs. We are the first to study these configuration policies at scale, and our 27-month longitudinal dataset lets us track their evolution over time.

Dependent misconfigurations and configuration diversity. Zhang et al. [37] identify dependent misconfigurations in 5G/4.5G RRC, where certain parameter combinations violate inter-dependency constraints in the 3GPP specification. We are orthogonal: rather than checking specification compliance, we ask whether operators' freely chosen values achieve good HO outcomes in practice — often they do not, most notably the A4-equivalent A5 misconfiguration deployed by AT&T and Verizon. A separate line of work [9, 10, 27, 28] proposes data-driven and learning-based approaches for automated cellular configuration selection; most relevant is Iridescence [10], which tunes 5G NSA configurations in the presence of deployment confounders — precisely the coverage confound we identify as a systematic obstacle to evaluating HO configurations. These operator-side tools address *how* to tune configurations automatically across the full configuration space; we address the prior question of *what* is misconfigured and *why*, providing the ground-truth characterization such tuning systems need to define their optimization targets. The contributions are complementary.

9 Conclusion

We conducted the first large-scale, longitudinal study of HO configurations in the 5G era, collecting a 27-month dataset of 48K HOs and 239K reporting events across all three major US operators. Our analysis reveals that current configurations are often overly aggressive and inefficient, with overhead concentrated in a few event types operators can and should revisit; despite large provisioned configuration spaces, actual HO behavior concentrates in a tiny dominant subset, tighter in 5G than LTE. A pervasive coverage confound means counter-intuitive performance outcomes are frequently explained by deployment density rather than configuration quality. Finally, TTT is the most reliable lever for ping-pong control, a critical A5 misconfiguration deployed by two operators negates a key design advantage, and Verizon's aggressive 5G→LTE demotion incurs a significant throughput cost invisible to RSRP-based optimization. We expect our findings and released dataset to help operators reduce overhead and guide the research community toward more effective, adaptive HO policies for next-generation networks.

Acknowledgments

We thank the anonymous reviewers and our shepherd for their helpful comments. This work was supported in part by NSF grants 2340283, 2312834 and 2415216.

References

[1] Online. [CoNEXT '26]Too Many Reports, Too Few Handovers: A Large-Scale Study of Configuration Inefficiencies in Operational 5G Networks. (Online). https://github.com/NUWiNS/conext26-handover-paper-dataset.

[2] 3GPP. [n. d.]. 3GPP Standard Specification. ([n. d.]). http://www.3gpp.org.

[3] 3GPP. 2012. *Evolved Universal Terrestrial Radio Access (E-UTRA); Mobility enhancements in heterogeneous networks.* Technical Report TR 36.839. 3rd Generation Partnership Project (3GPP). https://www.3gpp.org/ftp/Specs/html-info/36839.htm

[4] 3GPP. 2015. TS36.331: Radio Resource Control (RRC). (Mar. 2015). http://www.3gpp.org/ftp/Specs/html-info/36331.htm

[5] 3GPP. 2021. TS38.331: 5G NR: Radio Resource Control (RRC). (Mar. 2021). http://www.3gpp.org/ftp/Specs/html-info/38331.htm

[6] accuverxcalsolo [n. d.]. XCAL Solo. ([n. d.]). https://accuver.com/sub/products/view.php?idx=11.

[7] Haotian Deng, Chunyi Peng, Ans Fida, Jiayi Meng, and Y. Charlie Hu. 2018. Mobility Support in Cellular Networks: A Measurement Study on Its Configurations and Implications. In *Proc. of ACM IMC*.

[8] Claudio Fiandrino, David Juárez Martínez-Villanueva, and Joerg Widmer. 2022. Uncovering 5G Performance on Public Transit Systems with an App-based Measurement Study. In *Proc. of ACM MSWiM*.

[9] Changhan Ge, Zihui Ge, Xuan Liu, Ajay Mahimkar, Yusef Shaqalle, Yu Xiang, and Shomik Pathak. 2023. Chroma: Learning and Using Network Contexts to Reinforce Performance Improving Configurations. In *Proc. of ACM MobiCom*.

[10] Changhan Ge, Ajay Mahimkar, Zihui Ge, Romeo Fernandez, Joseph Maniaci, Shomik Pathak, and Maulik Shah. 2025. Iridescence: Improving Configuration Tuning in the Presence of Confounders for 5G NSA Networks. In *Proc. of ACM CoNEXT*.

[11] Moinak Ghoshal, Omra 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*.

[12] 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*.

[13] Moinak Ghoshal, Imran Khan, Z. Jonny Kong, Phuc Dinh, Jiayi meng, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2023. Performance of Cellular Networks on the Wheels. In *Proc. of ACM IMC*.

[14] Ahmad Hassan, Arvind Narayanan, Anlan Zhang, Wei Ye, Ruiyang Zhu, Shuowei Jin, Jason Carpenter, Z. Morley Mao, Feng Qian, and Zhi-Li Zhang. 2022. Vivisecting Mobility Management in 5G Cellular Networks. In *Proc. of ACM SIGCOMM*.

[15] Junxian Huang, Feng Qian, Alexandre Gerber, Z. Morley Mao, Subhabrata Sen, and Oliver Spatscheck. 2012. A Close Examination of Performance and Power Characteristics of 4G LTE Networks. In *Proc. of ACM Mobisys*.

[16] Michail Kalntis, Andra Lutu, Jesús Omaña Iglesias, Fernando A. Kuipers, and George Iosifidis. 2025. Smooth Handovers via Smoothed Online Learning. In *Proc. of IEEE INFOCOM*.

[17] Michail Kalntis, José Suárez-Varela, Jesús Omaña Iglesias, Anup Kiran Bhattacharjee, George Iosifidis, Fernando A. Kuipers, and Andra Lutu. 2024. Through the Telco Lens: A Countrywide Empirical Study of Cellular Handovers. In *Proc. of ACM IMC*.

[18] Imran Khan, Tuyen X Tran, Matti Hiltunen, Theodore Karagioules, and Dimitrios Koutsonikolas. 2024. An Experimental Study of Low-Latency Video Streaming over 5G. In *Proc. of IEEE MeditCom*.

[19] Konstantinos Kousias, Mohammad Rajiullah, Giuseppe Caso, Ozgu Alay, Anna Brunstorm, Luca De Nardis, Marco Neri, Usman Ali, and Maria-Gabriella Di Benedetto. 2022. Coverage and Performance Analysis of 5G Non-Standalone Deployments. In *Proc. of ACM WiNTECH*.

[20] Li Li, Ke Xu, Tong Li, Kai Zheng, Chunyi Peng, Dan Wang, Xiangxiang Wang, Meng Shen, and Rashid Mijumbi. 2018. A measurement study on multi-path tcp with multiple cellular carriers on high speed rails. In *Proc. of ACM SIGCOMM*.

[21] Li Li, Ke Xu, Dan Wang, Chunyi Peng, Kai Zheng, Rashid Mijumbi, and Qingyang Xiao. 2017. A longitudinal measurement study of tcp performance and behavior in 3g/4g networks over high speed rails. *Transactions of IEEE/ACM ToN* 25, 4 (2017).

[22] Yuanjie Li, Haotian Deng, Jiayao Li, Chunyi Peng, and Songwu Lu. 2016. Instability in Distributed Mobility Management: Revisiting Configuration Management in 3G/4G Mobile Networks. In *Proc. of ACM SIGMETRICS*.

[23] Yuanjie Li, Jiaqi Xu, Chunyi Peng, and Songwu Lu. 2016. A first look at unstable mobility managementincellular networks. In *Proc. of ACM HotMobile*.

[24] Yanbing Liu, Jingqi Huang, Sonia Fahmy, and Chunyi Peng. 2025. An In-Depth Look into 5G ON-OFF Loops in the Wild. In *Proc. of ACM IMC*.

[25] Yanbing Liu and Chunyi Peng. 2023. A Close Look at 5G in the Wild: Unrealized Potentials and Implications. In *Proc. of IEEE INFOCOM*.

[26] Zhutian Liu, Qing Deng, Zhaowei Tan, Zhiyun Qian, Xinyu Zhang, Ananthram Swami, and Srikanth V. Krishnamurthy. 2024. M2HO: Mitigating the Adverse Effects of 5G Handovers on TCP. In *Proc. of ACM MobiCom*.

[27] Ajay Mahimkar, Zihui Ge, Xuan Liu, Yusef Shaqalle, Yu Xiang, Jennifer Yates, Shomik Pathak, and Rick Reichel. 2022. Aurora: Conformity-based Configuration Recommendation to Improve LTE/5G Service. In *Proc. of ACM IMC*.

[28] Ajay Mahimkar, Ashiwan Sivakumar, Zihui Ge, Shomik Pathak, and Karunasish Biswas. 2021. Auric: Using Data-driven Recommendation to Automatically Generate Cellular Configuration. In *Proc. of ACM SIGCOMM*.

[29] Arvind Narayanan, Eman Ramadan, Jason Carpenter, Qingxu Liu, Yu Liu, Feng Qian, and Zhi-Li Zhang. 2020. A First Look at Commercial 5G Performance on Smartphones. In *Proc. of ACM WWW*.

[30] Arvind Narayanan, Xumiao Zhang, Ruiyang Zhu, Ahmad Hassan, Shuowei Jin, Xiao Zhu, Xiaoxuan Zhang, Denis Rybkin, Zhengxuan Yang, Zhuoqing Morley Mao, et al. 2021. A Variegated Look at 5G in the Wild: Performance, Power, and QoE Implications. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. 610–625.

[31] Yunzhe Ni, Feng Qian, Taide Liu, Yihua Cheng, Zhiyao Ma, Jing Wang, Zhongfeng Wang, Gang Huang, Xuanzhe Liu, and Chenren Xu. 2023. POLYCORN: Data-driven Cross-layer Multipath Networking for High-speed Railway through Composable Schedulerlets. In *Proc. of USENIX NSDI*.

[32] Chunyi Peng and Yuanjie Li. 2016. Demystify Undesired Handoff in Cellular Networks. In *Proc. of IEEE ICCCN*.

[33] Jing Wang, Yufan Zheng, Yunzhe Ni, Chenren Xu, Feng Qian, Wangyang Li, Wantong Jiang, Yihua Cheng, Zhuo Cheng, Yuanjie Li, et al. 2019. An active-passive measurement study of tcp performance over lte on high-speed rails. In *Proc. of ACM MobiCom*.

[34] Dongzhu Xu, Anfu Zhou, Xinyu Zhang, Guixian Wang, Xi Liu, Congkai An, Yiming Shi, Liang Liu, and Huadong Ma. 2020. Understanding Operational 5G: A First Measurement Study on Its Coverage, Performance and Energy Consumption. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication*. Association for Computing Machinery.

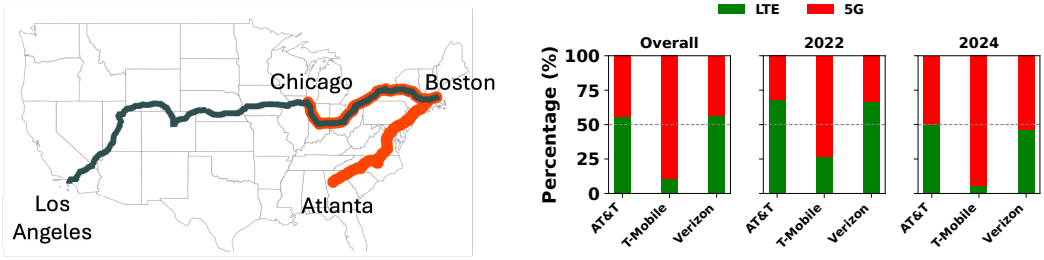
[35] Shichang Xu, Ashkan Nikravesh, and Z Morley Mao. 2019. Leveraging context-triggered measurements to characterize LTE handover performance. In *Proc. of PAM*.

[36] Zhehui Zhang, Yuanjie Li, Qianru Li, Jinghao Zhao, Ghufran Baig, Lili Qiu, and Songwu Lu. 2023. Movement-Based Reliable Mobility Management for Beyond 5G Cellular Networks. *Transactions of IEEE/ACM ToN* 31, 1 (2023).

[37] Zhehui Zhang, Yanbing Liu, Qianru Li, Zizheng Liu, Chunyi Peng, and Songwu Lu. 2023. Dependent Misconfigurations in 5G/4.5G Radio Resource Control. In *Proc. of ACM CoNEXT*.

Table 4. Description of common event parameters.

Parameter Name	**Description**
Hysteresis H	Prevents unnecessary and fast HOs. Sets a threshold below which the serving cell is considered better than a neighboring cell, even if the neighboring cell's signal strength is slightly stronger.
Threshold Θ	Defines the level at which a measurement is triggered. When the measured value crosses the threshold, a measurement report is generated.
Offset Δ	Added to the signal strength of a serving or a neighboring cell. Creates a margin in the comparison between the serving cell and the neighboring cell.
Time To Trigger TTT	Time during which specific criteria for the event need to be met in order to trigger a measurement report.
Trigger Quantity	Defines the measured parameter (RSRP or RSRQ) used to trigger an event.



(a) Measurement routes.

(b) Overall coverage.

Fig. 17. Measurement route and cellular coverage.

Appendix

A Ethics

This work does not raise any ethical concerns. The study was carried out by PhD students and faculty. We purchased multiple unlimited cellular data plans from one of the three major US carriers and our experiments comply with their customer agreements.

B Handover events and Parameters

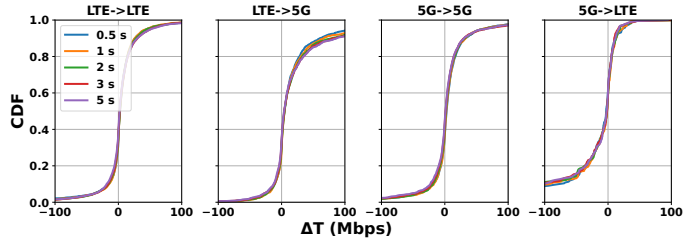
Table 4 describes LTE/5G-NR HO-related important event parameters.

C Route and Coverage

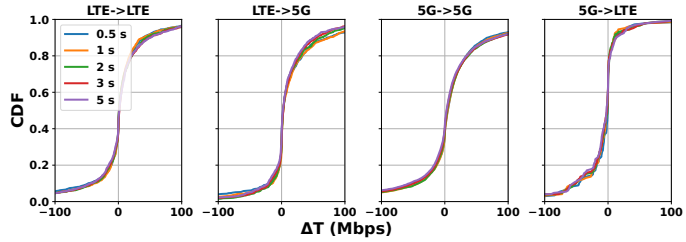
Fig. 17a shows the routes of the four driving trips. Fig. 17b shows the overall cellular coverage landscape for all three operators in our dataset. Coverage for a particular technology (LTE or 5G) is defined as the fraction of the number of XCAL logs (at a granularity of 1 s) for that technology over the total number of logs.

D Pre and post throughput difference at varying time granularity

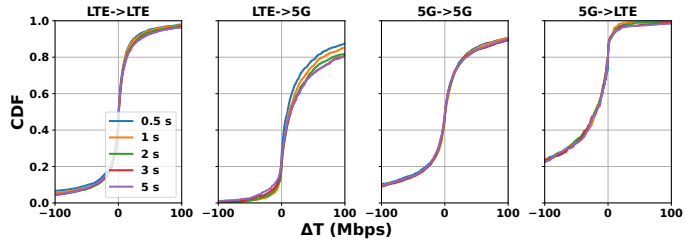
Fig. 18 shows ΔT across different time intervals ranging from 0.5 seconds to 5 seconds for all three operators and all HO types.



(a) AT&T.



(b) T-Mobile.



(c) Verizon.

Fig. 18. Pre and post throughput difference at varying time granularity.

Received 5 June 2026; accepted 10 September 2026