A Close Look at 5G in the Wild: Unrealized Potentials and Implications

Yanbing Liu Chunyi Peng
Department of Computer Science, Purdue University, West Lafayette, IN, USA
{liu3098, chunyi}@purdue.edu

***Abstract*—This paper reports our in-depth measurement study of 5G experience with three US operators (AT&T, Verizon and T-Mobile). We not only quantitively characterize 5G coverage, availability and performance (over both mmWave and Sub-6GHz bands), but also identify several performance issues and analyze their root causes. We see that real 5G experience is not that satisfactory as anticipated. It is mainly because faster 5G is not used as it can and should. We have several surprising findings: Despite huge speed potentials (say, up to several hundreds of Mbps), more than half are not realized in practice; Such under-utilization is mainly stemmed from current practice and policies that manage radio resource in a performance-oblivious manner; 5G is even less used where 5G is co-deployed over both mmWave and Sub-6GHz bands; Transiently missing 5G is not uncommon and its negative impacts last much longer. Inspired by our findings, we design a patch solution called `5GBoost` to fix the problems identified in legacy 5G operations. Our preliminary evaluation validates its effectiveness to realize more 5G potentials.**

I. Introduction

5G is rolling out rapidly across the globe, particularly in the US [1], [2]. Since the first rollout in April 2019, 5G in the US had served 10 million connections by 2020 and is projected to reach 100 million connections in 2022 and account for 63% of total mobile connections by 2025 [2]. Globally, 5G will reach 1.3 billion connections by the end of 2022 and serve a quarter of global mobile connections by 2025 [1].

5G is making big promises to be much faster (say, up to 20 Gbps peak data rates and 100+ Mbps on average [3]). To this end, network operators have been constantly investing in radio frequency (RF) spectrum resources and upgrading radio access technologies (RATs) to enhance network capacity and speed potentials. Network operators not only acquire new spectrum bands for 5G (say, 5GM over mmWave bands and 5GS over sub-6GHz bands)[1], but also repurpose the existing bands for advanced RATs (say, retiring 2G/3G bands for 4G and sharing 4G bands with 5GS). They also empower carrier aggregation (CA) to exploit increasing spectrum resources to serve a single device [7]. Instead of a single serving cell before, CA allows more than one cells to simultaneously serve the same device. As such, CA aggregates all spectrum blocks (each used by one serving cell, called a component carrier, or a RF channel) to offer radio access over wider spectrum. It flexibly increases the total spectrum resources assigned to the same device, thus promising to greatly boost data speed.

[1]In this paper, we use 5GM and 5GS to represent 5G over two frequency ranges: mmWave bands (>24 GHz) and Sub-6GHz bands (<6 GHz). They are called 5G+ & 5G by AT&T [4], 5G ultra wideband & 5G nationwide by Verizon [5], 5G ultra capacity & extended range by T-Mobile [6].

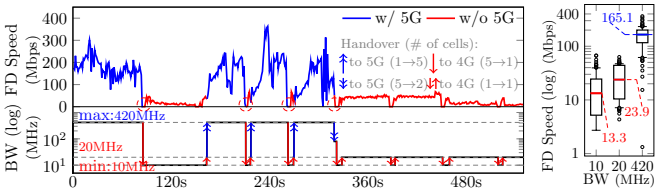

Fig. 1: A real-world instance with big speed potentials by 5G is observed in a 9-min static test (AT&T). However, 5G potentials are not realized at most time because current practice chooses to use 4G only in presence of good 5G coverage and performance. The number of serving cells repeatedly switches between 5 (with 5G) and 1 (without 5G), and finally ends with 1 (without 5G).

There is no surprise that 5G can download files much faster. It is indeed observed in recent studies (e.g., [8]–[11]) and our measurement study (see an example in Fig. 1). In this example, the download speed rockets to 165.1 Mbps (median, mostly around 100 – 300 Mbps) from 10 – 30 Mbps, as long as 5G (more precisely, 5GM) is used. *However, the problem is that 5G is not used all the time.* The set of serving cells repeatedly switches among three groups: I (five cells with 5G), II (one 4G cell only), and III (another 4G cell only), and finally ends with II/III (without 5G). Note that it is a static test (variances in radio channels are negligible). 5G is not used despite its presence and better performance. This implies that huge speed potentials provided by 5G are not realized at most time – not even close. More details (how and why) are elaborated in §IV.

In this work, we attempt to conduct a measurement study to characterize and examine (un)realized 5G potentials in the wild. Our aim is to give a close look into how they happen and pinpoint what should be responsible for these unrealized potentials (likely problematic and unanticipated). The identified causes shed light on immediate and long-term solutions to enhancing 5G and beyond.

Contributions. We have made four main contributions.

- **Measurement (§II).** We have conducted a 10-month measurement study of all three top-tier US operators (**`A`**, **`V`** and **`T`** for AT&T, Verizon and T-Mobile afterwards) over 13 representative regions (downtown, stadium, campus, residence) in two US cities (>19.8 km^2 in total, Table II). We have collected data traces in the experiments over 705 hours and 10,209 km, covering 1,368 5G cells and more than 133K cellset instances, 4.6M data speed samples, 46M RSRP/RSRQ measurements. To our best knowledge, this is one of the largest-scale studies

	No.	Description	Figure(s)
What (§III)	**F1**	All operators make 5G much faster (2x - 6.7x faster than 4G) but lean on distinct technologies. For **A** and **V**, 5GM is faster but 5GS not. For **T**, 5GS is faster (5GM not observed in our study).	Fig. 2
	F3	For **A** and **V**, 5G utilization is not too bad (despite improvement room) but faster 5G (here, 5GM) is.	Fig. 8
	F4	Do more and get less. Supporting both 5GM and 5GS lowers 5G utilization than supporting only one.	Fig. 6,7
	F6	Losing faster 5G transiently hurts data performance much longer.	Fig. 9,10
Why (§IV)	**F8**	Faster 5G can be missed at any step of a handover, but no valid configuration is a dominant contributor.	Fig. 12
	F9	Channel-specific policies unnecessarily limit the use of faster 5G.	Fig. 16
	F11	Multi-round configuration contributes to long-tail delays (a few seconds) before reaching faster 5G.	Fig. 14,17
Fix (§V)	**F12**	`5GBoost` increases the likelihood of using faster 5G and wider spectrum resources.	Fig. 18, 19
	F13	`5GBoost` at least doubles download speed in more than 50% of instances influenced by `5GBoost`.	Fig. 20
	F14	`5GBoost` effectively reduces the delay of those long-tail (> 2s) instances.	Fig. 22

TABLE I: Summary of our main results out of 14 findings (marked as F1 – F14).

to characterize 5G experience, except those measurement studies done by operators and professionals (e.g., [12], [13]). Our study is open to the research community. **The collected datasets and source codes are available at Github [14].**

• **Characterization (§III).** We quantitively characterize what 5G potentials look like and how much are (un)realized in practice. We find that the above example is not rare. Many instances observed in our measurement study confirm that 5G is indeed able to offer faster data speed (say, hundreds of Mbps) than 4G only but such speed gains are not often realized. Such under-utilization is commonly observed at most locations (more than 50% in most regions for **A** and **V**). We notice one surprising finding: doing more gets less. 5G is less utilized at the places with co-deployed 5GM and 5GS (doing more) than at those places supporting only one, say, 5GS (getting less). We observe that the device suffers with transient disruption although it eventually uses faster 5G. However, the transient impacts last longer, hurting data performance even when faster 5G is in use.

• **Cause analysis (§IV).** We further dive into an in-depth cause analysis. We examine how 5G potentials are missed and find that channel-specific policies and multiple-round configuration should take the blame. We find that the primary cell (PCell) manages radio resource in a performance-oblivious (even performance-unfriendly) manner. Some PCells (mostly at certain RF channels) limit the use of certain 5G cells or channels, thereby lowering the utilization of faster 5G. Moreover, we find that it takes longer (up to a few seconds) to get faster 5G because of multi-round configuration.

• **Solution (§V).** Inspired by our findings, we have proposed `5GBoost`, a quick fix solution with two patches to tackle the identified issues. We have conducted a preliminary trace-driven evaluation and validated its effectiveness to realize more 5G potentials. `5GBoost` at least doubles the median download speed for all three operators.

Table I summarizes our main findings. All the 14 findings (**F1** – **F14**) are elaborated in the rest paper.

II. Methodology and Datasets

We run active measurement on 5G phones with heavy or light traffic. Heavy traffic is used to measure data speed when we repeatedly download bulky files (500MB each) from Google Cloud. Light traffic is to ping Google every second to make radio connection active throughout the experiment. In each run, we record all the cells (both serving and candidate cells) observed at every test location, as well as their radio signal strength/quality (RSRP/RSRQ). We also record the sequence of the serving cellsets and how a cellset switches to another. We use MobileInsight [15], an open-source tool to capture 5G/4G signaling messages which contain rich information needed to understand how and why (§IV).

Due to inherent spatial diversity, we perform a *region-based* study which attempts to cover all the main locations accessible in the given region. Specifically, we run extensive driving tests along all the main roads for a full-region scan. Extra driving/walking/static experiments are later added at the places and routes of our interests. We test with 13 regions in two cities (19.8 Km^2 in total, Table II)[2]. These regions are selected with 5G coverage based on the official coverage maps provided by US operators [4]–[6] (one exception in R4 specified later). They cover representative areas like downtown, stadium, university campus, hospital, and residence. Unless specified, we use one phone model, Google Pixel 5, which supports all RATs: 5GM, 5GS and 4G. For specific purposes (say, A/B tests explained later), we also use Google Pixel 4a, which does not support 5GM, but low/mid-band 5GS. We run experiments sporadically from April 2021 to Jan 2022, with data collected over 705 hours and 10,209 Km (Table II). In total, we observe 1368 5G cells (5GM: 935 and 5GS: 433). In our study, **A** and **V** support both 5GM and 5GS but **T** supports only 5GS, not 5GM. This matches with **T**'s strategy to sidestep 5GM [16].

III. What Does 5G Experience Look Like?

In this section, we characterize 5G coverage, utilization and performance, and quantify how much potentials are unrealized. 5G is indeed faster but not often used as it can and should.

[F1] *All three operators offer much faster 5G, but through different technologies (*5GM *by* **A** *and* **V***, and* 5GS *by* **T***).*

Fig. 2 shows the boxplot (10/25/50/75/90-th percentiles) of download speed observed in all the test regions per RAT (5GM, 5GS and 4G). All three operators make 5G much faster than 4G, but through distinct technologies (actually, spectrum resources). **A** and **V** count on 5GM (not 5GS) but **T** leans on 5GS. From 4G to 5GM, **A** and **V** double download speed (**A**: 73.0 Mbps vs 36.3 Mbps, **V**: 108.7 Mbps vs 51.9 Mbps). For simplicity, we use the median speed unless

[2]City names are hidden for anonymity.

Operator	A (AT&T)					V (Verizon)						T (T-Mobile)		Total
Region	R1	R2	R3	R4	R5	R6	R7	R8	R9	R10	R11	R12	R13	
Area size (Km^2)	4.44	1.23	0.95	1.45	1.05	2.66	1.0	0.8	0.95	1.45	1.05	1.82	0.95	**19.8**
Duration (hour)	202.0	32.5	85.2	36.6	42.1	101.0	18.7	8.8	57.4	24.8	31.6	45.7	18.6	**705.0**
Distance (Km)	2812.3	841.5	974.7	660.8	550.3	1349.3	478.2	194.1	700.7	456.0	435.1	546.7	209.0	**10,209**
Num. of location grids	823	195	294	269	273	577	139	90	294	267	273	344	233	**4,071**
Num. of 5GM cells	375	113	38	**0**	66	117	27	12	93	46	48	**0** (N/A)	**0** (N/A)	**935**
Num. of 5GS cells	58	35	39	20	32	14	14	3	35	25	30	59	69	**433**
Num. of 4G cells	1759	1128	1481	590	1601	1376	929	493	2143	279	660	995	1041	**14,475**
Count of cellset instances	39253	8141	20713	9310	6642	14173	3538	716	10917	5849	6306	4992	3264	**133.8K**
Count of speed samples	1.1M	197K	638K	234K	15K	541K	111K	59K	582K	115K	192K	368K	218K	**4.6M**
Count of RSRPs/RSRQs	12.9M	2.1M	6.1M	1.4M	2.0M	9.5M	1.6M	626K	4.7M	1.3M	1.7M	1.5M	869K	**46.3M**

TABLE II: Statistics of datasets collected in all 13 regions for three US operators: A (R1-R5), V (R6-R11), and T (R12, R13).

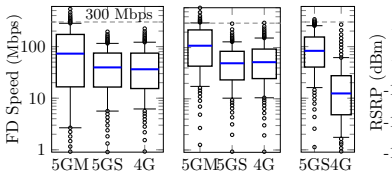

(a) A (b) V (c) T

Fig. 2: Download speed (log-scale).

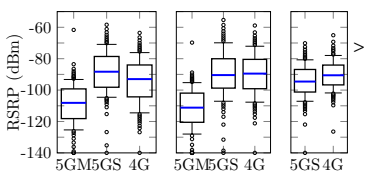

(a) A (b) V (c) T

Fig. 3: RSRP of A, V and T.

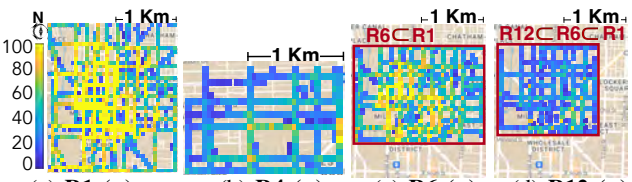

(a) R1 (A) (b) R4 (A) (c) R6 (V) (d) R12 (T)

Fig. 4: Maps of cell density observed in selected regions.

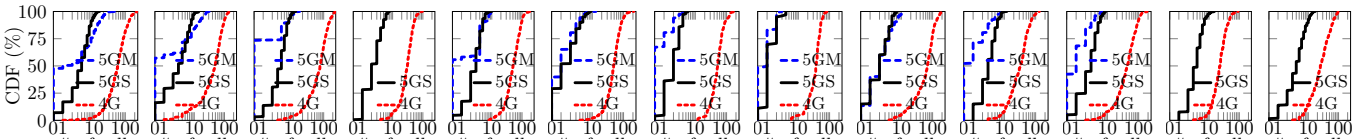

(a) R1 (b) R2 (c) R3 (d) R4 (e) R5 (f) R6 (g) R7 (h) R8 (i) R9 (j) R10 (k) R11 (l) R12 (m) R13

Fig. 5: Spatial distribution of cell density (RSRP $\geq$ -110 dBm) observed in all test regions: R1-R5 (A), R6-R11 (T) and R12-R13 (T).

specified. However, 5GS does not provide speed gains over 4G (**A**: 49.3 Mbps, **V**: 39.5 Mbps). We later show that it is because 5GS mostly operates over the same RF bands originally used for 4G through a technology called dynamic spectrum sharing (DSS) [17]; **A** and **V** offer comparable download speed by 4G and 5GS because both use similar spectrum resources. However, this is not the case for **T**. In contrast, **T** advances 5GS by using more spectrum resources and thus greatly increases data speed, with a 6.7-fold growth from 12.3 Mbps (4G) to 82.8 Mbps (5GS). We notice that 4G in **T** is much slower. This is because **T** repurposes more spectrum resources originally for 4G to run 5GS, which somehow sacrifices 4G performance. More details will be given later (**F9**, Table III).

We want to point out that 5GM or 5GS is not used alone without 4G in our study. As a matter of fact, 5G is added to the existing 4G network primarily in a Non-Standalone (NSA) mode, where a 4G cell acts as a master anchor and 5G cells are used together to offer secondary radio access [18]. NSA is the recommended choice to launch 5G at the start [19]. In our study, we see that **A** and **V** support NSA only, and NSA is the dominant choice to **T**. **T** claimed to support Standalone (SA, 5G as the master anchor) but SA is rarely observed only in R13 (**T**) at 1.4% of time. As a result, we focus on NSA 5G in this paper. 5GM or 5GS actually refers to 4G+5GM (NSA) or 4G+5GS (NSA), while 4G means 4G+NONE, namely, 4G only. The serving cellset consists of one 4G cell as the primary cell (PCell) and several secondary cells (SCells) using 4G and/or 5G. PCell not only provides mandatory radio access but also performs radio resource control to manage all radio access provided by SCells (4G and/or 5G).

[F2] *At places where 5G is launched, all the operators provide good 5G coverage and many choices to use 5G.*

We first check 5G/4G radio coverage. Fig. 4 shows the number of all the cells observed in four selected regions (not all the regions shown due to space limit). Fig. 5 shows the cumulative distribution function (CDF) of cell density with acceptable radio coverage per RAT. Fig. 5 only considers the cells with its median RSRP $\geq$ -110 dBm, which is used as the acceptable coverage criterion in this paper. CDF is the percentage of qualified locations with its acceptable cell density no larger than a given value. We divide all the accessible locations (along main roads) into grids. We test with different grid sizes (here, 0.001/0.0005/0.0001 in latitude and longitude, approximately, 110m $\times$ 90m, 55m $\times$ 45m, 11m $\times$ 9m) and most results are consistent with various grid sizes. We thus use medium-size grids (55m $\times$ 45m) unless specified.

Unsurprisingly, 4G > 5GS > 5GM, in terms of acceptable radio coverage. Everywhere is well covered with 4G. Tens of or even up to 100+ 4G cells are observed at more than 50% of locations. Such dense cell deployment matches with recent measurement studies (e.g., [20]). 5GS is almost everywhere but its cell density is lower. Each location is covered by several 5GS cells. This is attributed to the fact that all US operators quickly launch 5GS by repurposing their existing 4G bands (via DSS). All the operators use a small portion of 4G spectrum to run 5GS (Table III). **A** and **V** only repurpose two or three narrow channels (5MHz or 10MHz each); **T** is much more aggressive with seven channels and use wider channels (up to 60 – 100MHz). This reflects their distinct 5G strategies: **T** invests in 5GS but **A** and **V** count on 5GM more.

5GM is not seen everywhere. In all eligible regions (R1-R11), more than half of locations are not covered by 5GM

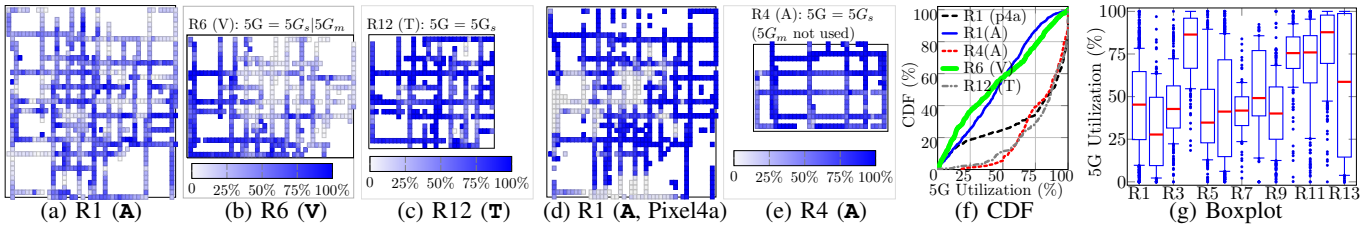

(a) R1 (**A**) (b) R6 (**V**) (c) R12 (**T**) (d) R1 (**A**, Pixel4a) (e) R4 (**A**) (f) CDF (g) Boxplot

Fig. 6: Spatial distribution of 5G utilization in the showcase settings (a-f) and in all the regions (g).

cells with acceptable radio coverage in 8 out of 11 regions, except R4, R6 and R9. R4 is a special exception, where **A** officially claims to deploy 5GM but no 5GM cells are observed in our study. It is mainly because of its limited deployment. Currently, 5GM is deployed on hotspot areas. This matches with 5GM cell towers observed in our study. They are equipped at the existing city infrastructures such as telephone poles and street light. We observed 5GM cell towers in a subregion (R1A in Fig. 7)), not throughout the whole region (R1). Clearly, the coverage ratio goes higher if we consider only the subregion with good 5GM coverage. Another reason is that 5GM coverage is indeed smaller. 5GM fades more quickly than 5GS and 4G because 5GM uses high-band ($>$ 24GHz) while 5GS and 4G use low/mid-band ($<$ 6GHz). As a result, we see that the RSRP values of 5GM cells are slightly smaller (Fig. 3). RSRQ results are similar and omitted due to space limit.

In a nutshell, everywhere is covered with rich radio access choices thanks to dense cell deployment. It is usually covered by several 5GM or 5GS cells (if applicable) plus several tens of 4G cells. As CA selects a group of cells to serve the device, there are a large number of combinations. We do see that the number of unique cellsets is at least one order of magnitude more than the number of cells (Table II).

[F3] *However, faster 5G is not often used at those places with good radio coverage. Such low 5G utilization is not rare.*

We next characterize 5G utilization in reality. As illustrated in our example (Fig. 1), the problem is not that 5G is not used, but 5G is not *often* used where it can. We characterize 5G utilization in a region as follows. At a given location, we first get its 5G utilization score as the ratio of using 5G in all the qualified runs. We further get the distribution of 5G utilization scores across all the locations with 5G coverage. As 5GM is not everywhere, we consider 5GM utilization only at those locations covered by 5GM. Note that for **A** and **V**, 5G can be 5GM or 5GS, (5G = 5GS | 5GM). For **T**, 5G = 5GS.

We first look into 5G utilization, and then examine 5GM and 5GS separately. Fig. 6g plots the 10/25/50/75/90-th percentiles of 5G utilization scores in all the regions. Every region is fully covered by 5G (at least 5GS), though some are not fully covered by 5GM. To better understand spatial patterns, we show 5G utilization scores under several settings (Fig. 6a – 6f). R1 (**A**), R6 (**V**) and R12 (**T**) are three special regions which overlap in the downtown of the same city (see Fig. 4). For the sake of comparison, we place the results over the same area size (2 km $\times$ 2.22 km, same as R1).

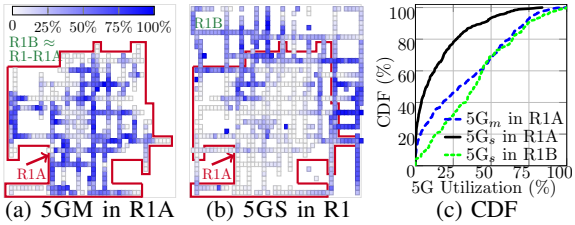

(a) 5GM in R1A (b) 5GS in R1 (c) CDF

Fig. 7: Utilization of 5GM in R1A (covered by 5GM and 5GS) and 5GS in R1A and R1B (covered by 5GS only).

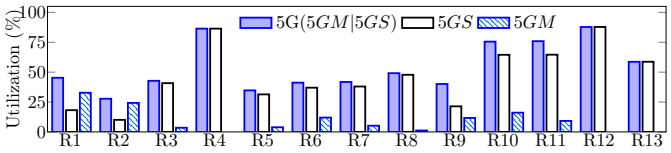

Fig. 8: 5G/5GS/5GM utilization in all the regions.

We have two interesting observations. *(1) 5G utilization is not very low in most regions.* **T** seems to do slightly better than other two operators, particularly in R12. Moreover, **T** uses 5G more often and more widely than **A** and **V** at the same locations (overlapping in R1, R6 and R12). Actually, **T** uses 5G at more than 50% of time at 304 out of 344 grids (88.3%) in R12. In contrast, **A** and **T** use 5G at >50% of time at 43.3% of (356 out of 823) grids in R1 and 42.5% of (245 out of 577) grids in R6. We admit that this comparison is slightly biased because R12 is a subregion of R6 and R1, all with good 5G coverage. Obviously, 5G utilization is region-dependent. 5G is less used in several regions like R2, R3, R5 – R9. At more than half of locations, 5G is used below 50% of time. Compared to **T**, there is more room for **A** and **V** to enhance its 5G utilization.

(2) faster 5G utilization is lower. Fig. 8 shows the median utilization scores in all test regions per RAT (5G, 5GM or 5GS). Note for **A** and **V**, 5GM is faster but 5GS not. However, 5GM is not often used in their regions except R1 and R2. Low 5GM utilization implies that abundant speed potentials enabled by 5GM might not be realized in practice.

[F4] *Do more and get less. 5G utilization becomes lower at places with both* 5GM *and* 5GS.

We notice this surprising result when comparing 5G utilization scores in R1 and R4 (**A**). R4 seems an exception to **A**, where 5G is quite often used. In R4, **A** uses 5G at more than 86.4% of time at half of locations, which is comparable to 87.8% in R12 (**T**) and even higher than 58.7% in R13 (**T**). We further validate it by running A/B tests with Pixel 5 and Pixel 4a (which supports 5GS only) in the same region R1. Surprisingly, the likelihood of using 5G goes much higher if only 5GS is supported (Fig. 6d vs. Fig. 6a). The

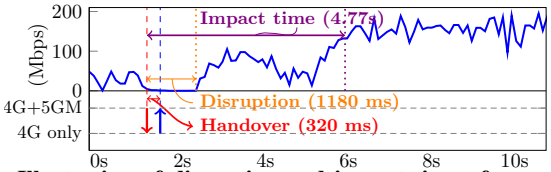

Fig. 9: Illustration of disruption and impact time of a transient miss, which is much longer than the miss lifetime.

median utilization score grows from 45.2% to 88.1% when the phone model downgrades from Pixel 5 to Pixel 4a. Ironically, deploying 5GM reduces 5G utilization when adding more resources for 5G. We believe that **A** and **V** do not intend to launch more to get less (5GM + 5GS $<$ 5GS).

It is the opposite of what we expect. 5GS and 5GM expect to complement each other to increase total resources to run 5G. We see that 5GM hurts 5GS in two counter-intuitive comparisons: Pixel 4a in R1 vs Pixel 5 in R1 and R4 (Pixel 5) vs. R1 (Pixel 5). By "*ignoring*" 5GM, 5G utilization (actually, 5GS utilization) grows. We next show that 5GS also hurts 5GM. This is illustrated in Fig. 7. We show 5GM and 5GS utilization in two subregions (R1A and R1B). R1A is covered by both 5GM and 5GS where R1B is covered by 5GS only. 5GS is used less in R1A than in R1B. 5GM is also used less where 5GS is used more (at the boundary of R1A).

[F5] *Faster 5G is missed transiently and persistently, both lowering data speed that a device can get.*

We find that 5G is missed in two manners: transiently and persistently. It is a transient miss if 5G is not used for a very short while (up to a few seconds) and then gets recovered. Otherwise, it is a persistent miss if 5G never comes back. We later show (§IV) that a persistent miss is due to structural factors (e.g., never considering 5G cells for radio access). As a result, it eventually converges to a state without 5G.

Both are observed in the example (Fig. 1). In this static test, two transient misses are observed at 210s and 262s. Each lasts for 4-7 seconds and then moves to Group I, which uses four 5GM cells, along with one 4G PCell. One persistent miss is observed at the end. After 340s, the device gets stuck into a ping-pong loop where it switches from one 4G cell to another and 5GM cells are not used any longer. It is not hard to understand that missing faster 5G hurts data performance. In the first example (Fig. 1, the download speed declines from 165.1 Mbps (using four 5GM cells plus one 4G cell) to 13.3/23.9 Mbps if using 4G only.

[F6] *The negative impact of a transient miss instance lasts much longer. It is not uncommon to observe >10s data speed degradation due to a transient miss.*

We want to highlight that the negative impact of a transient miss lasts much longer than its short lifetime. We use another example to illustrate its impact in Fig. 9. In this example, 5GM gets recovered within 320 ms. The device not only suffers from data disruption to recover faster 5G connection (here, 1.18s), but also needs much more time (here, 4.77s) to fully recover its high data speed. We define λ-impact time as the duration until data speed get recovered to a λ-fraction of 5G speed. Fig. 10 shows the impact time with λ = 60%, 70%

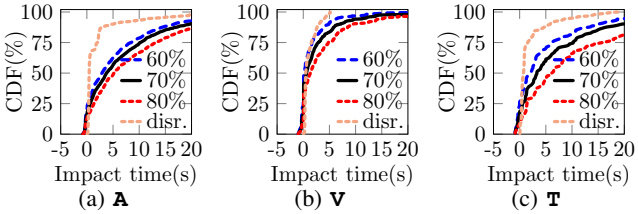

Fig. 10: Impact time caused by all transient miss instances.

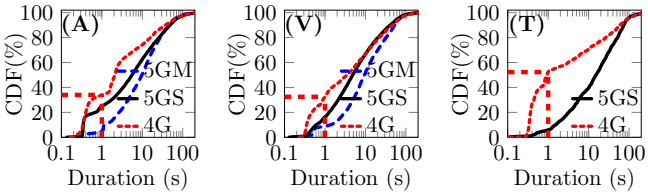

Fig. 11: CDF of the duration/lifetime of all the serving cellset instances. X-axis uses a log-scale. From left to right: A,V and T).

and 80%, as well as the disruption time. The 80%-impact time lasts longer than 10s in more than 10% of instances for **V**, and 30% of transient instances for **A** and **T**. It often requires >10x disruption time to recover from a transient miss.

This example is not a corner case. Fig. 11 plots the lifetime of all cellset instances observed in our study. There are many short-lived 4G-only instances; we see that it is shorter than 1s in in 34%, 33% and 52% of 4G-only instances for **A**, **V** and **T**. We further find that most happens together with switching a 4G PCell first and many will shortly switch to another 5G cellset. For example, for **A**, 5G is missed in 93.2% of instances if the PCell changes.

IV. Why Faster 5G is Not Used?

We next dive into root causes behind the under-utilization of faster 5G, which is 5GM for **A** and **V** and 5GS for **T**.

Cause Analysis. For each miss instance, we examine how a cellset switches to another and locate what prevents the switch to a faster 5G cell despite its presence. Fig. 12 shows our approach. A cellset switch is controlled by the PCell, following standard radio resource control (RRC) procedures (also called handovers) [21], [22]. It is primarily based on radio quality evaluation, namely, by comparing RSRP/RSRQ of the serving and/or candidate cells. The criteria are regulated by standard specifications [21], [22] (like events A1-A6, B1-B2) with tunable parameters such as thresholds and offsets for radio quality comparison. A procedure typically takes three steps: configuration, measurement+reporting (combined), and execution (say, adding or removing 5G cells as SCells).

If the switch successfully goes toward to a 5G cellset (actually, 4G+5GM or 4G+5GS), it must go through all the OK-κ states, each of which is the intermediate state after step κ in the finite-state-machine (FSM). In a miss instance, we locate the cause by checking where it goes "wrong". We find four causes behind missing faster 5G:

(E1) No valid configuration for faster 5G;

(E2) No measurement reporting despite the presence of qualified 5G candidate cells;

(E3) No command to add faster 5G cell(s);

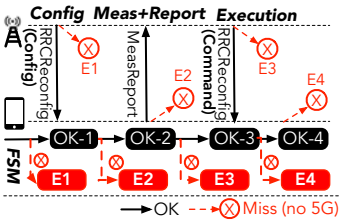

Fig. 12: FSM of a handover.

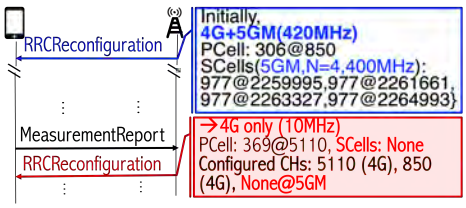

Fig. 13: An example of E1 for Fig. 1.

Fig. 14: An example of multi-round configuration.

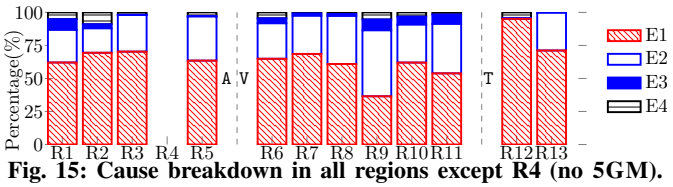

Fig. 15: Cause breakdown in all regions except R4 (no 5GM).

(E4) Faster 5G cell(s) are added but removed shortly.

E1-E3 result in missing 5G; E4 is responsible for transient 5G where 5G quickly disappears. Fig. 15 plots the breakdown of causes for missing 5GM in R1-R11 (**A** and **V**) and missing 5GS in R12 and R13 (**T**). R4 is an exception because 5GM is not observed. We have the following two findings.

[F7] *The cause breakdown is region-specific.*

[F8] *No valid configuration (E1) is a dominant cause.*

Many 4G cells do not configure measurement over RF channels used by faster 5G cells. As a result, these 5G cells are not considered for use. It is the top-1 factor in all the regions except R9 (**V**). E1 is responsible for more than 50% of instances without faster 5G in 11 out of 12 regions. E1 plays a more critical role in **T**; It is responsible for almost all the miss instances (95.2%) in R12.

Fig. 13 gives an illustrative example which is responsible for missing 5GM in the first example (Fig. 1). At the start, the serving cellset uses a 4G PCell (cell ID: 306, channel: 850[3]) and four 5GM cells. When it switches to a new 4G PCell (cell ID: 369, channel: 5110), it only configures measurement over 4G channels and remove all the configurations over 5G (here, 5GM) channels. A direct consequence is that 5GM will not be used. Originally, four 5GM cells run four channels (2259995, 2261661, 2263327, 2264993), each with 100MHz bandwidth (Table III); In another word, they together aggregate 400MHz over band n260 to run 5GM. Without 5GM configuration, all these 5GM cells will not be measured, reported and used. The total spectrum bandwidth sharply drops from 420 MHz to only 10 MHz (channel width in Table III).

E2 (no measurement reporting) is a top-2 cause in almost all the regions except R9. Note that we only consider E2 in all the instances without E1. E2 is mainly due to improper configurable parameters (say, thresholds and offsets for radio quality comparison). Regarding E3, we find that the device served by certain 4G PCells can hardly connect to 5GM or 5GS even with measurement reports. For example, we see that 4G PCells over band 12 or band 14 in **A** never work with 5G cells. It could be the case that **A** intends not to use 5G when those 4G PCells work. It implies strong channel-specific policies (see the next finding **F9**). E4 happens mostly with 5G failures. Specifically, we see that the device reports secondary-cell-group (SCG) failures to the PCell after it received the commands to add these 5G SCell(s) but the radio link to these 5G SCell(s) is too poor for physical transmission. Ideally, they should not happen because a 5G cell should be added by the PCell when its radio quality is acceptable. However, adding 5G cell(s) does not succeed every time. E4 contributes to more instances in some regions like R1 (**A**), R2 (**A**) and R9 (**V**).

The "Why" behind Why. We next attempt to understand the "why" behind current practice. Note that it is rooted in the logic which runs at the network side and is not fully visible to our study. We use a breakdown analysis to expose the policies and the network logic behind such polices.

[F9] *Channel-specific policies from operators unnecessarily limit the use of 5G over wider channels.*

We observe that all three operators invest wider spectrum resources to run 5G. Table III lists basic information of RF bands used by **A**, **V** and **T**. We focus on 5G bands and show their channel width and the number of channels used per band (for each operator). For 4G, we merge all the band information and show the total number of channels over all the 4G bands (**A**: 6, **V**: 7 and **T**: 6). Note that all 5G bands start with n*. All the 5GS bands are originally used for 4G, with the same band number. For example, band n66 and band 66 runs over the same spectrum, but n66 is for 5GS and band 66 is for 4G.

We notice that advanced RATs use wider channels. **A** and **V** rely on 5GM (band n260/n261) and its channel width is 100 MHz (up to 200/400MHz in the coming 5GM upgrade). **T** advances 5GS over one band n41 which offers 60-100 MHz for each channel. In contrast, other 5GS bands use narrow channels (5-20 MHz), which are the same as they are used for 4G. Unsurprisingly, we find that the aggregated amount of spectrum resources is the key to faster 5G. 5GM usually runs faster because it uses much wider spectrum (each cell is 100MHz, at most 4 5GM cells aggregated in our study). **T** offers higher speed not over all 5GS channels, but mainly through 5GS over band n41 with comparable channel bandwidth to 5GM.

Ideally, we expect to use wider spectrum as much as possible, likely boosting data speed. However, we find that it is not the case. Instead, we see that the operators intend to disable the use of those wide 5G channels when the 4G PCell runs over certain channels. Fig. 16 presents the ratios of all RATs associated with each channel which can be used for 4G PCells. In total, we see 12, 9 and 7 channels, a subset of all the 4G channels assigned to 4G PCells ("P" in Table III). Clearly, we see that operators impose channel-specific policies

[3]It is the RF channel number. It is associated with channel information such as working frequency, channel width, band and RAT [7], [23].

RAT	Band	(MHz)	# of CH
5GM	n260	**100**	12 (**A**), 1 (**V**)
	n261	**100**	8 (**V**)
5GS	n41	**60 - 100**	3 (**T**)
	n71	15-20	4 (**T**)
	n5	5-10	2 (**A**), 2 (**V**)
	n66	10	1 (**A**)
4G	**A** (6)	5-20	29 (12 for P)
	V (7)	5-20	61 (9 for P)
	T (7)	5-20	32 (7 for P)

TABLE III: 5G/4G band info.

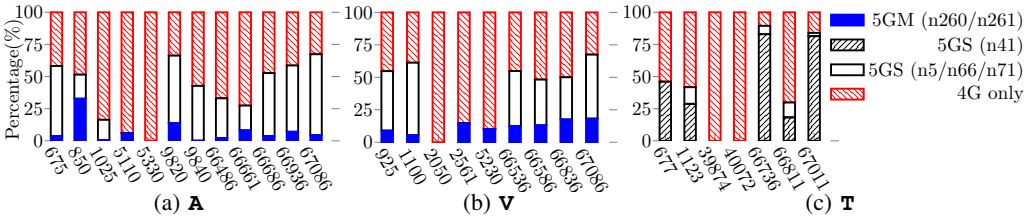

(a) **A** (b) **V** (c) **T**

Fig. 16: Breakdown of SCell RATs per 4G PCell channel.

(and constraints) to manage their radio access.

We see that some channels are never used with 5GS, 5GM or even both. Specifically, channel 5330 (band 14, **A**), 2050 (band 4, **V**), 39874 (band 41, **T**) and 40072 (band 41, **T**) work only with 4G. In a word, once any cell over these channels acts as a PCell, all the 5G access will be disabled. It mainly appears as E1 (with no configuration to measure any 5G channel). Apart from these four channels, the rest prefer 5GM or prefer 5GS. For instance, **A** never uses 5GS with channel 5110 (band 12). It has no configuration over 5GS channels and thus can be aggregated with 5GM and 4G. Moreover, we notice that the ratio of 5GM is quite low (5.9%). It is because of both E1 and E3. That is, in 65% of no-5GM instances using a PCell over channel 5110, 5GM channels are never configured for measurement (E1), as shown in the previous example (Fig. 13). In the rest 23% of instances, the device sends a 5GM measurement report to the PCell indicating at least one good 5GM candidate cell, but the PCell seems to ignore it and does not react to use any 5GM cells (E3). We observe similar things with channel 2561 and channel 5230 in **V**. In contrast, we also see that some channels are not friendly to 5GM, but consider 5GS and 4G only. In **A**, channel 1025 (1972.5 MHz, band 2) and channel 9840 (2357 MHz, band 30) never use 5GM. Channel 1025 is not friendly to 5GS with the similar problems (E1 and E3).

We see that all the channels that support all the RATs, prefer 5GS to 5GM, except channel 850 (**A**). 14 out of 15 channels (8 channels by **A** and 7 channels by **V**) see the 5GM utilization below 18.2%. Channel 850 (**A**, 1955 Mhz, band 2) is the only exception. It is more friendly to 5GM but the help is limited as it takes effect only when a cell over channel 850 acts as the PCell. In contrast, **T** does a much better job to support 5GS more (as 5GS is faster).

All three operators make such channel-specific policies not without rational. Some policies reserve 4G cells for different types of services and ease the management of configurations; They are backward compatible as their role in the past to enhance 4G experience. Some are not updated in time when they expand their 5G deployment. However, we want to point out that such policies are not performance-friendly, at least performance-oblivious. Such bindings make some 4G cells with high speed potentials, some not. But in practice, they are still selected to serve the device based on their radio coverage. While the operators (particularly, **A** and **V**) have heavily invested in their 5G network infrastructure cells and greatly enhanced their potentials to offer faster mobile broadband, such potentials are largely unrealized as speed gains to devices.

Long-tail delay due to multi-round configuration. We

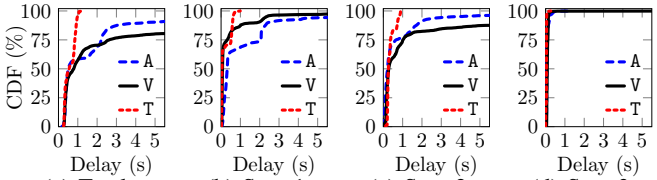

(a) Total (b) Step 1 (c) Step 2 (d) Step 3

Fig. 17: Delay breakdown at three steps: (1) configuration, (2) measurement+reporting and (3) execution (receiving the command from the PCell).

notice that 5G is temporarily missed even when it eventually converges to the desired 5G access (**F7**). We next explore how and why it happens.

[F10] *Long-tail delays (>5s) are observed in 9.6% and 18.8% of transient miss instances in* **A** *and* **V**.

We first check the time needed to complete a transient handover. Fig. 17a plots its distribution and the breakdown at three steps. We see the long-tail delays. The total delay is more than 5s in 9.6% (**A**) and 18.8% (**V**) of the instances, though it is less than 1s in 57.5%-58.4% of instances. We do not observe long-tail delays in **T**. The breakdown analysis shows that configuration at step 1 is one main delay source though it is more influenced at step 2 (measurement+reporting) in **V**. Configuration takes more than 2s in 26.4% of instances (**A**) and 9.9% of instances (**V**). The last step takes less than 500ms in 99.6% of instances and its influence is negligible.

[F11] *Multi-round configuration is responsible for long-tail delays towards faster 5G.*

We find that current practice sometimes runs configuration in multiple rounds. Fig. 14 gives an illustrative example with two-round 5GM configuration. No responses are observed upon the first configuration. Later, the PCell sends a new configuration which asks to measure other 5GM channels. There are many reasons for operators to configure in multiple rounds. This helps them to adapt to runtime dynamics. So, multi-round configuration is not rare in **A** and **V**. It appears in 19.2% and 7.6% of **A** and **V** instances, much higher than 0.3% in **T**; However, it does postpone the recovery from a transient miss to the desired 5G access. For **A** and **V**, 72.1% and 87.2% of the transient instances with multi-round configuration (Fig. 22) experience the configuration delay longer than 2s.

V. 5GBoost: Design & Evaluation

We now present our preliminary efforts to mitigate the issues identified in our study. To be compatible with standard mechanisms, we propose 5GBoost, a quick fix with two patches on top of the legacy 5G cell selection procedure. Fig. 21 depicts its main flow. 5GBoost adds two modules:

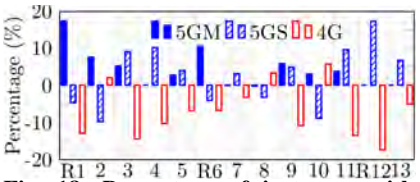

Fig. 18: Percentage of instances with serving RATs changed by `5GBoost`.

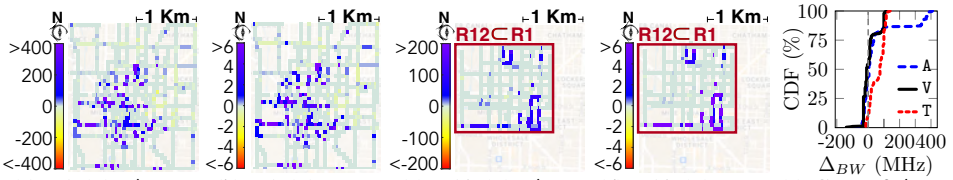

(a) R1 (**A**): Δ_{BW} (b) R1 (**A**): γ_{BW} (c) R12 (**T**): Δ_{BW} (d) R12 (**T**): γ_{BW} (e) CDF of Δ_{BW}

Fig. 19: Impact on the aggregated spectrum bandwidth (MHz).

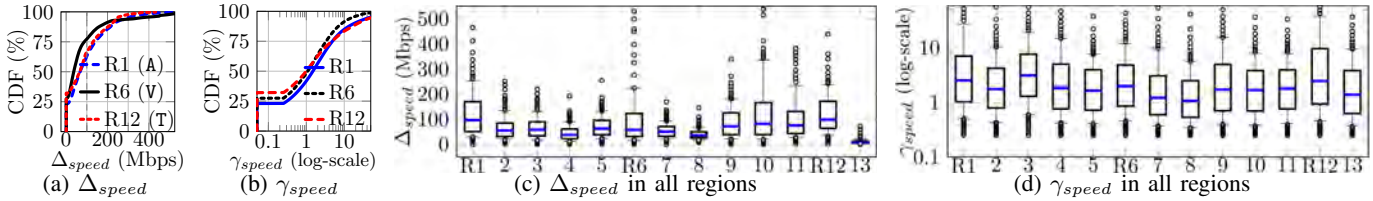

(a) Δ_{speed} (b) γ_{speed} (c) Δ_{speed} in all regions (d) γ_{speed} in all regions

Fig. 20: Download data speed gains with `5GBoost`.

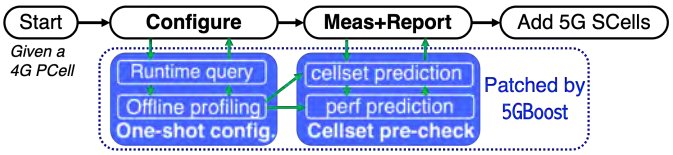

Fig. 21: The `5GBoost` design with two patches.

one-shot configuration and *cellset pre-check* to patch the legacy configuration and measurement/reporting operations.

◦ *One-shot configuration.* To accelerate cellset selection, `5GBoost` uses one-shot configuration to replace legacy multi-round configuration. We leverage historical data per location to profile 5G channels deployed, serving 5G cellsets as well as their possible performance (here, download speed). All the profiles are stored in a database at offline. The profiles can be constructed through crowdsourced measurements or third-party tools to get sufficient historical performance data. At runtime, we query the database to obtain available 5G channels and cells. Finally, all available 5G channels are configured within a single message (still via `RRCReconfiguration`). To handle configuration failures due to poor radio coverage, the configuration message will be retransmitted until the 5G cell addition command is received. This way, the device can immediately obtain 5G channels (cells) of interests, rather than back-and-forth for several times. It promises to reduce the delay and signaling overhead.

◦ *Cellset pre-check.* To enhance 5G utilization (or use wider spectrum to boost data speed), we patch a pre-check module before sending measurement reports of candidate 5G cells. We infer the possible 5G cellsets and look up the profiling database to estimate their aggregated bandwidth and potential performance. We observe that 5G serving cellset is determined by one leading 5G cell (actually, a primary cell of secondary group cells [21]). This makes possible to use one single 5G cell to infer the whole 5G cellset. We check whether the estimated bandwidth and data performance is below an acceptable level, here, a fraction of the optimal data performance among all the possible cellsets. It is called κ-optimal, where κ is a constant weight ($0 < \kappa < 1$, here, $\kappa = 0.8$). If so, reporting will be bypassed, either by dropping this message (by the device) or ignoring the received reporting message (by the network). `5GBoost` considers only the cellset with sufficient performance samples to combat network dynamics. By this means, we attempt to reduce the likelihood of selecting a 5G cellset that hurts network performance.

Evaluation. It is impossible for us to deploy `5GBoost` in commercial 5G networks to evaluate its effectiveness. Here, we have conducted a preliminary evaluation using a trace-driven study. For every handover instance collected in our real-world experiments, we run `5GBoost` to check whether and how it changes, say, selecting a different serving cellset or completing it in a shorter time or nothing changes. Note that `5GBoost` impacts some, but not all handover instances. If the serving cellset changes, we use historical data of the selected cellset to estimate its performance in this what-if study. We admit that our evaluation fails to take real-world dynamics and complexity into account. Here, we attempt to give a rough assessment on possible gains.

[F12] *`5GBoost` is effective in boosting the use of faster 5G and more spectrum resources.*

Fig. 18 shows the percentage of cellset instances changed by `5GBoost`. In most regions, we see that `5GBoost` helps to use 5G more, particularly faster 5G (5GM in R1-R11 except R4 for **A** and **V**, and 5GS in R12 and R13 for **T**). In R12 and R13, `5GBoost` helps **T** to use 5GS more. In R1 – R11 (except R4), `5GBoost` boosts the use of 5GM in 8 out of 10 regions, up to 17% in R1. We notice that `5GBoost` does not increase the use of 5GS in these regions. It is because 5GM can significantly increase data speed over 4G while 5GS cannot. As `5GBoost` is performance-oriented, it tends to select 5GM cells in most of cases. There is no much incentive for `5GBoost` to switch from 4G to 5GS.

`5GBoost` also increases the total spectrum resources aggregated to serve a mobile device. To evaluate its impact on aggregated spectrum bandwidth, we define two metrics: Δ_{BW} and γ_{BW}, as the absolute and relative gain (loss) imposed by `5GBoost`. That is, $\Delta_{BW}{=}BW_{\texttt{5GBoost}} - BW_{legacy}$ and $\gamma_{BW}{=}\Delta_{BW}/BW_{legacy}$. Fig. 19 uses R1 (**A**) and R12 (**T**) as two showcase regions to visualize spectrum gains (losses) at

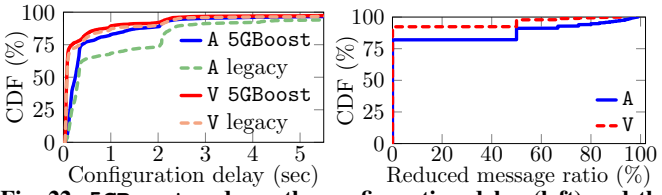

Fig. 22: `5GBoost` reduces the configuration delay (left) and the number of configuration signaling messages (right).

various locations. Note that Δ_{BW} can be negative (loss) but it is rare. Spectrum gains are observed at 20.9% and 17.7% of locations in R1 and R12. It is much more common than the losses at 3.7% and 1.7% of locations. At the rest locations, the use of spectrum resources is not impacted by `5GBoost`. More importantly, the gains are much larger than the losses. The median gain of Δ_{BW} (γ_{BW}) is 45 MHz (100%) and 80 MHz (425%) in R1 and R12, while the median loss is only 20 MHz (50%) and 10 MHz (8.3%) in those worsened cases. In the best case, `5GBoost` boosts spectrum width by 410 MHz and 140 MHz in R1 and R12, respectively. We further extend the evaluation to all the regions: R1-R5 (**A**), R6-R11 (**V**) and R12-R13 (**T**). Fig. 19e shows the CDF of Δ_{BW} in all regions; We exclude those locations not influenced by `5GBoost` (Δ_{BW}=0) and consider only the locations impacted by `5GBoost` in Fig. 19e. We clearly see that the loss is negligible compared with the spectrum gain, especially for **A** and **V**. For **A**, `5GBoost` adds $>$ 50 MHz spectrum in 19.3% of cases, and only loses >50 MHz in 1.2% of cases. The spectrum gain even reaches $>$ 300 MHz in 12.7% of cases, where `5GBoost` increases the number of 5GM cells from 0/1 to 4 (maximum in our study). The spectrum gains are more common in **T** regions (R12 and R13). `5GBoost` uses $>$ 50 MHz in 60.2% of cases and such loss is never observed. For **V**, `5GBoost` increases at least 100 MHz in 15.3% of cases.

[F13] *`5GBoost` at least doubles download speed in more than half of instances impacted by `5GBoost`.*

Similarly, we use Δ_{speed} or γ_{speed} to assess the absolute and relative speed gains by `5GBoost`. Fig. 18 shows that `5GBoost` greatly enhances 5G experience in terms of download speed. We first show their distributions in three showcase regions: R1 (**A**), R6 (**T**) and R12 (**T**). `5GBoost` increases data speed in 68.9%-77.0% of instances. The median speed gain is 90.5-97.9 Mbps. It at least doubles data speed in 58.1%, 51.6% and 49.9% of cases in R1 (**A**), R6 (**T**) and R12 (**T**), respectively. It is not hard to understand with **F12**. We observe such data speed gains in all 13 regions. Fig. 20c and Fig. 20d plot Δ_{speed} and γ_{speed} of cases impacted by `5GBoost`. We exclude the cases if nothing changes. `5GBoost` doubles data speed ($\gamma_{speed} \geq 100\%$) in more than half of runs in each region. In 10 out of 13 regions, `5GBoost` increases data speed by at least 50Mbps ($\Delta_{speed} \geq 50\ Mbps, median$).

[F14] *`5GBoost` reduces the long-tail delay (say, $>$ 2s) caused by multiple configuration rounds.*

`5GBoost` reduces the signaling message overhead and the cell selection delay during configuration phase (Fig. 22). It is attributed to one-shot configuration which reduces the delay, particularly the long tail caused by multiple configuration rounds. We see that `5GBoost` reduces the ratio of long configuration delays (>2s) from 26.6% to 11.1% in **A**. It is effective in 58.3% of long-tail instances. In the meanwhile, `5GBoost` saves more than half of signaling messages for configuration in 17.9% of **A** instances and 7.6% of **V** instances.

VI. Related Work

Recent years have witnessed active 5G measurement studies from both industry (e.g., [12], [13], [24]) and academia [8]–[11], [25]–[31]. Measurement studies from industry (by network operators and professionals like OpenSignal and OOkla Speedtest) mostly run at an extreme scale by leveraging their nationwide network infrastructure or crowdsourcing measurement from millions of phones. Such studies focus on characterizing real 5G experience without explaining why. Research by the academia dive into depth, but they run their measurement experiments usually at a small scale, e.g., at several 5G sites [8], [25] or over a small area [9], [29], unless they use huge data from network operators and service providers [10], [27], [31]. All early studies focused on measuring 5G radio coverage and performance (mainly, data speed) [8], [9], [25], [26]. For example, [8] was the first work to measure 5GM experience when several mmWave cells were just launched in Minnesota in 2019. [9] was a following measurement study which measured 5GS performance in a China's campus (5GM not supported). Recent studies started to measure beyond performance such as cell deployment [28], reliability [27], performance potentials [29], performance and power [30], and mobility management [31].

Our work is different from all the existing measurement studies. Our focus is not to characterize download speed, but investigate how much speed potentials are missed, understand why and attempt to solve the identified problems. The most relevant work is [20], which was the first time to disclose and measure missed potentials in 4.5G. Our difference and advance from [20] is our brand new findings and insights on 5G. Our study have covered both 5GM and 5GS by three US operators in 13 regions. As 5G advances with wider spectrum and more CA choices, potentials are missed in different forms.

VII. Conclusion

5G is advancing – and very fast. However, not all 5G benefits come without cost and pains. In this work, we have characterized 5G potentials realized and unrealized in a 10-month measurement study of all three US top-tier operators (**A**, **V** and **T**). We find that 5G is well deployed but not fully utilized as it can and should. The identified issues do not seem to disappear with upcoming 5G upgrades. But the good news is that they are ready to fix with software patches (enhancing radio resource control), with no need of upgrading physical network infrastructure.

Acknowledgment. We thank anonymous reviewers for their constructive comments. We are also grateful to Kai Ling, Kwang Taik Kim and Mung Chiang for their help at the early phase. This work was partially supported by NSF grants CNS-1750953 and IIS-2112471.

REFERENCES

[1] 5G America, "The Mobile Economy 2022," https://www.gsma.com/mobileeconomy/wp-content/uploads/2022/02/280222-The-Mobile-Economy-2022.pdf, Feb 28, 2022.

[2] GSMA, "The Mobile Economy North America 2021," https://www.gsma.com/mobileeconomy/northamerica/, accessed on July 25, 2022.

[3] 3GPP, "TS22.261: Service Requirements for the 5G System," 2021, v16.14.0.

[4] ATT, "Coverage Map," https://www.att.com/5g/coverage-map/, accessed on June 28, 2022.

[5] Verizon, "Coverage Map," https://www.verizon.com/coverage-map/, accessed on June 28, 2022.

[6] T-Mobile, "Coverage Map," https://www.t-mobile.com/coverage/coverage-map, accessed on June 28, 2022.

[7] 3GPP, "TS36.101: E-UTRA; User Equipment (UE) radio transmission and reception," 2021, v16.9.0.

[8] A. Narayanan, E. Ramadan, J. Carpenter, Q. Liu, Y. Liu, F. Qian, and Z.-L. Zhang, "A first look at commercial 5g performance on smartphones," in *Proceedings of The Web Conference 2020*, 2020, pp. 894–905.

[9] D. Xu, A. Zhou, X. Zhang, G. Wang, X. Liu, C. An, Y. Shi, L. Liu, and H. Ma, "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*, ser. SIGCOMM'20, 2020.

[10] X. Yuan, M. Wu, Z. Wang, Y. Zhu, M. Ma, J. Guo, Z.-L. Zhang, and W. Zhu, "Understanding 5g performance for real-world services: a content provider's perspective," 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 (SIGCOMM'22)*, 2022, pp. 101–113.

[11] X. Yang, H. Lin, Z. Li, F. Qian, X. Li, Z. He, X. Wu, X. Wang, Y. Liu, Z. Liao *et al.*, "Mobile access bandwidth in practice: Measurement, analysis, and implications," 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 (SIGCOMM'22)*, 2022, pp. 114–128.

[12] OpenSignal, "5G User Experience Report USA: Jan 2022," https://www.opensignal.com/reports/2022/01/usa/mobile-network-experience-5g, Jan 2022.

[13] Ookla, "OOKLA 5G MAP," https://www.speedtest.net/ookla-5g-map, accessed on July 21, 2022.

[14] "5G Measurement Datasets and Codes," https://github.com/mssn/INFOCOM23-5GMeas, 2023.

[15] MobileInsight, http://www.mobileinsight.net, 2022.

[16] "How and why T-Mobile sidestepped mmWave 5G," https://www.lightreading.com/5g/how-and-why-t-mobile-sidestepped-mmwave-5g/d/d-id/773678, Nov 23, 2021.

[17] 3GPP, "TS38.211: NR; Physical Channels and Modulation," 2021, v16.3.0.

[18] ——, "TS23.501: System Architecture for the 5G System," 2021, v16.9.0.

[19] 5G America, "3GPP Releases 16, 17 and Beyond," https://www.5gamericas.org/wp-content/uploads/2021/01/InDesign-3GPP-Rel-16-17-2021.pdf, Jan 2021.

[20] H. Deng, K. Ling, J. Guo, and C. Peng, "Unveiling the Missed 4.5G Performance In the Wild," in *Proceedings of the 21st International Workshop on Mobile Computing Systems and Applications*, ser. HotMobile'20, March 2020.

[21] 3GPP, "TS38.331: NR; Radio Resource Control (RRC)," 2021, v16.4.1.

[22] ——, "TS36.331: E-UTRA; Radio Resource Control (RRC)," 2020, v15.9.0.

[23] ——, "TS38.104: NR; Base Station (BS) radio transmission and reception," 2021, v16.7.0.

[24] OpenSignal, "Benchmarking the global 5G experience," https://www.opensignal.com/2021/09/09/benchmarking-the-global-5g-experience-september-2021, Sep 2021.

[25] A. Narayanan, E. Ramadan, R. Mehta, X. Hu, Q. Liu, R. A. Fezeu, U. K. Dayalan, S. Verma, P. Ji, T. Li *et al.*, "Lumos5g: Mapping and predicting commercial mmwave 5g throughput," in *Proceedings of the ACM Internet Measurement Conference (IMC'20)*, 2020, pp. 176–193.

[26] T. Liu, J. Pan, and Y. Tian, "Detect the bottleneck of commercial 5g in china," in *2020 IEEE 6th International Conference on Computer and Communications (ICCC)*. IEEE, 2020, pp. 941–945.

[27] Y. Li, H. Lin, Z. Li, Y. Liu, F. Qian, L. Gong, X. Xin, and T. Xu, "A nationwide study on cellular reliability: measurement, analysis, and enhancements," 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 (SIGCOMM'21)*, 2021, pp. 597–609.

[28] M. I. Rochman, V. Sathya, N. Nunez, D. Fernandez, M. Ghosh, A. S. Ibrahim, and W. Payne, "A comparison study of cellular deployments in chicago and miami using apps on smartphones," *arXiv preprint arXiv:2108.00453, Accepted by ACM WiNTECH 2021*, 2021.

[29] Q. Li and C. Peng, "Reconfiguring cell selection in 4g/5g networks," in *IEEE International Conference on Network Protocols*, 2021.

[30] A. Narayanan, X. Zhang, R. Zhu, A. Hassan, S. Jin, X. Zhu, X. Zhang, D. Rybkin, Z. Yang, Z. M. Mao, Q. Qian, and Z.-L. Zhang, "A variegated look at 5g in the wild: performance, power, and qoe implications," 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 (SIGCOMM'21)*, 2021, pp. 610–625.

[31] A. Hassan, A. Narayanan, A. Zhang, W. Ye, R. Zhu, S. Jin, J. Carpenter, Z. M. Mao, F. Qian, and Z.-L. Zhang, "Vivisecting mobility management in 5g cellular networks," 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 (SIGCOMM'22)*, 2022, pp. 86–100.