An In-Depth Look into 5G ON-OFF Loops in the Wild

Yanbing Liu
Purdue University
liu3098@purdue.edu

Jingqi Huang
Purdue University
huan1504@purdue.edu

Sonia Fahmy
Purdue University
fahmy@purdue.edu

Chunyi Peng
Purdue University
chunyi@purdue.edu

Abstract

5G is much faster than 4G, offering faster data transfer and better user experience overall. Intuitively, 5G should be used as much as possible. However, in this study, we unveil a surprising finding in operational 5G networks: 5G radio access may be in a *persistent ON-OFF loop* which repeatedly turns 5G on and then off. We conduct extensive measurement experiments with three US operators (T-Mobile, AT&T, and Verizon) in two US cities to characterize and analyze 5G ON-OFF loop instances in the wild. Surprisingly, we find that such 5G ON-OFF loops are not rare. They are widely observed at many places, significantly hurting data performance (from several hundreds of Mbps to tens of or even zero Mbps). We further dive into their causes and uncover that inconsistent triggers to turn 5G on and off co-exist in real-world settings, repeatedly releasing 5G radio access after getting 5G back. We identify three loop types each with distinct triggering events/causes (sub-types). Inconsistent policies and mechanisms on both network and device sides, as well as "improper" use of certain frequency channels, are responsible for the loops observed in this study. **Our datasets and artifacts have been released on Github [4] and MI-LAB [5].**

CCS Concepts

• **Networks** → **Mobile networks**; **Network reliability**; **Network measurement**.

Keywords

5G, SA, NSA, Radio Resource Control, Configuration, ON-OFF Loop

ACM Reference Format:
Yanbing Liu, Jingqi Huang, Sonia Fahmy, and Chunyi Peng. 2025. An In-Depth Look into 5G ON-OFF Loops in the Wild. In *Proceedings of the 2025 ACM Internet Measurement Conference (IMC '25), October 28–31, 2025, Madison, WI, USA.* ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/3730567.3764479

1 Introduction

We are entering the 5G era, which delivers much higher performance than previous generations. Globally, 5G is rapidly rolling out, with the number of 5G connections reaching a milestone of 2 billion in Q3 2024 and hitting 7.7 billion by 2028 [10]. 5G is significantly faster, with its average download speed over 3× – 6× faster than 4G and its peak download speed up to 20× faster (up to 1 Gbps) [25].

Intuitively, 5G should be used as much as possible, wherever it can offer better mobile broadband experience. However, we find that **this is not always the case in operational 5G networks**. A mobile device often loses its 5G radio access even when there are *no noticeable changes* to its 5G radio channels in use (more precisely, when 5G serving cells still offer high radio signal quality and fast data access). In these cases, 5G radio access is often recovered but is later lost again. As a result, 5G radio access enters an **ON-OFF** loop, as illustrated in Figure 1. Figure 1a gives the corresponding finite state machine (FSM) where radio access oscillates between two states – **5G ON** and **5G OFF** – limiting the performance potential of 5G networks in the wild. Figure 1b gives a real-world instance observed in our measurement study (additional instances are given in §4) and the resulting performance degradation. Here, our primary test phone (OnePlus 12R) runs a speed test using a bulk file download application in a stationary experiment (at P16 in campus area A1, Figure 7). This test location is covered well by 5G networks deployed by T-Mobile, which supports 5G Standalone (SA) [7] and only uses 5G to serve the test phone (here, OnePlus 12R) at this location. When 5G is used (5G ON), the download speed is around 200 Mbps or higher, but it quickly drops to zero when losing 5G radio access (5G OFF).

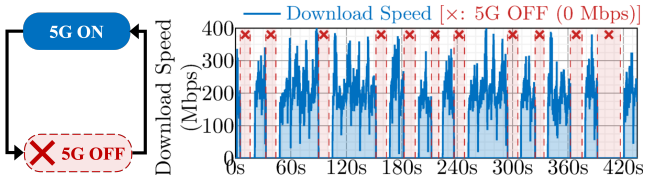

(a) State machine (b) A real-world instance

Figure 1: Illustration of one 5G ON-OFF loop (×: 5G OFF).

Such 5G ON-OFF loops are much worse than anticipated because they are **persistent** loops, not **transient** ones, which are temporary problems caused by network changes and environment dynamics such as radio signal quality fluctuations and movement (like ping-pong handovers [31]). Instead, 5G ON-OFF loops *persist* under the same (or quasi-same) environment. Both state transitions (5G ON → 5G OFF and 5G OFF → 5G ON) are invoked when real-world network conditions remain (largely) unchanged. This implies that persistent loops, unlike transient ones, are mostly *operational anomalies* caused by structural flaws and misconfigurations which result in the co-existence of all state transitions needed for this loop (two transitions in this example). We note that persistent loops have been previously studied in other network contexts, e.g., routing loops in the Internet (e.g., BGP route loops [11, 20] and persistent forwarding loops [21, 32]), persistent handover loops in 3G/4G networks [17, 38] and inter-operator switch loops [36]. Our work is inspired by these prior studies, but the identified loop problems are completely different.

In this work, we focus on **5G ON-OFF loops** where **5G radio access is repeatedly lost (5G OFF) in cases when 5G should be and can be ON** (5G is indeed used but not all the time). These persistent ON-OFF loops do not simply undermine stability with an

IMC '25, Madison, WI, USA

ACM ISBN 979-8-4007-1860-1/2025/10
https://doi.org/10.1145/3730567.3764479

		Finding	Description	Operators	Ref.+Results
Q1	※	F1, F2	5G ON-OFF loops are **much more common** than anticipated.	OP_T,OP_A,OP_V	§4.2, Figure 6,8,9
		F3, F4	Loops significantly hurt data performance and even suspend data service.	OP_T,OP_A,OP_V	§4.3, Figure 1,10,11
	※	F5	ON-OFF loops over 5G NSA are consistently observed across almost all phone models.	OP_A,OP_V	§4.4, Figure 12
	※	F6	ON-OFF loops over 5G SA are observed only with one phone model (OnePlus 12R).	OP_T	§4.4
Q2	※	F7, F8	**Inconsistent ON/OFF triggers create loops** (three types: S1, N1 and N2).	OP_T,OP_A,OP_V	§5, Figure 13
	※	F9	**"A few bad apples ruin all"**. One/few poor 5G SCells ⇒ losing all 5G SA cells.	OP_T	§5.1, Figure 14
	※	F10	**"4G ruins 5G"**. 4G fails ⇒ losing all 5G NSA cells.	OP_A, OP_V	§5.2, Figure 15
	※	F11, F12	**Inconsistencies in policies and mechanisms** change, but inconsistency persists.		
	※	F14, F15	Certain channels and policies should take the blame.	OP_T,OP_A,OP_V	§5.3, Figure 17-19
Q3	※	F16, F17	Key features to predict S1E3 loops (SCell modification failure) are identified.	OP_T	§6, Figure 20-22
		F18	Loop prediction is promising for certain loop types.		

Table 1: Summary of our main findings. ※ is used to highlight new and surprising findings.

oscillation among multiple states; rather, they harm 5G reliability and significantly hurt data performance, e.g., lose several hundreds of Mbps and even suspend data service. In this example (Figure 1b), 5G frequently switches on and off every several tens of seconds (11 times in 420 seconds and there is no end in sight). Clearly, these persistent loops hurt data performance with long-lasting effects. Such performance degradation is unnecessary because good 5G radio access is indeed available and offers high data speed.

Our goal is to **measure and analyze persistent 5G ON-OFF loops in operational 5G networks**, and answer the following questions: *(Q1) How likely and how often do such 5G ON-OFF loops occur in the wild? What are the resulting performance impacts? (Q2) Why do these loops occur? What are their causes? (Q3) What can be done to mitigate such loops in operational 5G networks?*

To this end, we conduct a measurement study with three major 5G operators – T-Mobile, AT&T and Verizon (abbreviated as OP_T, OP_A and OP_V) – in the US. Table 1 summarizes our key findings. We use ※ to mark findings that we believe to be particularly new and surprising. We make four main contributions in this work.

First, to our knowledge, we conduct the first measurement study to characterize 5G ON-OFF loops in the wild (§4.1). Our study covers both 5G deployment options: standalone (SA) and non-standalone (NSA), with extensive experiments over 11,000 minutes in 11 test areas ($\sim$ 19 Km^2) with three major US operators (OP_T, OP_A, OP_V).

Second, we reveal the surprising finding that 5G ON-OFF loops are quite common in reality (§4.2). This is different from the common expectation that 5G radio access is reliably used in operational 5G networks. Instead, repeatedly losing 5G access is widely observed at many test locations in our study. 5G ON-OFF loops occur in >50% of our tests, at more than half of test locations, and occur with 100% likelihood in quite a few places. Unsurprisingly, such loops significantly hurt data performance (§4.3). We further find that loops over 5G NSA (with OP_A and OP_V) are observed with several test phone models (with one exception) while loops over 5G SA (with OP_T) are observed with one specific phone model (device factors explained in §4.4).

Third, we conduct an in-depth cause analysis of the loops (§5). We identify three primary loop types, each consisting of several sub-types with distinct triggering events/causes. We unveil new forms of inconsistency where inconsistent triggers to turn 5G on and off co-exist in operational 5G networks, resulting in 5G ON-OFF loops. We further explore the reasons behind network operator policies and practice.

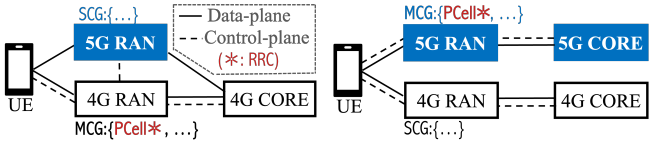

(a) 5G NSA (b) 5G SA

Figure 2: 5G network architecture and radio access using two deployment options: 5G SA and 5G NSA.

Finally, we make the first attempt to model the loop likelihood in practice (§6). Our results highlight the impact of radio conditions on triggering ON-OFF loops and indicate that there is potential to predict certain loop types at any location. We discuss the implications and limitations in §7.

Ethics and Release. This work does not raise any ethical issues. Datasets and code used in this work are released on Github [4]. A guide for reproducing our experiments and executing experiments with additional scenarios is available at [5].

2 5G Radio Access Primer

In this section, we introduce necessary background on 5G radio access and radio resource control. Appendix A (Table 6) lists all the abbreviations and acronyms used in this paper.

5G network architecture and radio access. Figure 2 depicts the 5G network architecture, which consists of two main subsystems: radio access network (RAN) and core network (CORE), similar to previous generation cellular networks. 5G is deployed with two options – 5G SA and 5G NSA [7]. 5G NSA leverages the existing 4G network infrastructure to quickly launch 5G. 5G NSA uses 4G RAN to offer master radio access and 5G RAN for secondary radio access to transfer user traffic. In contrast, 5G SA operates 5G independently of 4G infrastructure. 5G RAN offers master radio access and 4G RAN, if used, is for secondary radio access.

In a cellular network, a cell is the *logical* unit to offer radio access, running one radio access technology (RAT, here, 5G or 4G) over one frequency channel with its fixed channel width ranging from 5 MHz to 100 MHz [8]. Each cell covers a limited geographical area (say, a coverage sector) and physically resides at one cell tower, which accommodates many cells over various frequency channels and directional antennas . In a nutshell, radio access is provisioned through *a set of serving cells*, which is divided into two groups: master cell group (MCG) and secondary cell group (SCG) if both radio access technologies, or RATs (here, 5G and 4G) are used. The MCG is mandatory while the SCG is optional. A cell group consists

of one primary cell which is mandatory and several secondary cells (SCells) which are optional. PCell/PSCell is the primary cell of the MCG/SCG. PCell is the main control point responsible for radio resource control (RRC) such as establishing, managing/modifying and releasing RRC connections [6, 9].

Turning 5G ON/OFF via RRC. In this work, we define 5G ON as long as 5G radio resources are actively used, either in SA mode (5G as the master RAT, optionally with 4G) or in NSA mode (4G as the master RAT with 5G as secondary RAT) to offer radio access. Conversely, we define 5G OFF as long as no 5G resources are used, including 4G-only connectivity (4G without 5G) and IDLE (no active radio connectivity). Turning 5G ON/OFF is realized by standard RRC procedures regulated by 3GPP (TS38.331 [9] for 5G RRC, TS36.331 [6] for 4G RRC). Specifically, 5G is turned ON by establishing a new RRC connection with 5G serving cells or by adding 5G cells as SCells of the existing RRC connection. Each 5G cell to be added must meet the requirement of good radio signal quality, i.e., its RSRP/RSRQ measurement should be larger than a given threshold or stronger than the serving cell by a certain offset (all measurement and reporting events defined in [6, 9]). 5G becomes OFF when RRC releases an existing connection with 5G serving cells or removes all 5G serving cells of the existing RRC connection. It can be triggered by certain radio signal measurements (e.g., RSRP/RSRQ measurements smaller than one threshold) or radio link events (e.g., radio link failures). All RRC procedures involved in 5G ON-OFF loop instances will be elaborated later.

3 A Closer Look at the Example

In this section, we delve into the motivating example (Figure 1b) to understand how one 5G ON-OFF loop is persistent at one location (P16, Figure 7). The example motivates our larger-scale measurement study and analysis in operational 5G networks.

We run stationary experiments at this test location. For each run, we extract the sequence of serving cell sets (CS) using the methodology detailed in §4.1 and Appendix B. Table 2 lists five main 5G cells observed in this example. These cells operate over four frequency channels in two 5G bands (n41 and n25). Each cell is denoted by ID@FreqChannelNo, where ID is the cell identifier and FreqChannelNo is the radio frequency channel number (here, ARFCN for 5G and EARFCN for 4G) specified by 3GPP [8]. $5G_1$ and $5G_2$ run over two frequency channels – 521310 and 501390 – centered at 2607 MHz and 2507 MHz with wide channel widths (90 MHz and 100 MHz); The remaining three cells run over the same band (n25) with a much narrower channel width (10 MHz). $5G_3$ runs over one frequency channel (398410, centered at 1992 MHz) while $5G_4$ and $5G_5$ run over another channel (387410, centered at 1937 MHz). Table 2 lists their measured RSRP values (median ± deviation) through extensive measurements (>500 samples per cell) in many runs at the same location. RSRP is the default metric of radio signal quality in RRC procedures and RSRQ values are omitted unless specified. Clearly, this location has good 5G coverage.

Figure 3a shows the corresponding FSM, which oscillates between two states (5G SA and IDLE). Figure 3b shows how RRC procedures are repeatedly performed over time, using two ON-OFF cycles in the first 44 seconds of Figure 1b. Due to space limits, the detailed procedures along with all main signaling messages are given in Appendix B (Figure 24 - Figure 26).

5G Cell		Band	Ch. Freq	Width	RSRP ($\pm\sigma$)
$5G_1$	393@521310	n41	2607 MHz	90 MHz	-82 ± 9.6 dBm
$5G_2$	393@501390	n41	2507 MHz	100 MHz	-82 ± 9.8 dBm
$5G_3$	273@398410	n25	1992 MHz	10 MHz	-81 ± 9.7 dBm
$5G_4$	273@387410	n25	1937 MHz	10 MHz	-82 ± 7.0 dBm
$5G_5$	371@387410	n25	1937 MHz	10 MHz	-86 ± 9.5 dBm

Table 2: 5G cells in the example (Figure 1b) with OP_T (5G SA).

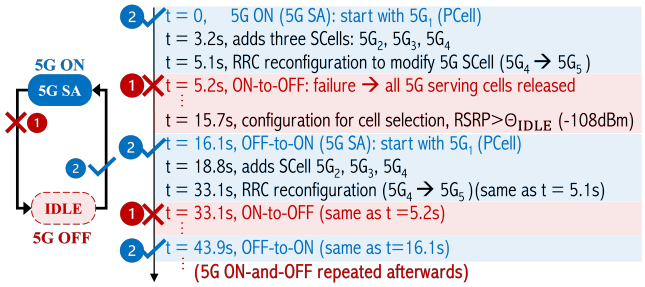

(a) FSM (b) RRC procedures over time (Ex: Figure 1b)

Figure 3: 5G ON-OFF transitions are repeated in the example.

5G SA ⇔ IDLE in the example (Figure 1b). The phone starts at one 5G ON state (5G SA) with $5G_1$ as the PCell. It is realized through one `RRC Connnection Establishment` procedure [9], which is used to set up one RRC connection when there is no active RRC connection (IDLE). Later (at t = 3.2 s), one `RRC Reconfiguration` procedure [9] is performed to add three 5G cells ($5G_2$, $5G_3$ and $5G_4$) as three SCells. Then, another `RRC Reconfiguration` procedure is performed to modify SCells but fails, resulting in the release of all 5G serving cells. Specifically, at t = 5.1 s, the phone receives another `RRC Reconfiguration` message, which commands the phone to replace $5G_4$ with $5G_5$ ($5G_4 \rightarrow 5G_5$), both over the same frequency channel (@387410). However, the phone fails to complete SCell modification, releasing all 5G serving cells and returning to the IDLE state (more in Figure 26). Note that the phone has to re-establish a new RRC connection because there is data to transfer (here, file download). After about 10 seconds (t = 15.7 s), the phone receives system information which is broadcast to guide the phone to select an appropriate cell to establish an RRC connection. In this example, the RSRP threshold is set to -108 dBm ($\Theta_{IDLE} = -108$ dBm) for selecting a 5G cell in band n41. As long as the RSRP of one 5G cell in band n41 exceeds -108 dBm, its signal strength is believed to be good enough for the phone to establish a 5G connection through this cell. In this example, the phone selects $5G_1$ as the PCell to re-establish another RRC connection. Again, three SCells (here, $5G_2$, $5G_3$ and $5G_4$) are later added after this RRC connection is established (within 3 seconds). For sake of simplicity, we show the state as **5G SA** as long as $5G_1$ is used, regardless of 5G SCells in use. Afterwards, the above OFF-then-ON process is repeated. At t = 33.1 s, the phone receives the message to perform one `RRC Reconfiguration`procedure for SCell modification ($5G_4 \rightarrow 5G_5$), which fails again. As a consequence, this RRC connection is released and later re-established again at t = 43.9 s. We find that 5G repeatedly turns ON then OFF in this persistent loop as long as the SCell modification $5G_4 \rightarrow 5G_5$ is invoked. Such SCell modification fails every time, resulting in losing 5G radio access and high data speed (about 200 Mbps or higher).

We note that persistent loops have been reported in prior studies [17, 18, 34, 37]. They observed handover loops which switch back and forth among several choices of serving cells in 3G/4G/5G networks due to problematic parameter configuration. Our work differs from them. More RRC procedures other than handovers are involved. Moreover, we present the first loop type observed in 5G SA, which has not been studied before. In this work, we focus on 5G ON-OFF loops and conduct extensive measurements with all three US 5G operators to unveil and analyze distinct ON-OFF loops.

4 Reality Check: Loops in the Wild

In this section, we characterize how likely and how often one ON-OFF loops occur in the wild. We find that, surprisingly, these loops are not corner cases but quite commonly observed in operational 5G networks, resulting in significant performance degradation and even service disruption.

4.1 Measurement Methodology

We conduct extensive measurement experiments with all three major operators (OP_T, OP_A and OP_V) in 11 test areas in two U.S. cities – West Lafayette, IN (C1) and Lafayette, IN (C2). C1 and C2 are two small Midwestern cities with the population of roughly 100,000 (including students) and 70,000. Figure 5 gives the maps of 11 test areas: A1– A5 (OP_T), A6 – A8 (OP_A) and A9 – A11 (OP_V). These test areas cover Purdue campus, downtown and residential zones, with area sizes ranging from 1 Km^2 to 2.9 Km^2.

In each test area, we randomly select multiple locations to check whether a 5G ON-OFF loop is observed. At each test location, we use one 5G phone to conduct a few stationary experiments (runs) at different hours of the day or on different days. Each run is a 5-minute speed test which is a bulk file download (500 MB each) from a remote server. We use tcpdump to capture traffic packets and record data throughput. We use Network Signal Guru [2] to collect signaling messages and extract key information including serving cells over time, as well as their RSRP/RSRQ measurements and the involved RRC commands and configuration parameters.

For each run, we extract the sequence of serving cell sets (see Appendix B for details). There are two types of sequences: (I) *no loop*, and (II) *loop* (Figure 4). A loop occurs if one subsequence (here, $\{CS_k, CS_{k+1}, \cdots, CS_{k+x}\}$) is repeatedly observed twice or more. The subsequence starts with 5G ON ($\{CS_k\}$), ends with 5G OFF ($\{CS_{k+x}\}$), and repeats. The loop is a persistent one if it ends in this loop (namely, no new cell sets out of the loop subsequence). Otherwise, it is a semi-persistent loop. In this study, we observe that persistent loops are dominant (§4.2).

Note that for this section, we perform a **sparse (coarse-grained)** spatial measurement study where all the selected locations are apart from each other and cover the whole area (e.g., see all test locations in A1 in Figure 7). This is different from a **dense (finer-grained)** spatial measurement which is later used to examine how likely an identified loop occurs at locations in close proximity (§6). Spatial correlation means that a loop observed at a certain location highly likely occurs at nearby locations (within several tens of meters and up to two hundred meters). To fairly characterize how 5G ON-OFF loops occur in a certain area, we avoid close locations which are likely impacted by such spatial correlation. We later check all the loop instances observed at these test locations and confirm that they are indeed independent per location.

Operator	OP_T	OP_A	OP_V
Period	December 26, 2024 – May 9, 2025		
City	West Lafayette (C1), Lafayette (C2)		
Area	A1 – A5	A6 – A8	A9 – A11
Area Size	9.7 Km^2	4.4 Km^2	5 Km^2
♯ Location	46	28	28
Total time	7,445 min	1,768 min	1,821 min
5G mode	**5G SA**	**5G NSA**	**5G NSA**
5G band	n25, n41, n71	n5, n77	n77
4G band	2, 12, 66	2, 12, 17, 30, 66	2, 5, 13, 66
♯ 5G/4G cell	**242**/113	**129**/386	**78**/365
♯ RSRP/RSRQ	27.5M	7.1M	10.9M
♯ CS sample	31,204	7,552	7,314
♯ CS (unique)	2,106	2,324	1,418
♯ ON-OFF loop	**1,353**	**397**	**332**

Table 3: Statistics of our basic dataset.

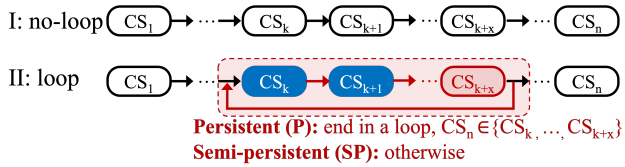

Figure 4: Two possible forms of serving cell set sequences: (I) no loop, (II) loop (persistent and semi-persistent).

We first conduct experiments with OP_T in one showcase campus area A1 (Figure 7). We randomly select 25 locations which are at least 200 meters apart (mostly > 300 m). For each location, we repeat at least 10 runs (up to 20 runs and more runs are added in case a loop is rarely or never observed). Given the findings in A1, we extend our measurement study to a larger-scale in more test areas with all three operators. In the remaining areas, we choose 5 – 10 locations per area with at least 5 runs per location (mostly 10 runs and up to 20 runs, depending on how likely a loop is observed at the test location). Note that our experiment scale in A2 – A11 is slightly smaller than in A1; We believe that such experiment scale is acceptable because consistent results are observed.

We also run driving tests in addition to stationary experiments because RRC procedures and key parameters configured by the PCell play an essential role in cause analysis (§5). To collect such information from all 5G/4G cells deployed in the test areas, we conduct driving experiments along all main roads until no new 5G/4G cells are observed.

Table 3 shows basic statistics of a five-month measurement study (sporadically from December 2024 to May 2025). We used OnePlus 12R, a 5G phone model released in Feb 2024 which supports advanced 5G features used by all three US operators. We later ran more experiments with six phone models (§4.4) and observe that the results are consistent across all phone models over 5G NSA (with OP_A and OP_V), but the results over 5G SA (with OP_T) are device-specific (loops are observed only with OnePlus 12R, which will be explained in §4.4). Moreover, we also ran extra experiments at dense locations in close proximity for a finer-grained spatial analysis (§6). Table 3 covers our experiments for a sparse spatial analysis, at 102 sparse locations (OP_T: 46, OP_A: 28, OP_V: 28) in 11

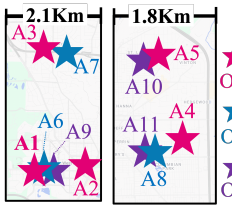

Figure 5: C1/C2 map.

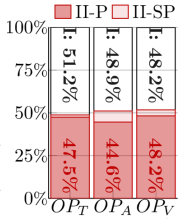

Figure 6: Ratio (%).

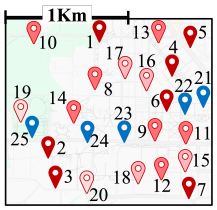

Figure 7: Map of A1.

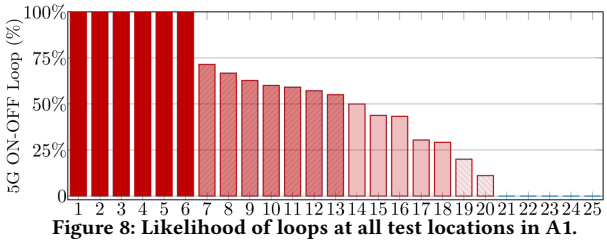

Figure 8: Likelihood of loops at all test locations in A1.

test areas (over 19 Km^2, > 11,000 minutes). All the experiments used OnePlus 12R unless specified. OP_T primarily supports 5G SA while OP_A and OP_V use 5G NSA. This indicates that OP_T has progressed more in 5G deployment than OP_A and OP_V. This matches reported results [25]. This is also evident from more 5G bands (here, n25, n41, n71) and more 5G cells than 4G cells used by OP_T. In contrast, OP_A and OP_V still largely count on 4G, with more bands for 4G and more 4G cells. In total, we observe more than 2,000 loop instances with 449 unique cells over 5G (OP_T: 242, OP_A: 129, OP_V: 78) and about 4.4K unique cell sets. We collect around 46K cell set samples and about 45M RSRP/RSRQ measurement samples.

4.2 Loops Commonly and Widely Observed

Our reality check shows that 5G ON-OFF loops are not uncommon. We have two related findings (**F1** and **F2**).

[F1] *5G ON-OFF loops occur much more often than anticipated. Moreover, they are mostly persistent.*

At a given location, we check whether a 5G ON-OFF loop occurs in each stationary run, and then calculate the ratio of runs without loops (I) and with loops (II-P, II-SP), as illustrated in Figure 4. We use all runs in stationary experiments at all sparse test locations and Figure 6 plots the percentage of no-loop and loop runs per operator. All loops are 5G ON-OFF ones unless specified. Surprisingly, loops are observed in around half of the runs with all three operators (OP_T: 48.8%, OP_A: 51.1%, OP_V: 51.7%). Remarkably, almost all loops are persistent (II-P). We rarely see semi-persistent loops (II-SP) with OP_T; For OP_A and OP_V, there is a small chance of exiting semi-persistent loops (OP_A: 6.5% and OP_V: 3.5%). In §5.3, we will see that these semi-persistent loops stem from the same causes as persistent loops, and exiting the loop occasionally happens due to runtime RSRP/RSRQ measurement dynamics, attributed to radio deployment by OP_A and OP_V. As a result, we do not distinguish between persistent and semi-persistent loops unless specified.

[F2] *5G ON-OFF loops are widely observed at a large portion of test locations. Loops are consistently observed with all operators in all test areas, despite operator- and area-specific heterogeneity.*

Showcase: A1 with OP_T. Figure 8 shows the likelihood of loops at 25 test locations in A1 (see the map in Figure 7). At each location, the likelihood is calculated as the ratio of the number of 5G ON-OFF loops over the total number of CS sequences. In A1, 5G ON-OFF loops are observed at 20 out of 25 test locations. Moreover, the loops are observed in more than half of the runs (likelihood > 50%) at 13 locations, and are observed every time at 6 locations (P1 - P6).

In more areas with three operators. Figure 9 plots the results with three operators in all test areas: A1 - A5 (OP_T), A6 - A8 (OP_A), A9 - A11 (OP_V). Figure 9a shows the loop ratio per area. Similar to the aggregated results in Figure 6, we see that loops

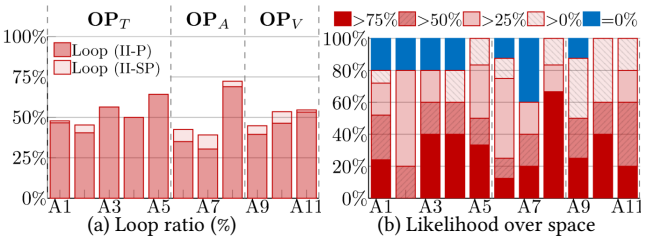

(a) Loop ratio (%) (b) Likelihood over space

Figure 9: 5G ON-OFF loop ratios in all test areas.

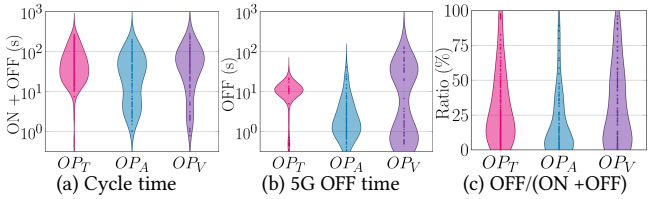

(a) Cycle time (b) 5G OFF time (c) OFF/(ON +OFF)

Figure 10: 5G OFF time impacts per operator.

occur quite often in all test areas (**F1**). As expected, the results vary across geographical areas, primarily due to real-world deployment heterogeneity and locations with poor radio conditions (§5.3).

We now quantify the likelihood of loops over space. In each area, we calculate the portion of test locations with loop likelihood γ within certain ranges. Figure 9b plots the breakdown with γ in four quartiles of [100%, 75%), [75%, 50%), [50%, 25%), [25%, 0%) and $\gamma = 0\%$ (no loops observed). Clearly, 5G ON-OFF loops are widely observed at a large portion of locations. Loops are observed at more than 80% of test locations in all areas except A7. The loop likelihood exceeds 50% at more than half of the locations in 8 out of 11 areas (except A2, A6, A7). For each operator, we observe area-specific diversity, particularly all test areas (A6 – A8) for OP_A and the *"outlier"* (A2) for OP_T. This is because the dominant loop type changes in these areas (Figure 16, §5.3).

4.3 Performance Impacts of ON-OFF Loops

5G ON-OFF loops significantly hurt data performance. We assess their performance impacts in terms of 5G OFF time ratio (**F3**) and data speed loss during 5G OFF (**F4**).

[F3] *5G ON-OFF loops occur quite often (every several tens of seconds) with a noticeable portion of 5G OFF time period.*

5G OFF time and ratio. We extract every ON-OFF cycle in all the loop instances and show the violin plots of the cycle (ON +OFF) time, the OFF time and the OFF/(ON +OFF) ratio in Figure 10. The results are merged per operator because the results in different areas are consistent unless specified in §5.3. We make three observations. First, loops occur quite frequently, mostly every several tens of seconds; the median cycle time is 41s, 26s and 49s for OP_T, OP_A

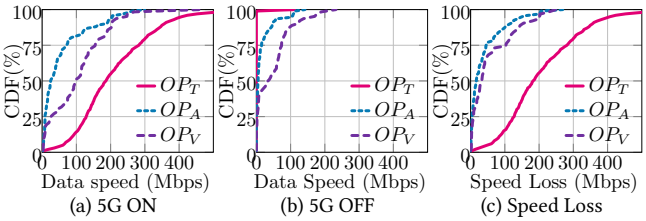

(a) 5G ON (b) 5G OFF (c) Speed Loss

Figure 11: CDFs of the median download data speed during 5G ON/OFF periods and data speed loss caused by 5G OFF.

and OP_V, respectively. Second, both OP_T and OP_V experience a substantial portion of 5G OFF time (> 22% for more than half of the loop instances). OP_A is least impacted and its OFF time period takes >7.4% in half of the loop instances. Third, impacted time varies across operators. For OP_T, its OFF time is mostly within 10 – 15 seconds but the entire cycle has a wide span. For OP_A, its OFF time is much shorter (mostly below 5 seconds). For OP_V, its OFF time is distributed in two ranges: (1) below 5 seconds, similar to OP_A (both running 5G NSA), and (2) around 30 seconds, which is attributed to one special loop sub-type (N2E2 with OP_V, §5.3). In the next section, we will demonstrate that operator-specific impacts are rooted in their distinct loop types (§5.3).

[F4] *5G OFF significantly hurts data performance and may even suspend data services. Performance loss varies across operators.*

Data speed loss during 5G OFF. We measure the median download data speed during 5G ON and 5G OFF periods in each ON-OFF cycle. Figure 11 plots the cumulative distribution functions (CDFs) of the download speed when 5G is ON/OFF, as well as data speed loss. Clearly, OP_T loses much more data speed once 5G is OFF. This is partly because OP_T offers the fastest 5G access (Figure 11a, with a median speed of 186.1Mbps, much higher than 24.9 Mbps and 97.5Mbps with OP_A and OP_V); Additionally, data service is almost suspended when 5G is OFF (Figure 11b); this is attributed to its loop type which becomes IDLE when 5G is OFF (§5.1). In contrast, data performance degrades less with OP_A and OP_V. This is mainly because 4G is still used in some loop instances when 5G is not used.

4.4 Across Phone Models

We now present additional experiments with different phone models to validate whether the observed 5G ON-OFF loops occur across phone models. We ran extra experiments with six phone models (OnePlus 13/13R/12R/10 Pro, Samsung Galaxy S23 Ultra and Google Pixel 5, Table 4) in three weeks from August 20 to September 9, 2025. For each phone model, we ran the same stationary experiment (namely, a bulk file download for at least five minutes) in each run. For each operator, we choose five test locations from our earlier measurement study and repeat at least five runs at each test location per phone model. We find that 5G ON-OFF loops with OP_A and OP_V are observed with almost all phone models (**F5**) but loops with OP_T are observed with OnePlus 12R only (**F6**).

[F5] *5G ON-OFF loops over 5G NSA are consistently observed across all test phone models (except OnePlus 10 Pro with OP_A).*

Figure 12 shows the results across all test phone models with OP_A and OP_V. 5G ON-OFF loops are frequently observed across

Phone Model	Release	Chipset (all by Qualcomm)	Android	3GPP
OnePlus 13R	Jan 2025	SM8650-AB Snapdragon 8 Gen 3	Android 15	V17.4.0
OnePlus 13	Oct 2024	SM8750-AB Snapdragon 8 Elite	Android 15	V17.4.0
OnePlus 12R	Feb 2024	SM8550-AB Snapdragon 8 Gen 2	Android 14	V16.6.0
OnePlus 10 Pro	Jan 2022	SM8450 Snapdragon 8 Gen 1	Android 12	V16.3.1
Samsung S23	Feb 2023	SM8550-AC Snapdragon 8 Gen 2	Android 15	-
Google Pixel 5	Sep 2020	SM7250 Snapdragon 765G	Android 11	V15.9.0

Table 4: Key specifications of all test phone models.

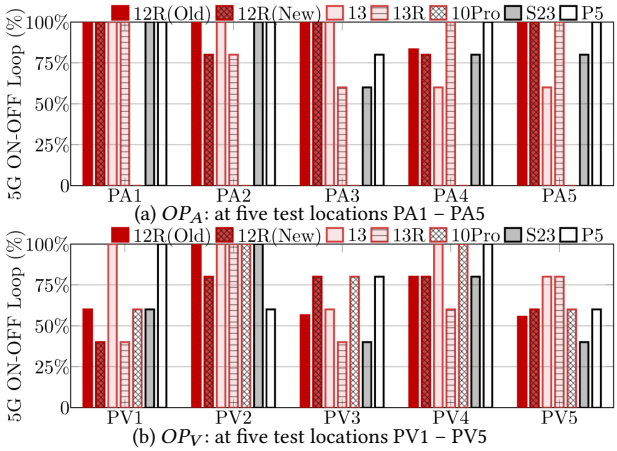

(a) OP_A: at five test locations PA1 – PA5

(b) OP_V: at five test locations PV1 – PV5

Figure 12: 5G ON-OFF loops across six phone models over 5G NSA.

all phone models (except OnePlus 10 Pro with OP_A). First, we compare the loop ratios using OnePlus 12R in old experiments (before May 2025) and new experiments (in Aug/Sep 2025). The loops are consistently observed at all test locations (with small variance impacted by the experimentation scale). Second, similar loop ratios are observed across all other phone models. The loop ratios at all the test locations exceed 40% (> 2 loop runs out of 5+ runs) using any test phone model (except OnePlus 10 Pro with OP_A). Third, we find that the exception happens because OnePlus 10 Pro uses 4G only (not 5G) with OP_A, which was reported by some AT&T users [1]. Interestingly, we observe new 5G ON-OFF loops with OP_T at certain locations in city C2, where OP_T supports 5G NSA, not 5G SA. These imply that 5G ON-OFF loops are not likely caused by device-side issues but highly likely stem from certain structural issues in common practice for 5G NSA (elaborated in §5.2).

[F6] *5G ON-OFF loops over 5G SA are observed only with OnePlus 12R because 5G connectivity varies with device models in our study.*

First, we observe 5G ON-OFF loops over 5G SA (with OP_T) with OnePlus 12R only. At all five test locations with 5G ON-OFF loops in old experiments, we observe the same loops in new experiments using OnePlus 12R phones. We tested with three OnePlus 12R phones and the findings are consistent. This demonstrates that observed loops are not specific to a certain device but are associated with at least this phone model (OnePlus 12R).

Second, we further analyze why the loops are not observed with other phone models. We find that all observed loops are associated with certain 5G SCells (elaborated in §5.1); however, these "problematic" SCells which result in 5G ON-OFF loops using OnePlus 12R are not used by other phone models. There are three cases which vary with phone models. (1) Mobile phones do not use any 5G SCell. In our study, both OnePlus 10 Pro and Pixel 5 phones, which are early 5G phone models (released before January 2022), support 5G

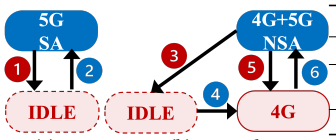

(a) S1 (b) N1 and N2

	5G	FSM	Sub-types	Trigger for 5G OFF	Trigger for 5G ON
S1	SA	❶↔❷	S1E1, S1E2, S1E3	One/few SCells "fail" ⇒ 5G MCG released (**F9**)	Good 5G candidate cells available and found (meeting the RSRP/RSRQ criteria)
N1	NSA	❸→❹→❻	N1E1, N1E2	4G PCell "fails" ⇒ 5G SCG released (**F10**)	
N2	NSA	❺↔❻	N2E1, N2E2	4G PCell handover, 5G SCG failure	

(c) Loop sub-types, triggers and causes

Figure 13: Three 5G ON-OFF loop types with their sub-types, triggers for 5G ON and 5G OFF and causes.

SA but do not support carrier aggregation over 5G SA, which is an advanced feature. These test phones use only one 5G PCell. (2) Mobile phones do not use 5G SCells which result in 5G ON-OFF loops. For instance, $5G_4$ and $5G_5$ over channel 387410 (band n25) are two "problematic" cells which result in the loop in the example (§3). We observe that some phones (here, OnePlus 13R) do not use these SCells. More specifically, compared to four 5G cells (1 PCell and 3 SCells) used by OnePlus 12R, 13R uses only two 5G cells (the same PCell and 1 SCell) and both using 4×4 MIMO. In contrast, the four cells used by 12R use 2×2 MIMO. This is mainly because OnePlus 13R supports more advanced physical features so that OP_T uses two serving cells for 13R phones. By comparing RRC signaling messages collected by Network Signal Guru, we further find that 13R runs a newer RRC release (`V17.4.0`, compared to `V16.6.0` used by 12R) [9]. We find that 12R phones receive only downlink configuration when adding any SCells over band n25 but 13R phones receive both uplink and downlink configuration. 13R phones do use additional parameters as SCell traffic feedback. We gauge that these changes might help 13R phones without using the problematic SCells and avoiding the ON-OFF switch. (3) Mobile phones do not use "problematic" 5G SCells. Here, OnePlus 13 and Samsung S23 Ultra phones are not supported by Network Signal Guru [2] and we cannot obtain RRC signaling messages to investigate why. We do observe that the 5G serving cells are different from those used by 12R phones at the same location. For example, the Samsung S23 phone uses the PCell over over band n71, different from the one used by 12R (over n41). As a result, the used SCells are different and any ON-OFF switch with OnePlus 12R is not observed.

Finally, we point out that loops are not likely caused by an implementation bug with OnePlus 12R, even though loops are observed only with this phone model. The loops depend on the selection of serving cells which is realized by RRC procedures. 5G SA uses more complex RRC policies/configurations and advanced physical technologies. The selection of serving cells is diversified (see different sets of serving cells with different phone models at the same locations). While device-specific factors impact the selection of serving cells over 5G SA, RRC policies/configurations are also responsible for "improper" cell selection. In the remainder of the paper, all loops discussed are with OnePlus 12R unless otherwise specified.

5 An In-Depth Look at Loop Causes

In this section, we conduct an in-depth analysis to understand why loops occur in operational 5G networks. We divide loop instances into three types based on their 5G ON/OFF state transitions (**F7**).

[F7] *Three 5G ON-OFF loop types - one for 5G SA (S1) and two for 5G NSA (N1 and N2) - are observed in the wild.*

We dive into each type and find that all loops share one common cause: *the triggers to turn 5G ON and OFF can co-exist under (quasi-) same network conditions.* Specifically, 5G cells are added back because good candidate cells are available and found for use because their RSRP/RSRQ measurements meet the pre-configured criteria. Releasing 5G cells (turning 5G OFF) is not always associated with RSRP/RSRQ measurement events. Loops are created with inconsistent ON/OFF triggering conditions (**F8**).

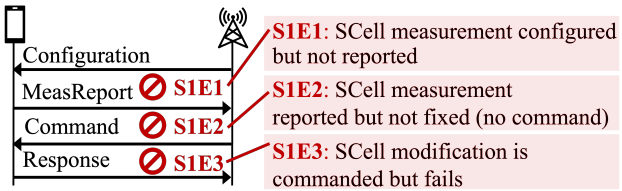

Figure 14: Three sub-types of S1E1, S1E2, S1E3 with unanticipated triggers for 5G OFF during the normal RRC operations.

[F8] *Inconsistent ON/OFF triggers create loops.*

We next elaborate on how this common problem (**F8**) occurs in practice. Figure 13 shows all three loop types each with several sub-types and distinct triggering events (causes). Due to space limits, we give real-world loop instances per sub-type in Appendix C and only highlight their key triggers in this section.

5.1 S1: 5G SA ⇔ IDLE

S1 is an ON-OFF loop type for 5G SA. In this study, all S1 loop instances are only observed with OP_T. As illustrated in Figure 13a, 5G ON means that 5G SA is used (with one 5G PCell and 5G MCG); 5G OFF means that 5G SA is not used and the state is IDLE (the PCell along with 5G MCG is released). As there is no valid PCell responsible for RRC, this RRC connection becomes IDLE without any serving cell.

S1 loops stem from the triggers to turn 5G OFF, namely, the 5G SA → IDLE transition (❶). Good 5G cells must be available at a location; otherwise, 5G SA cannot be recovered later. The problem lies in why 5G turns OFF despite the presence of good 5G cells (**F9**).

[F9] "A few bad apples ruin all". *Losing 5G radio access in S1 loops is not because all the 5G serving cells perform poorly but because one or a few 5G cells are problematic. When RRC fails to see/use/modify one or a few 5G SCells, 5G MCG is released as a whole. There are three sub-types with distinct triggering events: S1E1, S1E2 and S1E3.*

We observe that all the S1 loop instances release the whole 5G MCG (turn 5G off) upon an"exception" or "failure" with one or a few 5G SCells. Figure 14 shows three loop sub-types (S1E1, S1E2 and S1E3), each with its trigger event during normal RRC operation. An active RRC connection keeps running the following four steps: *configuration - measurement and reporting - command - command response* [9]. The PCell sends its *configuration* parameters to customize cell-specific operation criteria including whether and what to measure and report. If the configured criteria are satisfied,

the device measures nearby cells and reports their RSRP/RSRQ measurements. Upon receiving the measurement reports, the PCell runs its local logic to determine whether and how to change the serving cell(s) and update configurations. If applicable, the PCell sends a command to execute the change and the device finally responds with the execution status. For each loop sub-type, we highlight its triggering event to turn 5G off, which eventually results in a persistent 5G ON-OFF loop.

∘ *S1E1: no RSRP/RSRQ measurements of one or more 5G SCells.* We find that all the 5G serving cells are released as long as the RSRP/RSRQ measurements of one 5G SCell is not received in the measurement and reporting step. This is because if a serving cell provides good radio coverage, its RSRP/RSRQ measurement should be reported. Hence, not receiving RSRP/RSRQ measurements of a serving cell is a reasonable indicator of its poor radio quality. Here, the problem is that RRC should not handle one/few bad apples (poor radio coverage of one/few 5G SCells) by releasing the whole group (MCG) of all the serving 5G cells.

∘ *S1E2: poor RSRP/RSRQ measurements of one or more 5G SCells.* S1E2 is similar to S1E1. The only difference is that RSRP/RSRQ measurements of all the serving SCells are reported but the RSRP/RSRQ results of at least one 5G SCell (bad apple) are really poor (for instance, RSRQ ≈ -25 dB for one 5G SCell 390@3874010 in Figure 28). No commands are further observed upon receiving the measurement reports of all the serving cells. Similar to S1E1, S1E2 creates a problem by releasing all serving cells.

∘ *S1E3: SCell modification failure.* S1E3 occurs when a command to modify one or a few SCells is issued but fails. Specifically, one `RRC Reconfiguration` message is sent to modify 5G SCells upon receiving the RSRP/RSRQ measurements of one candidate 5G cell which has a stronger radio coverage than one of the serving SCell [9] (for example, $5G_4 \rightarrow 5G_5$ (273@387410 → 371@387410) in Figure 3). In S1E3 instances, the ON-OFF cycle is repeated as follows: 5G PCell first adds several 5G SCells into 5G MCG, finds a better candidate cell but loses all 5G radio access when attempting to replace it with a 5G SCell, and finally re-establishes 5G connection with the same 5G PCell.

Note that in all S1 instances, 5G is later turned on every time because an RRC connection is successfully re-established (❷) with good 5G cells available for use. Hence, releasing the whole 5G MCG to handle one or a few "problematic" SCells creates a loop and **a few bad apples ruin all** (**F9**). More interestingly, we see that almost all the bad apples use one 5G channel 387410 (**F14**, §5.3). This is why we do not see S1 loops with other phone models because they do not use the bad apples at the same locations as OnePlus 12R does.

5.2 5G NSA: N1 and N2 Loops

We observe two ON-OFF loop types with 5G NSA. Figure 13b plots their state transitions:

∘ N1: 5G NSA ⇔ IDLE* (IDLE+4G), via ❸→❹→❻→❸···

∘ N2: 5G NSA ⇔ 4G, via ❺↔❻.

In this work, all the N1 and N2 loop instances are observed with OP_A and OP_V. Compared to 5G SA, 5G NSA uses 4G as its master radio access and 5G as its secondary radio access. This means that losing 5G radio access does not necessarily release the RRC connection, which is different from 5G SA. There are two 5G OFF cases: (N1) losing both 4G and 5G (IDLE, via ❸) and (N2) losing 5G only (via ❺). To get 5G radio access back, both use `RRC Connection Reconfiguration` to add 5G serving cells (❻). Procedure ❻ must be performed with one valid 4G PCell and thus all N1 instances must first establish an active RRC connection (❹) once IDLE. That is, in all N1 instances, the state quickly switches from IDLE to 4G only and then to 5G NSA. We use IDLE* to represent IDLE plus transient 4G. We further analyze all loop instances and identify four sub-types with distinct triggering events to turn 5G off, as depicted in Figure 15.

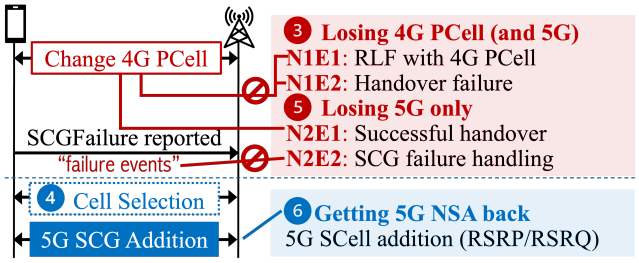

Figure 15: Four loop sub-types: N1E1, N1E2, N2E1, N2E2, with their distinct triggering events to turn 5G off (lose 5G NSA).

We observe that good 5G cells are available for use in this case. Otherwise, 5G radio access could not be recovered later. Similar to S1 loops (for 5G SA), we are interested in why the UE repeatedly loses 5G radio access in the presence of good 5G cells. It turns out that 5G is innocently blocked due to 4G failure (**F10**) and inconsistent conditions (policies and mechanisms) involved in normal RRC procedures (**F11**).

[F10] "4G ruins 5G". *Losing 5G radio access is not because 5G SCG performs poorly but because 4G MCG does. More precisely, failing to change 4G PCell results in losing all 5G serving cells.*

This problem is rooted in RRC mechanisms specified by 3GPP [6]. 4G PCell is the main control node responsible for RRC. As a consequence, RRC has to release all other serving cells (including 5G SCG) once it loses 4G PCell. In this study, we observe two failure cases (N1E1 and N1E2).

∘ *N1E1 (4G PCell radio link failure) and N1E2 (4G PCell handover failure).* Although all 5G serving cells are performing well in these cases, 4G PCell "fails" to perform well. In N1E1, the serving PCell experiences a radio link failure (RLF) and thus has to release both 4G and 5G radio access (see an instance in Figure 30). Unlike the case of N1E1, N1E2 receives a normal handover request to change 4G PCell but fails to complete this handover procedure (see an instance in Figure 31). As a result, 5G radio access turns off until a good 4G PCell is recovered.

We next analyze N2 loops which switch back and forth between 4G+5G (5G ON) and 4G (5G OFF). Interestingly, inconsistency in policies and mechanisms changes with new forms over time, but inconsistency itself never disappears. In this study, we observe two inconsistency forms (N2E1 and N2E2, **F11**) and also find that another parameter inconsistency identified in prior work is not observed any longer (**F12**).

[F11] "Inconsistency in policies and mechanisms". *Repeatedly losing 5G radio access stems from inconsistent policies and mechanisms to turn 5G on and off, particularly for 4G PCell handover (N2E1) and SCG failure handling (N2E2).*

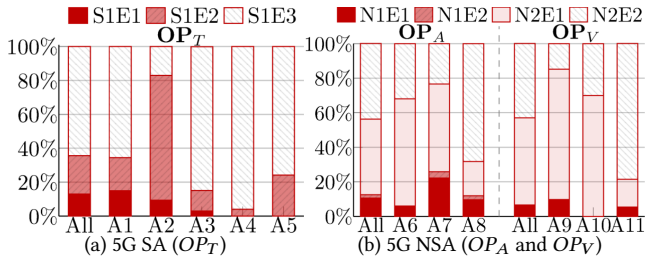

(a) 5G SA (OP_T) (b) 5G NSA (OP_A and OP_V)

Figure 16: Loop breakdown in all areas with three operators.

∘ *N2E1: Repeated 4G PCell handovers.* Compared to N1E2, N2E1 successfully changes the 4G PCell (through a successful handover). Surprisingly, we find that 4G handovers are unstable in stationary experiments. The serving PCell switches back and forth because of inconsistent policies in selecting one PCell (**F8**). We find that network operators employ cell-specific policy constraints that block or discourage 5G use with certain 4G cells (over one 4G channel: channel 5815 by OP_A, in our experiments **F14**). One 4G cell over this "5G-disabled" channel (380@5815 in Figure 32) is preferred in a handover procedure because its RSRQ is stronger than another candidate cell (event A3), but this cell is not allowed to work with 5G cells so that it quickly switches to another cell (380@5415 in this instance) once the RSRP/RSRQ of 5G candidate cells are reported. However, 4G PCell later switches back to the original cell over the"5G-disabled" channel (here, 380@5815) because its RSRQ measurement meets the handover criterion (offset stronger than 380@5415, event A3). Such inconsistency in PCell selection policies results in repeated 4G handovers, as well as 5G ON-OFF loops.

∘ *N2E2: Repeated 5G SCG failure handling.* We also observe 5G⇔4G loops due to repeated SCG failure handling. These loops stem from inconsistent mechanisms to turn 5G ON and OFF. Turning 5G OFF is triggered by a radio link failure (RLF) event with 5G serving cells, detected with lower-layer events (pre-defined by 3GPP [6]) such as random access failures, maximum number of retransmissions reached in a short time interval, timer expiration, and synchronization errors. However, turning 5G ON is realized by 5G SCell addition triggered by an RSRP/RSRQ measurement reporting event (here, event B1, the RSRP/RSRQ of one 5G candidate cell becomes better than a pre-defined threshold). Due to the distinct nature of the two mechanisms, it is almost impossible to avoid inconsistent conditions. In some real-world settings where both conditions to turn 5G OFF and ON are satisfied, an ON-OFF loop occurs. We observe that N2E2 loops are reported in a recent study (using 5G measurement in 2023) [18], but unlike [18], N2E2 loops are not rare in our study (§5.3).

We also note that there may exist other inconsistent conditions that result in new 5G ON-OFF loops. A prior study [37] reported another N2 loop sub-type caused by inconsistent A2-B1 parameters, where 5G candidate cells are added as serving cells when their RSRP/RSRQ measurements become better than a pre-defined threshold Θ_{B1} (B1 event) and later 5G serving cells are released when their RSRP/RSRQ becomes worse than another pre-defined threshold Θ_{A2} (A2 event). As long as $\Theta_{B1} < \Theta_{A2}$, loops may occur when the RSRP/RSRQ measurements of 5G cells are between these two thresholds. Interestingly, such A2-B1 loops which were observed in early 5G measurement studies (in 2021 – 2023)[37] are not encountered in our study because the thresholds used for events A2 and B1 have been changed by network operators (OP_A and OP_V). All loop sub-types identified in this study except N2E2 are new and have never been reported before.

[F12] *One N2 loop sub-type reported in prior work [37] (due to inconsistent A2-B1 parameters) is not observed in this study (thresholds corrected by mobile network operators and/or vendors).*

5.3 Loop Breakdown

We now present the breakdown among loops observed in our study and explore why they occur in practice.

[F13] *S1E3 is dominant for 5G SA (with an exception in A2), while N2 (N2E1 and N2E2) is dominant for 5G NSA.*

Figure 16 plots the loop breakdown per operator and per area. Overall, S1E3 is the dominant loop sub-type for 5G SA (OP_T) while N2 is dominant for 5G NSA (OP_A and OP_V). S1E3 accounts for 64.4% of loop instances, compared to 13.0% (S1E1) and 22.6% (S1E2). The results are consistent in all test areas except in A2; S1E2 is widely observed in A2 due to distinct (much worse) 5G radio coverage in A2 (Figure 17b). S1E1 and S1E2 are more likely observed at places where "bad apples" have worse 5G radio coverage, based on their 5G OFF triggers in §5.1.

N1 is rarely observed with OP_A and OP_V. Moreover, N1E2 is not observed with OP_V in our study. This is expected since 4G is well-provisioned by both operators. Failing to change the 4G PCell is rare (§5.2). This explains why data performance degrades less with OP_A and OP_V than with OP_T (**F4**, Figure 11b). Both operators still use 4G, and do not completely turn all radio access off as with OP_T (5G SA). We see that the breakdown of N2E1 and N2E2 varies in three selected areas. N2E2 is more prevalent in A8 (OP_A) and A11 (OP_V) but N2E1 is more common in other areas. This is also attributed to 5G radio coverage in test areas. In A8 and A11, 5G radio coverage is significantly worse with a higher likelihood of SCG failure handling. Radio coverage diversity contributes to area-specific heterogeneity (**F2**, Figure 9).

We next explore the reasons for the loops. Based on the aforementioned cause analysis, every loop instance is centered on its "problematic" serving cell which invokes one RRC procedure to turn 5G OFF. In all S1 instances, the problematic cell is one or a few "bad apples:" specifically, one 5G SCell which is never measured (too bad to be measured) (S1E1), or a cell with the poorest RSRP/RSRQ measurement out of all reported serving cells (S1E2), or a cell which is used to replace the existing serving cell but invokes a SCell modification failure. In all N1 instances, the problematic cell is the 4G PCell which experiences an RLF (N1E1) or a handover failure (N1E2). In N2E1 instances, the problematic cell is the 4G PCell out of its handover loop which does not work with 5G. In N2E2 instances, the problematic cell is the 5G cell that initially triggers an SCG failure but is later added back because its RSRQ/RSRQ is good enough for SCell addition. We analyze all the "problematic" cells in all the loop instances and investigate the reasons behind the problems.

We find that RRC policies and configurations are not cell-specific, but channel-specific. A network operator likely uses the same configuration for all the cells over the same channel in certain geographic areas. We thus conduct a channel analysis to check whether "problematic" cells operate over specific channels. Interestingly, we

channel	Usage breakdown					SCell modification failure ratio
	no-loop	loop	S1E1	S1E2	S1E3	
126270	1.6%	1.0%	0.0%	0.0%	1.6%	0.0%
387410	22.3%	77.1%	98.0%	56.5%	80.0%	12.3%
398410	21.0%	10.1%	1.0%	27.7%	5.8%	0.7%
501390	28.3%	5.6%	1.0%	6.2%	6.4%	0.7%
521310	26.8%	6.2%	0.0%	9.6%	6.2%	1.1%

Table 5: Usage and failure ratio per channel with OP_T.

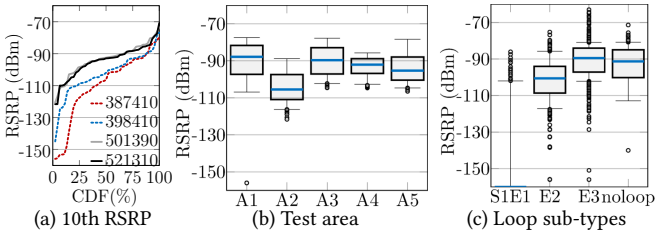

(a) 10th RSRP (b) Test area (c) Loop sub-types

Figure 17: RSRP measurements of cells on channel 387410.

find that every operator has one primary "problematic" channel (OP_T: 387410, OP_A: 5815 and OP_V: 5230) (**F14**).

[F14] *For OP_T, all the "problematic" cells primarily use one 5G channel (387410). For OP_A and OP_V, 4G channels (OP_A: 5815 and OP_V: 5230) are problematic.*

∘ OP_T. Interestingly, we find that almost all bad apples (which result in losing 5G access) operate over one specific channel 387410. Table 5 shows the usage breakdown per channel. Five 5G channels are used by OP_T and all the channels except 126270 are evenly observed in the no-loop instances. In contrast, only one channel 387410 is dominant (77.1%) in the loop instances. This is largely consistent across all three loop sub-types, but it is slightly different in S1E2: another channel 398410 (also on the same 5G band n25) contributes to about 25% of S1E2 instances. Later, we will show that radio quality of some cells over this channel is not good (Figure 17a). We further compare the failure ratio of SCell modification per channel used by the newly added SCell. Note that SCell modification failure is the triggering event to turn 5G OFF in S1E2 loops. Table 5 shows that the failure ratio over channel 387410 is significantly higher (12.3%), an order of magnitude higher, than those over other 5G channels. We attempt to understand why 387410 is the main problem. We do not find any channel-specific configurations and policies, which are the same as those over another channel (398410). The only difference is that the radio signal quality of 5G cells over channel 387410 is significantly worse than those on other channels (Figure 17). Figure 17a plots the CDF of the 10th percentile RSRP value across all test locations. Figure 17b and Figure 17c plot the median RSRP of the serving cells over channel 387410 per area and per type. RSRP measurements are significantly lower in A2, which contributes to many more S1E2 instances in A2 (Figure 16a). RSRP measurements are also much lower in S1E1 and S1E2 instances. This is not the case for S1E3. RSRP measurements are comparable in S1E3 and no-loop instances. This implies that SCell modification happens not because the RSRP/RSRQ of one SCell is poor, but likely because there exists a better candidate cell.

∘ OP_A *and* OP_V. We perform similar channel comparison for OP_A and OP_V. We focus on N2E1 and N2E2 loops.Both operators use many more channels for 4G than for 5G (OP_A: 3 and OP_V:2).

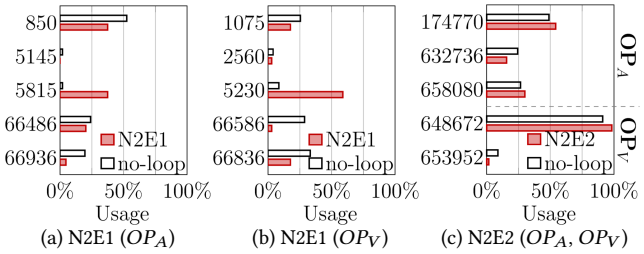

(a) N2E1 (OP_A) (b) N2E1 (OP_V) (c) N2E2 (OP_A, OP_V)

Figure 18: Usage breakdown per channel with OP_A and OP_V.

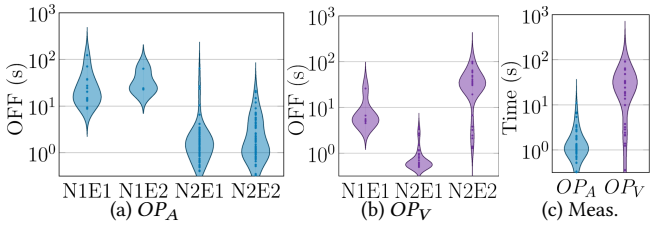

(a) OP_A (b) OP_V (c) Meas.

Figure 19: 5G OFF time varies with loop types (OP_A and OP_V).

Figure 18 plots the breakdown comparison for top-5 4G channels and all 5G channels. We see that two 4G channels stand out in N2E1 instances. They are rarely used in no-loop instances but account for a substantial portion of loop instances. Specifically, channel 5815 (centered on 742 MHz, band 17) is rarely used in 1.6% of no-loop instances but is responsible for about 40% of N2E1 instances with OP_A. Similarly, another 4G channel 5230 (centered on 753 MHz, band 13) used by OP_V accounts for more than half of N2E1 instances. We further examine all the PCell handovers when both channels are involved. We see that all 5G serving cells are released as long as 4G PCell switches to any cell over channel 5815 (OP_A) or 5230 (OP_V). This indicates that both operators employ policies to disable 5G use along with the use of certain 4G channels (**F15**).

[F15] *Both operators (OP_A and OP_V) employ certain policies that impact 5G use and the 5G OFF time.*

We observe a number of channel-specific policies in our study. Specifically, we find that OP_A uses the following two polices for channel 5815: (1) any 4G PCell over channel 5815 never works together with any 5G cell but it still allows 5G cell measurement (configured at the configuration step); (2) Upon receiving the RSRP/RSRQ measurement of any 5G cell, the 4G PCell over channel 5815 immediately switches to another cell over channel 5145 (with the same cell ID), despite no RSRP/RSRQ measurement of the new cell. As a result, 4G PCell handover may switch to another cell which is weaker than the old PCell over channel 5815 or even a cell that performs poorly (with handover failure or RLF). We find that such policies are responsible for all N1 (N1E1 and N1E2) instances with OP_A. However, OP_V uses different policies for channel 5230. The main difference is that channel 5230 is allowed to work with 5G, though all 5G cells are still released once the 4G PCell switches to the one over channel 5230. However, 5G can be quickly recovered through SCG addition as the PCell over this channel is allowed to use 5G. It results in transient 5G OFF in N2E1 instances with OP_V.

Figure 19 shows that the OFF time are distinct with both operators. Clearly, 5G OFF time is much shorter (mostly within 1 second, up to 5 seconds) with all N2E1 instances with OP_V. In contrast,

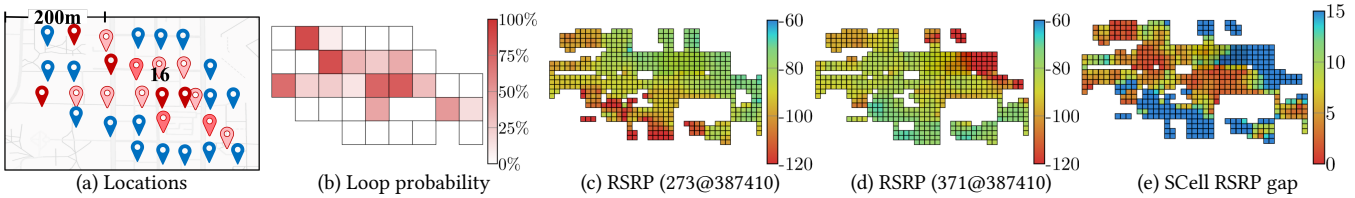

(a) Locations (b) Loop probability (c) RSRP (273@387410) (d) RSRP (371@387410) (e) SCell RSRP gap

Figure 20: The probability of one given S1E3 loop instance impacted by RSRP measurements of key cells.

5G OFF time is longer with OP_A due to different policies used by the operators. We make an interesting finding on 5G OFF time in N2E2 instances with OP_V. OP_V runs different policies to recover 5G SCG, which results in significantly longer 5G OFF durations compared to OP_A. For OP_A, 5G measurements are reported within 3 seconds after losing 5G in 90% of N2E2 instances. In contrast, for OP_V, 66% of N2E2 instances wait for more than 30 seconds before 5G measurements begin. We find that these delays are often multiples of 30 seconds (e.g., 30s, 60s, 90s). This is because OP_V employs the following policy and practice. The UE typically does not start measurement immediately with the initial configuration. Instead, it waits for the updated configuration information from the network. We see that OP_V sends this configuration every 30 seconds so it takes much longer for RSRP/RSRQ measurement (Figure 19c). As a result, the OFF time in N2E2 instances is much longer and 5G OFF time is more diverse (Figure 10b).

Unfortunately, we cannot determine why the two operators use such policies and practice. We assume that they want to advocate 4G usage in certain scenarios to manage their 4G and 5G usage. Some may be operational slips; for instance, OP_V does not guide the UE to perform 5G measurement as soon as 5G radio access is lost. However, such polices and practice slow down 5G access recovery, under-utilizing 5G. We also find that such policies and practice contribute to ON-OFF loop complexity. Compared to S1 loops in 5G SA, loops in 5G NSA are more impacted by runtime dynamics. For instance, a 4G PCell out of the loop might be selected because its time-changing RSRP/RSRQ is occasionally stronger than the preferred 4G PCell. This is why semi-persistent loops are observed with OP_A and OP_V.

6 A Final Look at RSRP/RSRQ Impacts

In this section, we attempt to quantify the impacts of runtime RSRP/RSRQ measurement fluctuations and environment dynamics on 5G ON-OFF loops. If technically feasible, this can help predict how likely a 5G ON-OFF loop occurs in the wild. Given one loop instance identified at one location, we start with a finer-grained spatial analysis to model how likely this identified loop occurs and varies in its close proximity. Considering that loop instances at nearby locations share the same structural factors (causes and triggers in §5), we examine the impacts of varying RSRP/RSRQ measurements on loop likelihood.

Showcase study around P16. We first use one showcase study to model how the loop probability changes in proximity, for a given loop observed at one location. Here, the loop probability is mainly affected by radio conditions that change over space. Different from sparse spatial analysis used for our reality check (§4), we run stationary experiments only at locations near the given site. We choose several loop instances and perform such finer-grained spatial measurements for every loop instance. Here, we use the first example at location P16 (Figure 1b) to show technical feasibility.

Figure 20 plots our finer-grained spatial analysis results in this showcase study. We run stationary experiments at over 30 locations near P16 and Figure 20b shows the measured probability of this loop instance (S1E3). The loop probability gradually drops to zero at the edge of this test region (see all blue nodes in Figure 20a and white grids in Figure 20b). We observe that location P16 is not the one with the highest loop likelihood. This S1E3 loop is more commonly observed at other locations. Test locations used in our reality check are more randomly selected to cover the whole test area. This S1E3 loop instance is caused by SCell modification failure between two SCells – 273 and 371 – both on channel 387410. We conduct walking experiments to collect RSRP/RSRQ measurements for both cells across the test region. Cell 273@387410 is stronger in the northwest corner while cell 371@387410 is stronger in the southeast. We observe strong spatial correlation where this loop likely occurs when the RSRPs of two involved SCells on channel 387410 are close (**F16**).

[F16] *The probability of S1E3 loops is high only at locations where the RSRP of SCells on channel 387410 are close.*

We plot the relationship between loop probability and the RSRP gap at each test location in this showcase study. At the locations where their RSRP gap between the two SCells is less than 6 dBm, the loop probability exceeds 50%. We further calculate their Spearman correlation coefficient which is -0.65. This confirms that the SCell RSRP gap is a key factor that impacts the loop probability.

Moreover, we find that a small RSRP gap between these two SCells is necessary but not sufficient for this S1E3 loop. This is because whether the target SCells (273 and 371 on channel 387410) are used depends on 5G PCell in use. We find that the target SCells (273 and 371) are used if and only if the PCell belongs to a specific group (here, cells 104 and 393 on channels 501390 and 521310). Therefore, the use of target SCells is solely determined by whether one target PCell is being used, which is further determined by the RSRP gap between target PCells and all other candidate cells on channel 501390 and 521310.

[F17] *The use of target SCells is determined by the target PCell, which influences the S1E3 loop probability.*

Figure 21b shows the target SCells are highly likely used when the target PCell has a stronger RSRP (with the RSRP gap larger than 6 dB). Note that matches with RRC configuration (the offset for event A3). This relationship follows a logistic-like curve: when the RSRP gap is very high or very low, the SCell usage ratio approaches 100% or 0%, respectively. When the RSRP gap is close to zero, the usage ratio of the target SCells is approximately 50%.

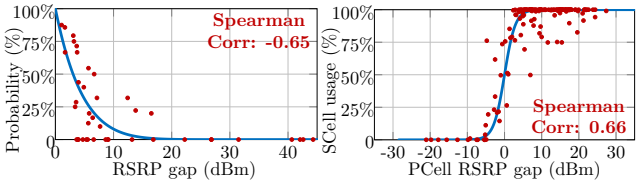

(a) SCell RSRP gap (b) PCell RSRP gap

Figure 21: Two impact factors of S1E3 loop probability.

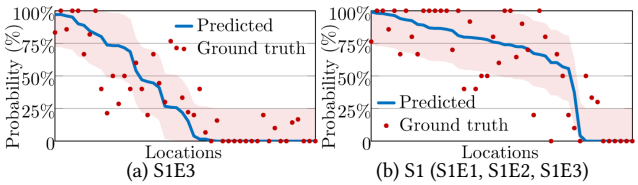

(a) S1E3 (b) S1 (S1E1, S1E2, S1E3)

Figure 22: Prediction results and ground truth of loop probability at all locations in coarse-grained analysis.

Extension to other S1E3 loop instances. We next extend the model learned in a finer-grained spatial analysis into a general one which predicts the S1E3 loop probability at any given location. We find that both findings (**F16** and **F17**) hold true for all S1E3 loop instances. This model uses RSRP measurements of all cells and available cellset combinations. We extend this model as follows.

First, for each possible cellset combination i, we calculate the RSRP gap between the target PCell and other candidate PCells, denoted as Δ_i^p, and predict the usage ratio of the cellset combination, denoted as u_i. Inspired by the results in our fine-grained spatial analysis (Figure 21a), we use a logistic function for loop prediction: $u_i = f_1(\Delta_i^p) = 1 \Big/ \left(1 + e^{-k\Delta_i^p}\right)$. Here, k is a learnable parameter.

Second, we calculate the RSRP gap Δ_i^s between target SCells, and use it to predict the probability p_i of S1E3 loops with this cellset combination. Driven by the relationship shown in Figure 21b, we construct the following function for modeling: $p_i = f_2(\Delta_i^s) = \max\left((1 - \Delta_i^s/t)^n, 0\right)$, where t and n are learnable parameters.

Finally, we combine the predicted usage ratio u_i and loop probability p_i of all cellset combinations to compute the overall loop probability P at each location as: $P = f(\Delta_i^p, \Delta_i^s) = \sum_i u_i\, p_i$. To train this model, we use the loop probability and RSRP data collected from our fine-grained spatial analysis. The parameters k, t, and n are optimized by minimizing the mean squared error (MSE) between the predicted and observed loop probabilities.

To evaluate the effectiveness of our prediction model, we apply it to predict the loop likelihood at all test locations used in our reality check (§4). We compare the predicted S1E3 loop probability at each location with the ground truth obtained from stationary experiments to assess the prediction accuracy. Figure 22 shows the predicted probabilities and most are within the ±25% error bounds. At more than half of locations, the prediction error is even below 10%. This demonstrates that it is promising to quantify the impacts of runtime RSRP/RSRQ measurements to predict the S1E3 loop probability at different locations (with different 5G SCells).

[F18] *Leveraging runtime RSRP/RSRQ measurements to predict other S1 loops is promising.*

Extension to other S1 loops. We find that the above prediction approach can be generalized for S1 loops but not for N1 and N2 loops. This is because we observe strong correlation with the RSRP features and the likelihood of S1E1 and S1E2 loops, but the correlation is noisy for N1 and N2 loops. The cause analysis shows that the impact of RSRP/RSRQ measurements on N1 and N2 loops are subtle (and complicated). Given the limited number of training samples, we are unable to support complicated models with many tunable parameters.

We extend the above loop prediction framework to other S1 loop types and predict the overall loop likelihood. To model the probability of S1E1 and S1E2, we replace one feature from the SCell RSRP gap (used for S1E3) to the RSRP of the worst SCell. For each cell set combination at each location, we independently compute the probabilities of S1E1, S1E2, and S1E3, and then aggregate them to obtain the overall probability of S1 loops. Figure 22b shows the prediction accuracy of the overall model. The gap between predicted probability and ground truth is below 25% and 30% at 67.4% and 82.6% of locations, respectively. This shows potential to generalize our loop prediction for all S1 loops. We can estimate how one S1 loop observed at one location occurs in proximity of another (e.g., while walking). This complements our sparse spatial analysis because measuring the loop likelihood at every location through many stationary runs is time consuming.

7 Discussion

Our work is at an early stage as it is the first measurement study to uncover and characterize 5G ON-OFF loops in the wild. There are a few open and remaining issues, as well as a number of limitations.

Device dependence. For both 5G deployment options (SA and NSA), 5G ON-OFF loops are device-dependent. Essentially, every loop is created by a combination of several procedures which impact the selection (addition and release) of 5G serving cells. All these procedures are involved, with step-by-step cooperation between the network and the device. Specifically, the measurement and reporting on the mobile device side is impacted by the configuration from the PCell, and the selection decision made by the PCell is impacted by the measurement results reported by the mobile device. In this work, we observe that the loops over 5G NSA (with OP_A and OP_V) are less device-dependent while those over 5G SA (with OP_T) are quite device-dependent. For OP_A and OP_V, all the test devices react similarly and the configuration from the network side plays a decisive role for the the loops. In contrast, for OP_T, device responses are different across test phone models, resulting in distinct sets of serving cells at the same location, which further impacts the RRC configurations and the selection of future serving cells. In our study, 5G ON-OFF loops with OP_T are observed only with one out of six phone models (OnePlus 12R). This implies that device heterogeneity multiplies the complexity of cell selection, at least for 5G SA used by OP_T. We plan to investigate the impact of devices and RRC configurations for 5G SA in our future work.

Location dependence. Not surprisingly, loops are location-dependent. Due to time limits, we were unable to test at more locations and in more cities in the US. Our most recent experiments in August - September 2025 included new locations, e.g., in the campus area A1, A6, A9. We observed 5G ON-OFF loops at new

locations. We omit these results because our test scale is insufficient (only 2 or 3 runs per location) to rigorously quantify and compare the loop likelihood. We observe that 5G ON-OFF loops are not limited to a few test locations, but likely to be observed at many locations. A larger-scale study is required to rigorously assess location dependence.

Other experimental settings. This work focuses on stationary experiments with bulk file download to characterize 5G ON-OFF loops with three U.S. operators. To better quantify the real-world impact of loops, it is also important to examine other experimental factors such as applications, mobility and operators, in addition to devices and locations. We ran extra experiments with other applications (file upload, video streaming, and live streaming) and observed 5G ON-OFF loops regardless of the application type. This is because all these applications need to continuously transfer heavy traffic and require an active RRC connection at all times. We also tested with walking experiments, particularly around one test location or between two locations with 5G ON-OFF loops identified in the stationary experiments. We observed 5G ON-OFF loops in close proximity, then the loops were gone or changed. This is mainly due to varying RSRP/RSRQs as explained in §6. We believe that the problems are not limited to US operators because the issues we identified are not unique to certain operators but common for 5G technologies (5G SA and 5G NSA).

To this end, we have released our measurement and analysis scripts [5], along with our datasets [4], so that measurements can be easily reproduced and expanded at a larger scale with more network operators, phone models, locations, mobility, and applications. In the mean time, we plan to report our findings to 5G network operators to explore remedies that correct network configurations.

8 Related Work

Loops or instability in cellular networks. Two recent studies [18, 37] are the most relevant to our work. They present two 5G ON-OFF loop sub-types. Zhang et al. [37] investigated dependent misconfigurations when selecting multiple serving cells in 5G/4.5G networks and reported one N2 loop (5G NSA ↔ 4G) caused by inconsistent ON/OFF triggering conditions with uncoordinated A2/B1 thresholds. This loop sub-type is not observed in our study (inconsistent A2/B1 thresholds have been likely corrected by network operators, **F12**). Liu et al. [18] examined SCG failure handling in 5G NSA and reported another N2 loop (N2E2 in our work). N2E2 loops are observed in our study but they are not as rare as presented in [18]. Our work differs from these studies in several respects. First, prior studies consider other problems and only have by-product findings on 5G ON-OFF loops, whereas loops are our focus. Second, we conduct extensive measurement experiments to characterize real-world loop instances at a much larger scale and observe many more loop instances and types. We dive into their root causes and quantify loop likelihood in the wild. Our study is more recent and covers 5G SA. Several of our findings are new and surprising.

Loops or handover instability in cellular networks have been studied in earlier work [17, 22, 28, 31, 34, 36, 38]. The first work was traced back to ping-pong handovers, where a mobile device repeatedly switches between two cells within a short period of time, causing a "ping-pong" effect in 2G handovers [31]. Ping-pong handovers are regarded as transient loops because they are caused by user movement (around the boundary of two cell coverage). Persistent handover loops were exposed in 3G/4G networks [17] and then validated in other studies [28, 34, 38]; These loops were rooted in misconfigurations and uncoordinated parameter settings for reporting events. Yuan et al. [36] reported one loop in Google Fi, which stems from conflicting polices for inter-operator switch. Our work is inspired by this work but our problem is different. Prior studies focus on unstable mobility management which oscillates between two or more sets of serving cells, not completely turning 5G off. We focus on 5G ON-OFF loops which are more harmful than handover instability.

5G measurement on radio access and mobility. Another related but different topic is 5G measurement studies in recent years [12, 13, 15, 18, 19, 23, 24, 27, 29, 33, 35]. These studies measure a wide variety of 5G deployment and performance issues including early 5G radio rollout [23, 33], data performance [15, 24, 29], performance on the go [12, 27], radio access failures [18], mobility management [13], and carrier aggregation [35]. Some studies revealed RRC configuration and operations [13, 35] and RLF in practice [18, 27]. Our findings match theirs, but they do not focus on 5G ON-OFF loops due to RRC operation.

Loops and misconfigurations in other networks. Routing loops in the Internet have been studied over the past two decades (e.g., [11, 20, 21, 32]). These loops share some common issues in policy misconfiguration and structural problems that result in routing instability. Misconfigurations have been studied in other network contexts like data centers [14, 30], SDN [26] and DNS [16]. Although this work also considers persistent loops, the problem context is entirely different. Our problem context is 5G/4G RRC, which uses more complex policies and configuration to manage location-dependent radio access and unfortunately loses 5G access even when available and offering higher performance.

9 Conclusion

In this work, we unveil surprising findings about 5G ON-OFF loops in operational 5G SA and 5G NSA networks. Loops stem from inconsistent triggers to turn 5G on and off. We report unexpected inconsistencies responsible for various loop types that we observe in our study. Our work is still at an early stage, only uncovering the tip of the iceberg. It will be worthwhile to scale up our measurement study to extensively measure 5G ON-OFF loops in the wild, particularly as 5G is constantly evolving. This is critical since ON-OFF loops unnecessarily hurt data performance.

Acknowledgment. We thank our shepherd and all anonymous reviewers for their constructive feedback, which has substantially improved the paper. We also thank Ziyu Li and Sujala Awasthi for running some experiments to collect loop instances in the wild. This work has been partially supported by NSF grants CNS-2112471.

References

[1] 2023. Does the OnePlus 10 pro work on ATT now? https://www.reddit.com/r/ATT/comments/12tgi0x/does_the_oneplus_10_pro_work_on_att_now/.
[2] 2024. Network Signal Guru. https://play.google.com/store/apps/details?id=com.qtrun.QuickTest.
[3] 2025. 5G NR ARFCN calculator. https://5g-tools.com/5g-nr-arfcn-calculator/.

[4] 2025. 5G ON-OFF Measurement Datasets. https://github.com/mssn/5GMeas-Dataset.
[5] 2025. Readme for 5G ON-OFF loop measurements in the wild. http://milab.cs.purdue.edu/onoff_instruction.
[6] 3GPP. 2025. TS36.331: E-UTRA; Radio Resource Control (RRC). V17.12.0.
[7] 3GPP. 2025. TS37.340: E-UTRA and NR; Multi-connectivity; Overall description; Stage-2. V18..5.0.
[8] 3GPP. 2025. TS38.104: NR; Base Station (BS) radio transmission and reception. V17.17.0.
[9] 3GPP. 2025. TS38.331: NR; Radio Resource Control (RRC). V17.12.0.
[10] 5G America. 2024. Global 5G Connections Reach Nearly Two Billion. https://www.5gamericas.org/global-5g-connections-reaches-nearly-two-billion/.
[11] Lixin Gao and Jennifer Rexford. 2001. Stable Internet routing without global coordination. *IEEE/ACM Transactions on Networking (TON)* 9, 6 (2001), 681–692.
[12] 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 *Proceedings of the 2023 ACM on Internet Measurement Conference*. 678–695.
[13] Ahmad Hassan, Shuowei Jin, Arvind Narayanan, Ruiyang Zhu, Anlan Zhang, Wei Ye, Jason Carpenter, Z. Morley Mao, Zhi-Li Zhang, and Feng Qian. 2022. Vivisecting Mobility Management in 5G Cellular Networks. In *SIGCOMM'22*.
[14] Dilip A Joseph, Arsalan Tavakoli, and Ion Stoica. 2008. A policy-aware switching layer for data centers. In *ACM SIGCOMM Computer Communication Review*, Vol. 38. ACM, 51–62.
[15] Rostand A K. Fezeu, Claudio Fiandrino, Eman Ramadan, Jason Carpenter, Lilian Coelho De Freitas, Faaiq Bilal, Wei Ye, Joerg Widmer, Feng Qian, and Zhi-Li Zhang. 2024. Unveiling the 5G Mid-Band Landscape: From Network Deployment to Performance and Application QoE. In *ACM SIGCOMM 2024*. 358–372.
[16] Siva Kesava Reddy Kakarla, Ryan Beckett, Behnaz Arzani, Todd Millstein, and George Varghese. 2020. GRoot: Proactive Verification of DNS Configurations. In *Proceedings of the 2020 ACM SIGCOMM Conference (SIGCOMM'20)*. 310–328.
[17] 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 *ACM International Conference on Measurement and Modeling of Computer Science (SIGMETRICS'16)*.
[18] Yanbing Liu, Gunpeng Guo, and Chunyi Peng. 2024. Demystifying Secondary Radio Access Failures in 5G. In *HotMobile*.
[19] Yanbing Liu and Chunyi Peng. 2023. A Close Look at 5G in the Wild: Unrealized Potentials and Implications. In *INFOCOM*.
[20] Ratul Mahajan, David Wetherall, and Tom Anderson. 2002. Understanding BGP misconfiguration. *ACM SIGCOMM Computer Communication Review* 32, 4 (2002), 3–16.
[21] Markus Maier and Johanna Ullrich. 2023. In the loop: A measurement study of persistent routing loops on the IPv4/IPv6 Internet. *Computer Networks* 221 (2023).
[22] Vikash Mishra, Debabrata Das, and Namo Narayan Singh. 2020. Novel algorithm to reduce handover failure rate in 5G networks. In *2020 IEEE 3rd 5G World Forum (5GWF)*. 524–529.
[23] 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 *(WWW'20)*.
[24] Arvind Narayanan, Xumiao Zhang, Ruiyang Zhu, Ahmad Hassan, Shuowei Jin, Xiao Zhu, Xiaoxuan Zhang, Denis Rybkin, Zhengxuan Yang, Zhuoqing Morley Mao, Qian Qian, and Zhi-Li Zhang. 2021. A variegated look at 5G in the wild: performance, power, and QoE implications. In *Proceedings of the 2021 ACM SIGCOMM Conference (SIGCOMM'21)*. 610–625.
[25] OpenSignal. 2023. Benchmarking the Global 5G Experience: June 2023. https://www.opensignal.com/2023/06/30/benchmarking-the-global-5g-experience-june-2023.
[26] Heng Pan, Zhenyu Li, Penghao Zhang, Kave Salamatian, and Gaogang Xie. 2020. Misconfiguration checking for SDN: data structure, theory and algorithms. In *IEEE 28th International Conference on Network Protocols (ICNP'20)*.
[27] Yueyang Pan, Ruihan Li, and Chenren Xu. 2022. The first 5G-LTE comparative study in extreme mobility. *Proceedings of the ACM on Measurement and Analysis of Computing Systems (SIGMETRICS)* 6, 1 (2022).
[28] Chunyi Peng, Yuanjie Li, Zhuoran Li, Jie Zhao, and Jiaqi Xu. 2016. Understanding and Diagnosing Real-World Femtocell Performance Problems. In *IEEE International Conference on Computer Communications (INFOCOM)*.
[29] Muhammad Iqbal Rochman, Vanlin Sathya, Damian Fernandez, Norlen Nunez, Ahmed S Ibrahim, William Payne, and Monisha Ghosh. 2023. A comprehensive analysis of the coverage and performance of 4G and 5G deployments. *Computer Networks* 237 (2023), 110060.
[30] Peng Sun, Ratul Mahajan, Jennifer Rexford, Lihua Yuan, Ming Zhang, and Ahsan Arefin. 2014. A Network-State Management Service. In *SIGCOMM'14*.
[31] Daniel Wong and Teng Joon Lim. 1997. Soft handoffs in CDMA Mobile Systems. *IEEE Personal Communications* 4, 6 (1997), 6–17.
[32] Jianhong Xia, Lixin Gao, and Teng Fei. 2007. A measurement study of persistent forwarding loops on the Internet. *Computer Networks* 51, 17 (2007), 4780–4796.

CS	Cell Set
MCG	Master Cell Group
NSA	Non-StandAlone (one 5G deployment option)
PCell	Primary cell of the master cell group (MCG)
PSCell	Primary cell of the secondary cell group (SCG)
RAN	Radio Access Network
RAT	Radio Access Technology (here, 5G or 4G)
RLF	Radio Link Failure
RRC	Radio Resource Control
RSRP	Reference Signal Received Power
RSRQ	Reference Signal Received Quality
SA	StandAlone (one 5G deployment option)
SCG	Secondary Cell Group
SCell	Secondary Cell
UE	User Equipment

Table 6: Abbreviations and acronyms used in this paper.

[33] 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 *(SIGCOMM'20)*.
[34] Shichang Xu, Ashkan Nikravesh, and Z Morley Mao. 2019. Leveraging Context-Triggered Measurements to Characterize LTE Handover Performance. In *PAM*.
[35] Wei Ye, Xinyue Hu, Steven Sleder, Anlan Zhang, Udhaya Kumar Dayalan, Ahmad Hassan, Rostand AK Fezeu, Akshay Jajoo, Myungjin Lee, Eman Ramadan, et al. 2024. Dissecting carrier aggregation in 5G networks: Measurement, QoE implications and prediction. In *ACM SIGCOMM 2024*. 340–357.
[36] Zengwen Yuan, Qianru Li, Yuanjie Li, Songwu Lu, Chunyi Peng, and George Varghese. 2018. Resolving Policy Conflicts in Multi-Carrier Cellular Access. In *ACM International Conference on Mobile Computing and Networking (MobiCom'18)*.
[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 *CoNext*.
[38] Xiaohui Zhao, Hanyang Ma, Yuan Jin, and Jianguo Yao. 2018. Measuring instability of mobility management in cellular networks. *IEEE Network* 32, 5 (2018), 138–144.

A Abbreviations and Acronyms

Table 6 lists all the abbreviations and acronyms used in this paper.

B Methodology with Examples

In this appendix, we present our approach of retrieving the serving cell sets from RRC signaling messages received on the test phone. This is an important step in our measurement methodology (§4.1). We then analyze the sequence of serving cell sets and check whether a loop occurs.

A serving cell set is updated by adding, modifying and removing one or more serving cells as shown in Figure 23. The serving cell set CS_{k+1} at time t_{k+1} is retrieved by merging CS_k and all the changes from t_k to t_{k+1}, $CS_{k+1} = CS_k + \Delta CS_{k+1}$. There are three types of changes (Figure 23): ❶ the PCell (of the MCG) changes; ❷ the PCell does not change but SCells of the MCG change; ❸ the MCG does not change but the SCG changes. All the changes are made through RRC procedures.

We use several examples to illustrate how to parse RRC signaling messages to extract ΔCS_k and update CS_k over time. All the raw messages are captured by Network Signal Guru [2]. These examples (Figure 24 - Figure 26) together show how the serving cell sets change during one ON-OFF cycle in the first loop example (Figure 1b). This is one S1E3 loop instance with OP_T (5G SA) at one test location (P16 in A1). Table 2 lists five main 5G cells observed

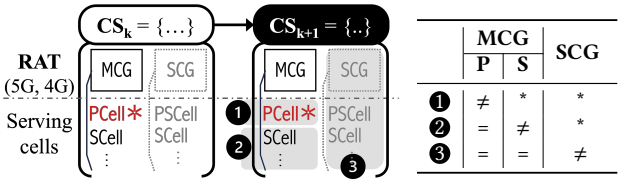

Figure 23: The serving cell set (CS) changes in three forms. P and S are the PCell and SCell of the MCG, and * is a wildcard.

in these examples. In addition to the five main cells, another cell 104@501390 is observed and used in these examples. Each cell is denoted by ID@FreqChannelNo, where ID is the cell identifier and FreqChannelNo is the radio frequency channel number (here, ARFCN for 5G and EARFCN for 4G) [8]. An online tool [3] can be used to calculate the actual frequency given a channel number. There are six 5G cells are observed over four 5G channels: 521310, 501390, 398410 and 387410.

• **From IDLE to 5G SA 1 (Figure 24)**. In this example, the serving cell set changes from ∅ → {393@521310, one 5G PCell } (❶). This is realized through `RRC Connnection Establishment` procedure to set up an RRC connection when the RRC state is IDLE. The procedure is specified by 3GPP TS38.331 ([9] for 5G SA). Key parameters are acquired from the `Master Information Block (MIB)` and `System Information Block (SIB)`, including candidate cells and the selection criteria. The UE then measures all candidate cells and checks whether there exists any candidate cell which meets the specified selection criteria (e.g., RSRP/RSRQ larger than a pre-configured threshold). Then, the UE chooses an appropriate cell (here, 393@521310) to establish the RRC connection with three messages: `RRC Setup Request`, `RRC Setup`, and `RRC Setup Complete`. Note that the PCell information is found in all RRC messages but skipped in the last two messages (Figure 24). If the cell is seen but not used, its `NR Cell Global ID` is invalid (= 0).

• **From 5G SA 1 to 5G SA 2 (Figure 25)**. This follows the previous example and adds three 5G SCells into the MCG (❷). The change from CS_1 to CS_2 is realized through one `RRC Reconfiguration` procedure [9]. `RRC Reconfiguration` is a vital RRC procedure which is widely used to modify the existing RRC connection for various purposes such as configuring measurements, modifying radio bearers, or modifying (adding/releasing) serving cells. In this example (Figure 25), the `RRC Reconfiguration` procedure is used to add three 5G cells, which can be extracted from the `RRC Reconfiguration` message with one `sCellToAddModList` field listing all the cells to add (here, `273@387410`, `273@398410`, `393@501390`). Every SCell is assigned an `sCellIndex`. Upon receiving this command, the UE starts adding these three cells and then responds with an `RRC Reconfiguration Complete` message after it finishes 5G SCell addition (three 5G SCells used along with one PCell).

• **From 5G SA 2 to IDLE (Figure 26)**. In this example, the serving cell set experiences two changes. The first change from CS_2 to CS_3 is also realized through an `RRC Reconfiguration` procedure and the difference is modifying one 5G SCell. Specifically, the `RRC Reconfiguration` message is configured with both `sCellToAddModList` and `sCellToReleaseList` fields. The former lists the cell to add (namely, `104@501390` with `sCellIndex` = 4) and the latter lists the cell to release (here, `sCellIndex` = 3, namely, `393@501390`). Both cells run over the same channel and the first change is one intra-channel SCell modification (❷). The second change also runs one `RRC Reconfiguration` procedure but the modified SCell fails this time. This RRC procedures asks the UE to replace one 5G SCell (`273@387410`, `sCellIndex` = 1) with another intra-channel cell (`371@387410`, `sCellIndex` = 3). Although this procedure ends with an `RRC Reconfiguration Complete` message, the exception occurs immediately (within around 5 ms). In fact, we see that all the data transmission is suddenly disrupted in the air. Moreover, we see that Network Signal Guru experiences an exception without any logs for a few seconds. During this exception, this RRC connection is released (with all the serving cells) and the serving cell set returns to CS_0 (IDLE). Note that this cell set change is not extracted from RRC signaling messages exchanged in the air (no transmission in the air) but learned from the network state. `MM5G` stands for 5G MM (Mobility Management) which is mainly involved in the Registration process. Here, the `MM5G` state is `DEREGISTERED` because it believes that there are no cells available for use. In this example, we end with the `RRC Setup Request` message to show how long the idle state lasts (about 11 seconds). This message is used to start an `RRC Connnection Establishment` procedure as shown in Figure 24. The serving cell set gets back to CS_1 (5G SA) after the `RRC Connnection Establishment` procedure completes and repeats the ON-OFF loop afterwards.

More about this loop example. We check all the loop and non-loop instances at this test location and make three observations. First, all the loop instances are caused by 5G SCell modification failure (at channel 387410). Second, 5G SCells at other channels (501390 and 398410) can be successfully modified. Third, 5G SCell modification may not happen every time, depending on the RSRP/RSRQ comparison of intra-channel cells. In a nutshell, the ON-OFF cycle continues as long as 5G SCell modification (over channel 387410) fails, regardless of 5G SCell modification at other channels. In the same loop instance, we see that the serving cell set attempts to change from CS_2 to CS_4 directly.

From 5G SA to 5G NSA We use the same approach to extract the serving cellset sequence over time for 5G NSA. The only difference is that 5G NSA uses 4G RRC (not 5G RRC) and all RRC procedures are similar but specified by TS36.331 [6] (for 4G RRC, used by 5G NSA). Similar to 5G SA, when the UE is in the IDLE state, it first obtains cell selection parameters from `MIB` and `SIB`. It then selects a 4G cell as the target PCell, and initiates an `RRC Conn. Establishment` procedure to establish a 4G connection through `RRC Connection Setup Request`, `RRC Connection Setup`, and `RRC Connection Setup Complete`. To add a 5G SCG, the network sends an `RRC Conn. Reconfiguration` message configuring a B1 event to instruct the UE to measure and report nearby 5G candidate cells. In response, the UE sends a measurement report containing all qualified 5G cells. Typically, the 5G candidate cell with the strongest RSRP or RSRQ in the report is selected as the PSCell and is configured in the `spCellConfig` field of the `RRC Reconfiguration` message. Following this procedure, the UE transitions to a 5G NSA cellset with both a 4G MCG and a 5G SCG.

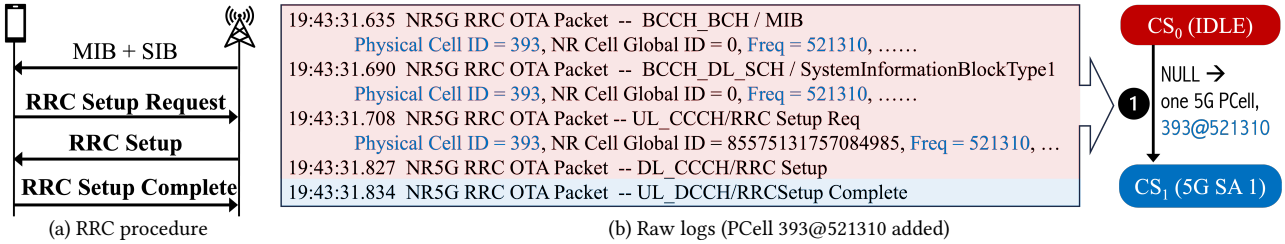

(a) RRC procedure (b) Raw logs (PCell 393@521310 added)

Figure 24: One `RRC Connnection Establishment` example where the PCell changes from null to a valid one (393@521310, 5G SA).

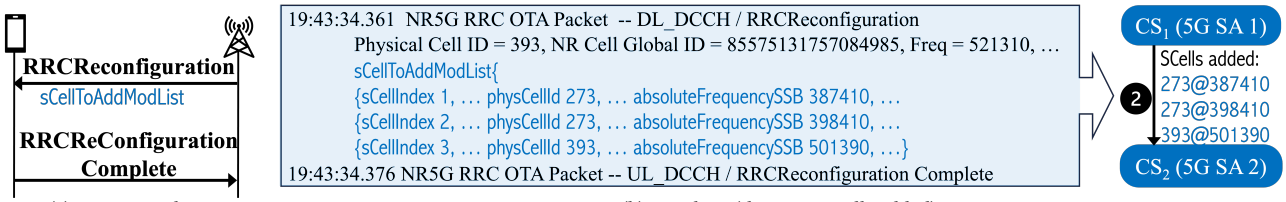

(a) RRC procedure (b) Raw logs (three 5G SCells added)

Figure 25: One `RRC Reconfiguration` example where three 5G SCells are added (along with 5G PCell 393@521310).

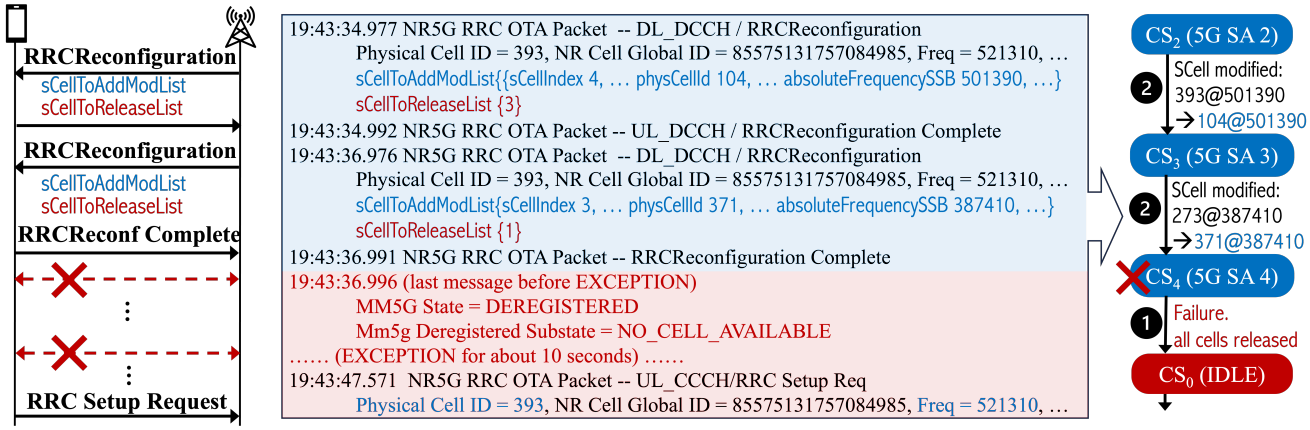

(a) RRC procedure (b) Raw logs (one 5G SCell successfully modified and another failed))

Figure 26: The last example with two `RRC Reconfiguration` procedures (one success, one failure) where 5G cells are released.

As discussed in §5.2, a UE may lose 5G SCG for several reasons, including 4G PCell radio link failure, PCell handover and SCG failure handling. When a UE receives an `RRC Conn. Reconfiguration` message containing the field of `rrcReestablishmentRequest`, a PCell radio link failure has occurred. In this case, the UE releases both 4G MCG and 5G SCG and returns to IDLE state. If the UE receives an `RRC Conn. Reconfiguration` message including `mobilityControlInfo` field but without `spCellConfig` field, this is a PCell handover without SCG reconfiguration, resulting in the loss of 5G SCG. Finally, when the `SCGFailureInformation` field is present in the `RRC Conn. Reconfiguration` message, this signals an SCG failure, and the network immediately sends another `RRC Conn. Reconfiguration` message instructing the UE to release 5G SCG.

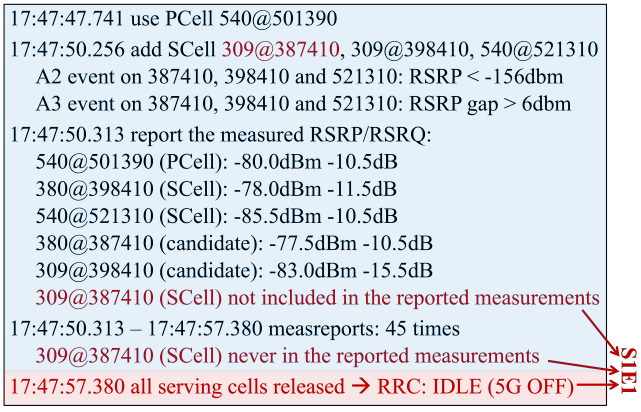

Figure 27: One S1E1 instance (bad apple: 309@387410).

C Loop Instances

In this appendix, we give real-world instances for each loop subtype.

S1E1. Figure 27 gives a loop instance where 5G turns off due to no RSRP/RSRQ measurements of SCells in 5G SA mode. In this instance, `309@387410` is the only serving cell that does not appear in the measurement and report step. The radio quality of this cell is too poor to be measured. Typically, a UE detects and releases a poor SCell by comparing its RSRP measurement value against the A2 event threshold. However, without measurement results, the UE

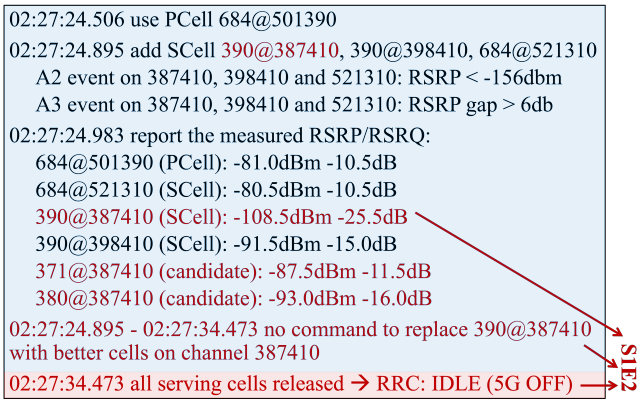

Figure 28: One S1E2 instance (bad apple: 390@387410).

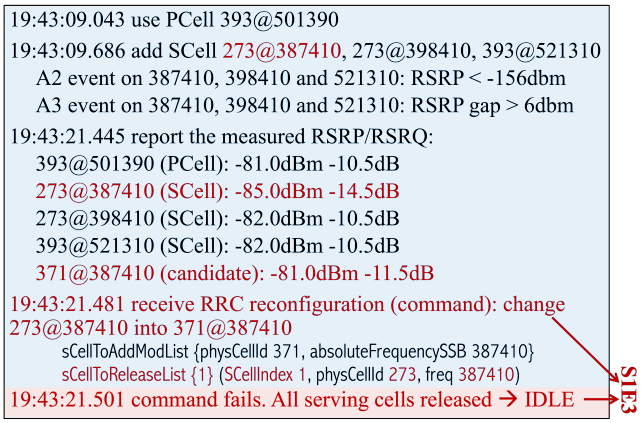

Figure 29: One S1E3 instance (bad apple: 273@387410).

cannot release this poor cell. As a result, the failure on this SCell leads to the release of the entire cell group.

S1E2. Figure 28 illustrates a loop instance where 5G connectivity is lost due to poor RSRP/RSRQ measurements from one or more 5G SCells. In this example, the measurement results show that the serving cell `390@387410` has poor radio quality (-108.5 dBm RSRP and -25.5 dB RSRQ). Despite the presence of two significantly better candidate cells (`371` and `380`) on the same channel, the UE does not receive any handover commands to replace this poor serving cell. Eventually, this leads to a radio link failure and the UE releases all serving cells.

S1E3. Figure 29 presents a loop instance where 5G is lost because RRC fails to modify a 5G SCell. In this instance, the failure occurs immediately after the UE receives an `RRC Reconfiguration` message instructing it to switch the SCell from `273@387410` to `371@387410`. In our per channel analysis in §5.3, we have identified that for OP_T, the handover failure probability of SCells on channel `387410` is much higher than on other channels.

N1E1. Figure 30 presents a loop instance due to radio link failure on 4G MCG in 5G NSA mode. In this instance, a UE detects a radio link failure on 4G MCG while being served by PCell `191@66936`. The UE sends a reestablishment request (`otherFailure` as `reestablishmentCause`) to the network side and releases all 4G and 5G serving cells. Later, the UE reestablishes the 4G connection on `238@5815`. However, the cells on channel

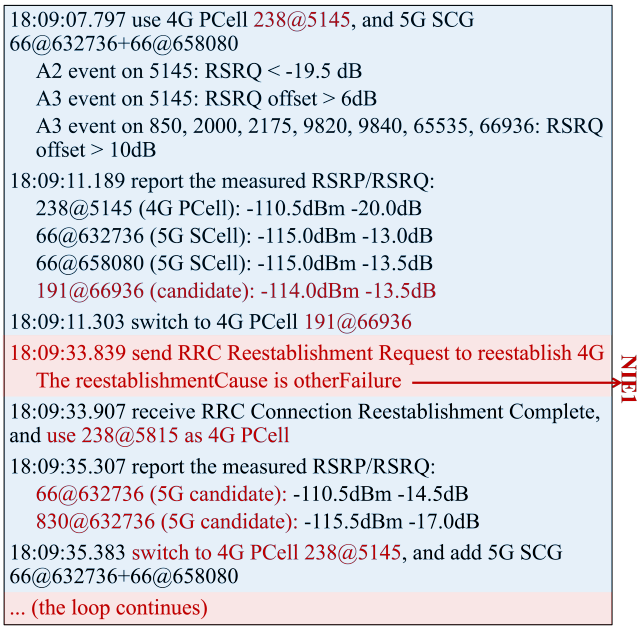

Figure 30: One N1E1 instance (RLF on 191@66936).

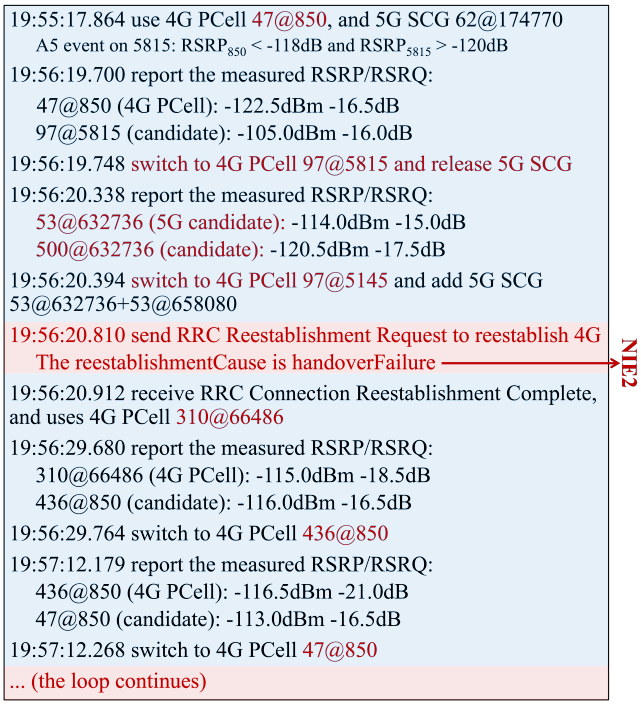

Figure 31: One N1E2 instance (handover failure on 97@5415).

5815 cannot be used with 5G due to OP_A's policy. As a result, when the UE reports available 5G cells, it is redirected back to the original PCell `238@5145`, which supports 5G SCG. This creates a handover loop among the three PCells.

N1E2. Figure 31 shows a loop instance due to handover failure on 4G MCG in 5G NSA. The UE receives a handover command to switch PCell from `97@5815` to `97@5145`. However, the handover

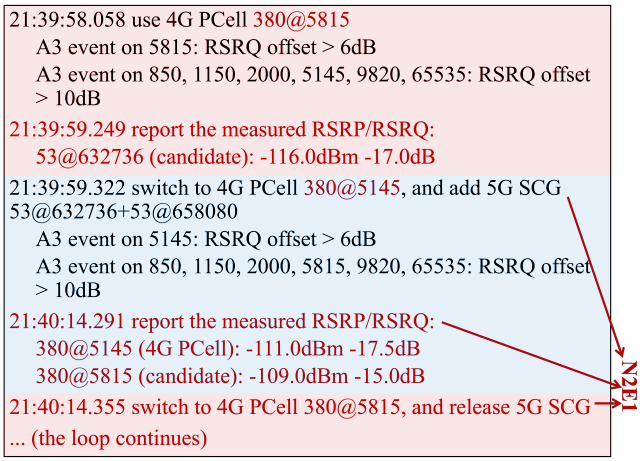

Figure 32: One N2E1 instance (policies on channel 5815).

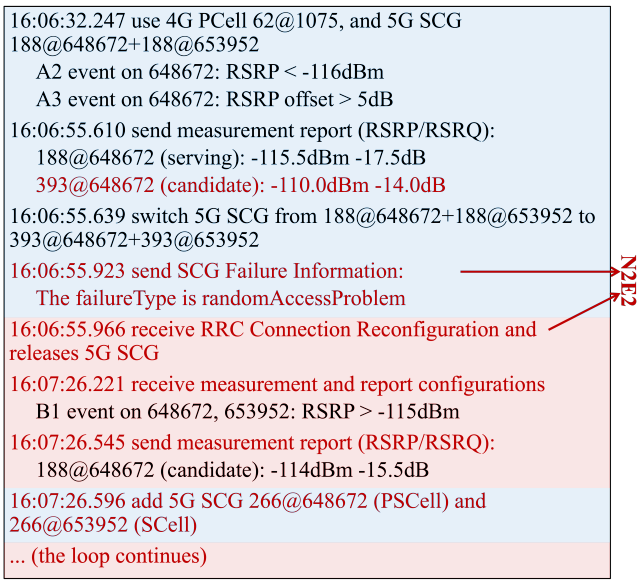

Figure 33: One N2E2 instance (policies on SCG recovery).

fails to complete, triggering the UE to send a reestablishment request to the network with `reestablishmentCause` set to `handoverFailure`. The UE subsequently performs multiple PCell handovers but ultimately returns to the original PCell, leading to a handover loop and repeated failures.

N2E1. Figure 32 shows a N2E1 loop instance where 5G SCG is lost due to a 4G PCell handover. In this case, a UE is initially served by a cellset with 4G MCG and 5G SCG, and when 4G PCell hands over from `380@5145` to `380@5815`, the UE loses the 5G SCG. More importantly, the handover criteria between these two cells are asymmetric. As discussed in §5.2, the handover from `380@5815` to `380@5145` is triggered when an available 5G cell is reported, while the handover from `380@5145` to `380@5815` is triggered by the A3 event when the RSRQ of 380@5815 is stronger. This leads to repeated handovers between these two cells and a persistent 5G ON-OFF loop. Consequently, the UE has to re-run the SCG addition procedure again and again, leading to frequent transient 5G disruptions.

N2E2. Figure 33 illustrates a N2E2 loop instance caused by SCG failure in NSA mode. In this instance, when a UE switches 5G SCG from `188@648672+188@653952` to `393@648672+393@653952`, the UE detects random access failure and reports `SCGFailureInformation`. In response, the UE receives an `RRC Conn. Reconfiguration` message to instruct it to release 5G SCG. After 30.3 seconds, the UE receives a new 5G configuration, measures and reports available 5G cells, and finally recovers the 5G SCG. Therefore, the SCG addition procedure is significantly delayed by configuration and measurement, leading to a long-time 5G OFF period.