A Comprehensive Analysis of the Coverage and Performance of 4G and 5G Deployments

Muhammad Iqbal Rochman[1†], Vanlin Sathya[†], Damian Fernandez[*], Norlen Nunez[*], Ahmed S. Ibrahim[*], William Payne[†], and Monisha Ghosh[‡]

[+] *Department of Computer Science, University of Chicago, Illinois, USA.*
[*] *Department of Electrical and Computer Engineering, Florida International University, Florida, USA.*
[‡] *Department of Electrical Engineering, University of Notre Dame, Indiana, USA.*

Abstract

5G New Radio (5G NR), has now been deployed in all available bands: low (< 1 GHz), mid (1 - 6 GHz), and high (> 24 GHz), each with a trade-off between coverage and throughput/latency performance. The preceding 4G Long Term Evolution (4G LTE) networks also catching up to 5G mid-band by deploying in the unlicensed 5 GHz (using License Assisted Access/LAA) and the 3.5 GHz Citizens Broadband Radio Service (CBRS). We present a comprehensive analysis of 4G and 5G deployments in Chicago and Miami, focusing on coverage, throughput, and latency performance in low-, mid-, and high-bands, under several scenarios: outdoor, outdoor-to-indoor, and under high temperature. To measure deployed networks, we utilized a scalable methodology with commercial and custom apps to collect detailed signal data (*e.g.,* signal strength, cell ID, throughput) on user device (*i.e.,* smartphones). The analysis based on our measurements yields the following findings: (i) when comparing the throughput and latency performance of optimized mid-band 4G networks to 5G (both standalone (SA) and non-standalone (NSA)), they exhibit comparable performance, and (ii) it's important to note that mmWave 5G has the capability to deliver multi-Gbps throughput, even in NSA mode. However, this high-speed performance is susceptible to limitations imposed by factors such as distance, body interference, obstructions (such as Low-e glass), and overheating, which can render its performance less reliable. Therefore, even though 5G exhibits considerable potential in its early stages, additional efforts are required to guarantee that the stated goals of 5G are met.

Keywords: 5G, 4G, measurement, throughput, latency, coverage, outdoor-to-indoor, thermal

1. Introduction

Over time, the use cases of mobile cellular traffic have undergone a significant transformation, expanding beyond their traditional functionalities of voice calls, text messaging, and video communication. With the rapid advancement of technology, new applications such as augmented reality (AR) and virtual reality (VR) have emerged, ushering in a new era of mobile experiences. These new applications demand higher throughput and lower latency than can be delivered over current 4G networks. Thus, next generation 5G New Radio (5G NR) [1] cellular networks are being rapidly deployed to meet the increasing demands. The deployment of 5G networks encompasses a wide range of frequency bands, including low-band (< 1 GHz), mid-band (1−6 GHz), and high-band (> 24 GHz), each offering distinct advantages and trade-offs. Low-band 5G provides a solid foundation for wider coverage but lower throughput, thus ensuring a reliable connection in rural areas where coverage is the primary concern. Mid-band 5G, including C-Band (3.7 - 3.98 GHz), offers higher throughput compared to low-band, but lower coverage due to the propagation characteristics of

[1]Email: muhiqbalcr@uchicago.edu Addr: 5730 S Ellis Ave Rm 375, Chicago, IL 60637
This work was supported by NSF under grant CNS-1618836.

the band. It provides a good balance between coverage area and capacity, offering improved user experiences while still maintaining a reasonable coverage radius. Lastly, high-band, or millimeter wave (mmWave), 5G delivers multi Gbps speeds and extremely low latency. However, mmWave signals have limited range and are more susceptible to obstructions. They require a dense network of small cells to ensure consistent coverage and capacity in specific areas, such as dense urban environments, and cause consumer devices to overheat.

To ease deployments, 5G networks are deployed in two modes: Standalone (SA), where operators use only 5G for the entire network stack, and Non-Standalone (NSA), which uses a 4G channel as a primary anchor channel thus easing the transition from 4G to 5G SA. Presently, about 75% of 5G deployments in the US are NSA, but SA is picking up deployments in limited areas as well [2]. Recent measurements of 5G NSA mmWave deployments in major US cities demonstrated that indeed 5G mmWave can deliver extremely high throughput in the range of 1 - 2 Gbps [3, 4, 5]. These high throughputs are enabled by aggregating up to eight 100 MHz mmWave channels, depending on network and device capabilities. However, there are research questions still unanswered regarding the performance of 5G deployments in SA and NSA. Firstly, as current NSA deployments use the 4G primary channel as the signalling channel, the latency performance is not much improved compared to 4G. Secondly, 5G throughput deployments in low- and mid-band, along with outdoor-to-indoor performance, have not been analyzed yet. Finally, 5G mmWave networks, which utilize up to 8 channel aggregation on user devices, can place strain on the device's radio front-end and antenna system, leading to elevated skin temperatures and potentially result in throughput throttling.

In this paper, we present a thorough analysis and comparison of various deployment scenarios, utilizing a scalable methodology to measure cellular networks deployed in Chicago and Miami. We employed end-user devices (smartphones) to capture signal conditions and network performance on the point of view of the end-device. On each device, we utilized apps specifically aimed for focused data captures over a prolonged period of time: SigCap, FCC ST, NSG [5, 6, 7, 8], and Qualipoc [9]. Each apps comes with its trade offs, described in detail in §3.

We believe that this is the first such comprehensive study of 4G and 5G in multiple bands. Additionally, we address multiple research questions where we apply our device-based measurements. The contributions of this paper are the following:

1. We present analyses of measurements of 4G and 5G deployments in Chicago conducted between May and June 2021. In particular, we focus our analysis on the coverage, throughput, and latency performance between various deployment frequencies and technologies (4G, 5G-SA, and 5G-NSA) for the three major US operators (AT&T, T-Mobile, and Verizon).

2. We further focus our measurements on the mmWave deployments of 5G in Miami which were conducted between January and June 2021. This particular measurement campaign sought to answer the following question: what is the impact of **distance**, **human body blocking**, and **multiple devices** on 5G mmWave throughput performance?

3. As prior measurements were conducted outdoors, we then focused on measuring the degradation of 5G mmWave outdoor-to-indoor performance due to building loss. We conducted detailed measurements in a University of Chicago building near a Verizon mmWave base station (BS) in July 2021, and show that the electromagnetic-blocking capability of the Low-e glass installed in the building greatly affects the indoor throughput performance. Specifically, we quantify the throughput as a function of window opening gap size.

4. Finally, we present an analysis of the degradation of 5G mmWave throughput over a prolonged time period, due to the skin temperature rise resulting from sustaioned Gbps throughput. This effect has not been previously analyzed due to commonly-used speed test applications such as Ookla [10] and FCC Speed Test [11] running only over 5 - 10 seconds, which is not representative of user behavior. We utilized a thermal camera and Android APIs to show the correlation of skin temperature rise in the device with throughput reduction.

2. Related Work

5G NR specifications proposed by 3GPP standards allow the operators to operate on a wide spectrum of frequencies [12]. High-band 5G in mmWave, along with the mid-/low-band sub-6 GHz counterparts, make up the current 5G market for high capacity and low latency. We pay close attention to mmWave 5G due to its ultra-high bandwidth which attracts emerging bandwidth-hungry applications. On the other hand, mmWave is very sensitive to factors such as mobility and blockage due to its much shorter wavelength, making the upper-layer network management (e.g., bit-rate adaptation of video streaming) more challenging. Despite numerous studies on modeling and simulation of mmWave links [3, 4, 13, 14], the impact of mmWave on commercial 5G performance, as well as mobile application Quality-of-Experience (QoE) is largely under-explored. Also, the different 5G deployment modes, e.g., standalone (SA) vs. non-standalone (NSA), mobility and data management - are far more complex than what is generally analyzed or simulated. Measurement studies in the literature have shown wild fluctuations in 5G throughput, and even worse, service outages. Authors in [15] studied the current 5G deployments adopted by three major U.S. carriers

Challenges of measuring real-world 4G and 5G deployments. With the increasing number of 5G mmWave deployments in many cities, the emphasis is shifting to quantifying 5G mmWave performance using commercial deployments and user equipment (UEs). Recent literature [3, 4, 13, 14] has demonstrated the feasibility of achieving very high throughput (up to 2 Gbps) with consumer smartphones over commercially deployed 5G mmWave. These high throughput values are enabled by aggregating up to eight 100 MHz mmWave channels, depending on network and device capabilities. However there are still a number of challenges associated with guaranteeing QoS in 5G mmWave: beam-tracking, beam management, building blockage, and rain attenuation, to name a few. Advanced techniques, based on machine learning and artificial intelligence, have been proposed for addressing these limitations, for example in [16, 17]. At the lower end of the spectrum, frequency range 1 or FR1 which includes licensed bands below 6 GHz and the unlicensed 6 GHz band, 5G can provide extended, robust coverage but with lower throughput due to the limited available bandwidth[2]. Furthermore, the existing 5G NSA deployment has the potential for higher latency due to the additional overhead introduced by dual connectivity (DC). Using NSA configuration, the primary 4G channel and the secondary 5G channel might originate from base stations (BSs) that are not co-located.

Simultaneously, we are witnessing a growing trend in the deployment of 4G networks within the mid-band, utilizing License Assisted Access (LAA) specification to access the unlicensed 5 GHz frequencies as secondary aggregated channels. Most studies of LAA and Wi-Fi coexistence are based primarily on theoretical analyses, simulations, and limited experiments. With widespread deployments of LAA beginning in major cities in the US, it is now possible to verify deployment parameters and their impact on coexistence performance. Our first measurements and deployment statistics in various locations in Chicago where LAA networks have been deployed by AT&T, T-Mobile and Verizon in close proximity to Wi-Fi networks were reported in [6]. Additional controlled experiments between 5 GHz LAA and Wi-Fi deployments in downtown Chicago and the University of Chicago [6, 8, 7] showed an average cellular throughput of 150 Mbps when 60 MHz in the unlicensed band (using three aggregated 20 MHz channels in 5 GHz) is used along with a primary 15 MHz - 20 MHz bandwidth primary channel: this is 6× the average throughput of the licensed primary band alone.

The shared, licensed 3.55 - 3.7 GHz Citizens Broadband Radio Service (CBRS) band has also been used by major US operators as secondary aggregate channels to boost capacity. The standalone nature of CBRS deployment is mainly used for private 4G and 5G deployments [18] for mission critical applications in warehouses, logistics, ports, hospital, etc. Adjacent channel coexistence issues between 5G C-band (3.7 - 4.2 GHz) and 4G CBRS were reported in [19] which demonstrated the potential of C-band channels (excluding the coexistence issues) attaining an average throughput of 300 Mbps over a 60 MHz bandwidth. Thus, current cellular deployments, which feature a multitude of technologies and aggregated bands, have grown notably intricate. These complexities pose substantial challenge when attempting to replicate them

[2]At the time of measurement, 5G channel aggregation in FR1 was not implemented. On the other hand, multiple channels can be aggregated in frequency range 2, or FR2, for all 5G mmWave-enabled phones.

for research purposes, even within a large-scale test-bed such as the National Science Foundation's (NSF) Platforms for Advanced Wireless Research (PAWR) [20]. Therefore, measurements on deployed networks are crucial to understand its behaviour in the real world.

5G Outdoor-to-Indoor performance. Authors of [21] performed a measurement campaign in mmWave focusing on the Outdoor-to-Indoor (OtI) performance in different environments and concluded that a user can achieve a maximum of 2.5 Gbps downlink (DL) throughput in 90% of indoor locations with a link distance of up to 68 m. However, these conclusions are based on measurements using continuous wave channel sounders and non-commercial UEs, where the antenna placement and orientation can be quite different from those on commercial devices. Further, the throughput predictions were based on measured signal strengths, not on actual data transmission, since the experiments were conducted using channel sounders. Thus, our work in this paper reports the first results from OtI measurements in a location where an outdoor 5G mmWave BS by Verizon is utilized for detailed indoor measurements in a dormitory building in close proximity. Our measurement results, using a deployed network and commercial devices, offers a different conclusion from [21] regarding achievable throughput indoors from outdoor 5G mmWave: >1 Gbps DL throughput on 5G mmWave can only be sustained if the UE is line-of-sight (LoS) with the BS and the window is open. We note that our measurements were carried out in a newer building with low-E glass windows. Since we did not have access to other indoor locations in close proximity to a deployed 5G mmWave BS, we were unable to compare performance with different types of glass as was done in [21].

Prior analysis of 5G thermal throttling. Most recently, authors of [13] present detailed measurements of 5G mmWave deployments by two major commercial 5G operators in the US in diverse environments using smartphone-based tools. The measurement-driven propagation analysis demonstrated performance differences due to terrain, frequency of operation, antenna pattern, etc. However, the relationship between device temperature and sustained 5G mmWave throughput was not explored. In particular, we seek to demonstrate that the drop in throughput is indeed due to thermal effects. According to the 3GPP standard [22], a UE can provide information to the BS about its thermal state via the the RRC_CONNECTED message field. Upon receiving such a message from the UE, the BS will respond by temporarily reducing the number of aggregated data streams until the thermal warning messages are no longer received. Prolonged thermal throttling may even lead to handover to 4G, until the skin temperature drops to below a pre-specified threshold.

3. Overview of Data Collection Methodology

Measurements of deployed 4G and 5G networks were collected over several months in 2021 and 2022: (i) from May to June 2021 in Chicago, (ii) from January to June 2021 in Miami, (iii) July 2021 in Chicago, and (iv) May 2022 in Chicago, each with varying objectives that will be described in detail in §4, §5, §6, and §7, respectively. In all our experiments, we utilized commercial devices (smartphones) equipped with measurement apps to capture 4G and 5G signal data and performance metrics. We chose this methodology instead of using professional drive-test equipment since our methodology is more scalable and potentially suitable for crowd-sourcing. We have used a similar methodology to map broadband accessibility in a recent feasibility assessment by the Federal Communications Commission (FCC) in Colorado [23].

We use four Android applications summarized in Table 1, each of which supply varying degrees of information:

SigCap [24], developed at the University of Chicago, is an Android app that passively collects Global Positioning System (GPS) and wireless signal data through Android APIs without requiring root access. The API extracts data directly from the modem chip and hence conforms to the relevant standard specifications. However, the currently available APIs have some limitations: (i) inability to distinguish between secondary and neighboring 4G channels and (ii) limited 5G information. SigCap collects data every 10 seconds, which is the minimum interval allowed by the API to conserve power. While SigCap can also capture Wi-Fi data, we only focus on 4G and 5G signal analysis in this paper.

Each data record from SigCap consists of the following parameters: time-stamp, GPS coordinates, information on all 4G received channels (Physical Cell Id (PCI), E-UTRA Absolute Radio Frequency Channel

Table 1: Measurement Apps' Features

Features	SigCap	FCC ST	NSG	QualiPoc
LTE Cell Information	**All cells:** PCI, EARFCN, Band, RSRP, RSRQ, RSSI + Primary cell bandwidth	**Primary cell only:** PCI, EARFCN, Band, RSRP, RSRQ	PCI, EARFCN, Band, Bandwidth, RSRP, RSRQ, RSSI, SINR, CQI, MIMO mode, RB allocation	PCI, EARFCN, Band, Bandwidth, RSRP, RSRQ, RSSI, SINR, CQI, MIMO mode, RB allocation
5G Cell Information	5G-RSRP and 5G-RSRQ	5G-RSRP and 5G-RSRQ	PCI, NR-ARFCN, Band, Bandwidth, Beam ID, 5G-RSRP, 5G-RSRQ, SINR, CQI, MIMO mode, RB allocation	PCI, NR-ARFCN, Band, Bandwidth, Beam ID, 5G-RSRP, 5G-RSRQ, SINR, CQI, MIMO mode, RB allocation
Throughput-related metrics	No	Application-level uplink/downlink throughput, latency	Application, RLC, MAC, and PHY layer uplink/downlink throughput	Application, RLC, MAC, and PHY layer uplink/downlink throughput
Root access	No	No	Yes	Yes

Number (EARFCN), band number, Reference Signal Received Power (RSRP), Reference Signal Received Quality (RSRQ) and Reference Sgnal Strength Indicator (RSSI), 4G primary channel bandwidth, 5G Status, 5G RSRP, and 5G RSRQ. 5G PCI and bandwidth information are currently unavailable. Using the above data, we generate heat maps of 4G and 5G RSRP as shown in Fig. 1b, 1c, and 1d, by defining 10 m square grids and averaging the collected RSRP of the deployment we are interested in (4G Licensed, LAA, CBRS, 5G FR1, and 5G FR2) over the grid.

FCC Speedtest (FCC ST) [11] measures uplink/downlink throughput, and round trip latency to the speed test server with the lowest latency. Our measurements went through two servers in Chicago and one server in Miami. We confirmed that both servers in Chicago exhibited similar performance. The throughput and latency numbers reported by the app are end-to-end and include losses introduced by the back-haul.

On each downlink and latency test, the app records signal conditions at the beginning and the end of the test. The signal information is similar to SigCap, but at the time of measurements, only the primary 4G channel and secondary 5G channel is recorded, unlike SigCap which records all available channels. Without this information, we cannot map LAA and CBRS which are implemented as secondary carriers. Also, FCC ST cannot distinguish between 5G FR1 and FR2. Our testing confirms that the collected signal data is similar to SigCap, since the same Android APIs are used by both, but SigCap displays more information. Furthermore, FCC ST cannot be run very frequently, whereas SigCap can passively collect data every 10 secs and hence allows us to create maps with greater temporal and spatial detail. Both apps store their data as a JavaScript Object Notation (JSON) file which can be easily converted to a Comma Separated Values (CSV) file for analysis.

Network Signal Guru (NSG) [25] provides more extensive information compared to the previous two: 4G and 5G frequency, bandwidth, numerology, duplex mode, throughput on several network layers, 5G beam index, Sgnal-to-Interference-and-Noise-Ratio (SINR), block error rate (BLER), modulation, and the number of allocated resource blocks (RBs). However, unlike SigCap and FCC ST, it requires root access which is not available on all phones. Furthermore, the data collected by NSG can only be used for a detailed study of some few cases due to the difficulty in exporting data. For heat-maps and statistical analyses, we use SigCap and FCC ST data.

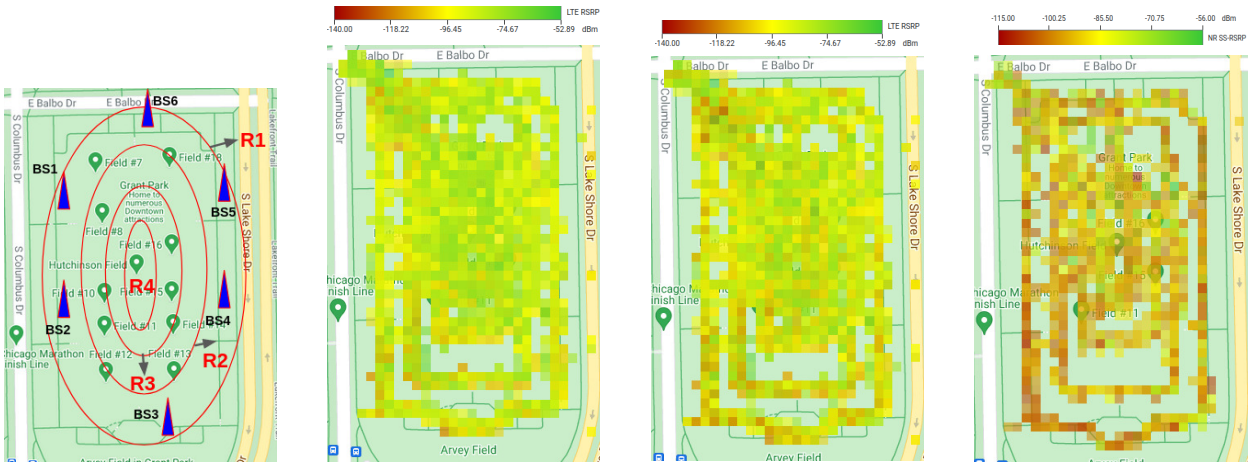

(a) Experiment Location (b) 4G Licensed (c) 4G Lic., LAA and CBRS (d) 5G FR1 and FR2

Figure 1: Hutchinson Field Overview: Verizon, T-Mobile and AT&T 4G, 4G+LAA/CBRS, and 5G Coverage

QualiPoc (QP) [26] also provides extensive information and requires root access, similar to NSG. However, it was only available to phones provided by its vendor at the time (*i.e.,* Samsung S21 Plus). It provides similar information as NSG, but provides easy CSV data export for our analysis. We started using QP in 2022 for the thermal performance measurements described in §7.

Utilizing the aforementioned four applications allows us to extract detailed data from the deployments and analyse cellular network performance in this paper.

4. Coverage, Throughput, and Latency Performance of 4G and 5G Deployed Networks in Chicago

In this section, we present a detailed analysis of our measurement campaign, focusing on coverage, throughput, and latency, conducted in an open field environment in Chicago. The measurements were taken at Hutchinson Field, a prominent area within Grant Park, spanning approximately 0.1 km^2 (shown in Fig. 1a). Table 2 shows dense deployments of Verizon's 4G Licensed, LAA, CBRS, and 5G, followed by fewer deployments by AT&T and T-Mobile. Hutchinson Field is used as a concert venue throughout the year, but on a daily basis it also attracts a multitude of individuals engaging in various activities such as leisurely strolls, picnics, and baseball games. Our measurements enabled us to assess the network's capacity to cater to the demands of a substantial user base, ensuring that the findings from our analysis are applicable to real-world scenarios and reflect the network's performance under significant load.

Table 2: Operator Deployment in Hutchinson Field, Chicago and Miami, Florida (TDD bands in bold).

Operator	Deployment	5G Freq.	5G Op. Bands	4G Op. Bands (LAA:46,CBRS:48)
Verizon	4G+LAA & CBRS, 5G	Low, High	n5, **n260**	2, 4, 5, 13, **46**, **48**, 66
T-Mobile	4G, 5G	Low, Mid	**n41**, n71	2, 4, 7, 12, 66
AT&T	4G+LAA	Low	n5	2, 4, 12, 14, 30, **46**, 66
Verizon (Miami)	4G+LAA, 5G	High	**n261**	2, 4, 13, **46**, 66

4.1. Methodology and Deployment Overview

Table 3 shows the mobile devices used for the measurements and their capabilities. All phones are equipped with SigCap and FCC ST. Since only Pixel 3 and Pixel 5 have root capability, we ran NSG only

Table 3: Devices used for 4G and 5G Measurements in Chicago and Miami

Location	Mobile Device	Network Support
Chicago	2 × Google Pixel 2	4G Licensed Only
	2 × Google Pixel 3	4G Lic., LAA, CBRS
	3 × Google Pixel 5	4G Lic., LAA, CBRS, 5G
Miami	2 × Google Pixel 5	4G Lic., LAA, CBRS, 5G

Table 4: LTE Operating Bands and Frequencies

Band #	Designation	Downlink Freq. Range (MHz)	Bandwidth (MHz)		
			AT&T	T-Mobile	Verizon
B2	mid-band	1930 - 1990	15	15	5
B4	mid-band	2110 - 2155	10	5	20
B5	low-band	869 - 894	N/A	N/A	10
B12	low-band	729 - 746	15	5	N/A
B13	low-band	746 - 756	N/A	N/A	10
B14	low-band	758 - 768	10	N/A	N/A
B30	mid-band	2350 - 2360	10	N/A	N/A
B46	mid-band	5150 - 5925	20	N/A	20
B48	mid-band	3550 - 3700	N/A	N/A	20
B66	mid-band	2110 - 2200	10	15	20

on those phones. We present downlink and uplink throughput data from FCC ST which should be considered an application layer throughput. Additionally, the FCC ST also collects round trip latency data for our analysis. Our AT&T, T-Mobile, or Verizon SIMs includes unlimited data plans[3]. Data was collected by walking with the devices in the four different regions, with different radii, as shown in Fig. 1a: Outer Region Round 1 (R1), Inner Region Round 2 (R2), Inner Region Round 3 (R3) and Inner Region Round 4 (R4).

We present only the latest data collected during May and June 2021, during the afternoon hours with few people (around 20) in the park. In total, we collected 8,353 SigCap data points, each containing information about 4G and 5G signals associated with a GPS coordinate. Specifically, there are 44,683 4G, 22,620 LAA/CBRS, and 3,097 5G data points in the measurement set. In addition, we collected 1,333 FCC ST measurements (708 4G, 386 5G and 239 mixed, where the technology changed during the test), with each containing uplink/downlink throughput and latency results. Using signal data collected by SigCap, we generated coverage maps of 4G, 4G+LAA/CBRS, and 5G as shown in Fig. 1b, 1c, and 1d, respectively. The data we collected is available on our website [27], and is available for download.

4.1.1. 4G deployments in Hutchinson Field

All of the operators that we studied have extensive deployments of 4G in low-band (Bands 5,12,13,14) and mid-band (Bands 2,3,4,7,30,46,48,66). Table 4 shows the corresponding frequencies for the LTE bands. All the operators use 2×2 Multiple Input Multiple Output (MIMO) transmission for 4G, except for T-Mobile, which occasionally uses 4×4 MIMO.

We observe the deployment of AT&T and T-Mobile macro-cell BSs (with low-band channels) outside Hutchinson Field, specifically in the greater Grant Park area. On the other hand, Verizon has a focused deployment of macro-cell BSs and small-cell CBRS and LAA BSs within Hutchinson Field. First, three General Authorized Access (GAA) [28] CBRS (Band 48) channels in 3.56, 3.58, and 3.6 GHz. Second, LAA (Band 46) channels on frequencies that overlap two sets of Wi-Fi-equivalent channels: {36, 40, 44} in

[3]Our subscribed Verizon plan stated that there is a throttling after 50 GBytes for 4G data, and no throttling for 5G data. For AT&T, there is a throttling after 100 GBytes to 4G and 5G data. For T-Mobile, 50 GBytes for 4G and 5G data. In our experiments, we have taken care of data usage using multiple SIMs to avoid data cap throttling.

Table 5: NR Operating Bands and Frequencies

Band #	Designation	Downlink Freq. Range (MHz)	Bandwidth (MHz)		
			AT&T	T-Mobile	Verizon
n5	low-band	869 - 894	5	N/A	10
n41	mid-band	2496 - 2690	N/A	20, 80	N/A
n71	low-band	617 - 652	N/A	20	N/A
n260	mmWave	37000 - 40000	N/A	N/A	100
n261	mmWave	27500 - 28350	N/A	N/A	100

U-NII-1 and {157, 161, 165} in U-NII-3. AT&T has similarly deployed small-cell BSs with LAA channels within the field, with frequencies overlapping two sets of Wi-Fi-equivalent channels: {149, 153, 157} and {157, 161, 165} in U-NII-3. The LAA and CBRS channels are always aggregated in groups of three 20 MHz channels (total of 60 MHz).

The aforementioned LAA channel schemes may lead to coexistence problems: channel 157 overlaps the two sets of AT&T LAA channels and there is also a full overlap between the U-NII-3 channel sets of AT&T and Verizon. Additionally, we observe a dense deployment of AT&T Wi-Fi access points (APs) across the entire 5 GHz unlicensed band in the measurement area. Thus, LAA/LAA and LAA/Wi-Fi coexistence problem may occur. However, the coexistence problem are out of the scope of this paper: we will investigate this in future work.

4.1.2. 5G deployments in Hutchinson Field

Blue triangles in Fig. 1a shows six Verizon's 5G mmWave BSs mounted in lampposts with an average distance between the BSs of 140 m (460 ft). Each BS are equipped with Ericsson radio, with multiple antenna panels, each has a separate PCI with multiple beam indices. There are no AT&T and T-Mobile 5G deployments inside the field; all of their 5G BSs are deployed outside the field. While all there operators have deployed 5G in NSA mode, T-Mobile have an additional SA mode which can be "forced" by using NSG to only connect to 5G cells.

Table 5 shows the corresponding frequencies for the NR bands. T-Mobile and Verizon used the low-band n71 and n5, respectively, with 20 MHz bandwidth. AT&T similarly used n5 but with a bandwidth of 5 MHz. The limited bandwidths of low-band channels are expected due to limited spectrum and to maximize coverage. Therefore, the low-band channels are expected to have a lower throughput performance compared to mid-band channels that have possible 40 to 100 MHz bandwidth. Additionally, due to the limitation of the Pixel 5 phone being able to aggregate only one 5G channel in FR1, the low-band 5G performance is worse than the mid-band 4G at the present time, since 4G can aggregate up to four channels.

In mid-band, T-Mobile have deployed in n41 using 20 and 80 MHz bandwidths. However, only one 5G FR1 channel can be aggregated by the Pixel 5 phone, leading to a diminished performance compared to 4G at the time of measurement. In mmWave, Verizon has densely deployed in n260 (39 GHz) using at most four carriers, each 100 MHz wide[4]. The higher number of aggregated channels predictably leads to a higher bandwidth and a vastly improved throughput compared to mid-band 5G. Using NSG, we observed that Verizon aggregates mmWave channels only if they were transmitted from the same mmWave panel, *i.e.,* they have the same PCI. Finally, all 5G transmissions for all operators used a maximum of 256-QAM modulation, less than the maximum 1024-QAM modulation supported due to the device limitation.

4.2. Performance Comparison

4.2.1. Statistical Analysis of RSRP and RSRQ

RSRP and RSRQ values from SigCap are used to create cumulative distribution function (CDF) plots for each operator. We split the values based on primary channels or other (*i.e.,* non-primary) channels.

[4]AT&T 5G mmWave has been measured in other areas of downtown Chicago but not in Hutchinson Field as of June 2021.

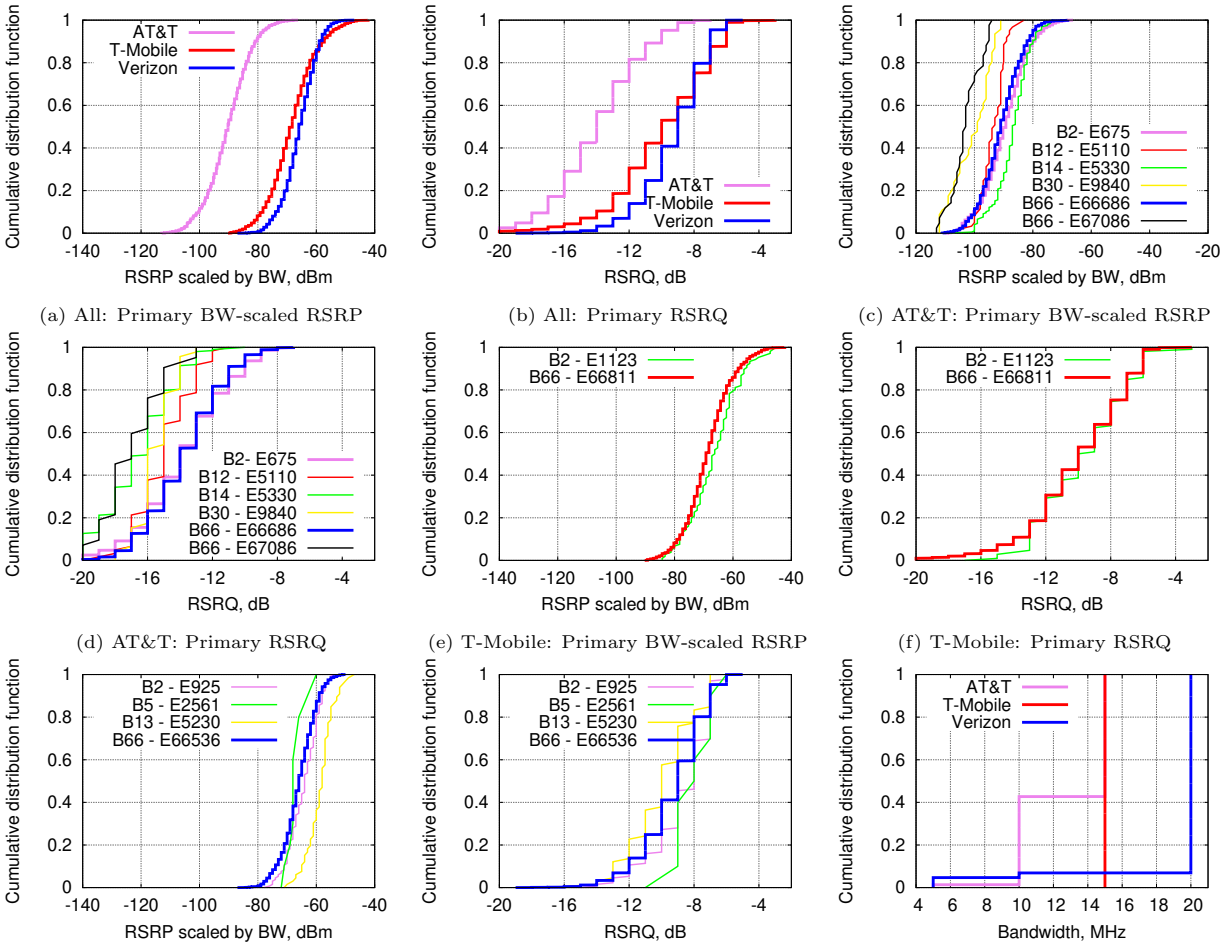

(a) All: Primary BW-scaled RSRP

(b) All: Primary RSRQ

(c) AT&T: Primary BW-scaled RSRP

(d) AT&T: Primary RSRQ

(e) T-Mobile: Primary BW-scaled RSRP

(f) T-Mobile: Primary RSRQ

(g) Verizon: Primary BW-scaled RSRP

(h) Verizon: Primary RSRQ

(i) All: Primary BW

Figure 2: AT&T, T-Mobile, and Verizon in Hutchinson Field: CDF of Primary Channel RSRP, RSRQ, and Bandwidth

Additionally, we present bandwidth-scaled RSRP metric to indicated both coverage and throughput performance at the same time. We calculate bandwidth scaling as $RSRP_{dBm} + 10 * log_{10}(BW_{MHz})$. However, this metric is only available on the primary channel bandwidth since the API has no reliable information on the total aggregated bandwidth. Table 4 and 5 shows bandwidth usage of all operators on each 4G and 5G bands, respectively.

Fig. 2a shows similar BW-scaled RSRP between T-Mobile and Verizon, while AT&T's is around 20 dB lower. Likewise, Fig. 2b shows a higher RSRQ for T-Mobile and Verizon, with AT&T around 4 dB lower. These CDFs indicate the better 4G coverage performance of T-Mobile and Verizon compared to AT&T's, which is further confirmed by throughput analysis presented in the next sub-section.

Next, we present the CDFs of BW-scaled RSRP sorted by Band (B) and EARFCN (E), to show each operator's channel selection performance. Fig. 2c shows the CDF for AT&T, which uses 5 Bands (2,12,14,30,66) as its primary channel, with highest occurrence in bolder lines for B2 (E675, 57% of data) and B66 (E66686, 33% of data), while Fig. 2d shows the RSRQ counterparts. We noticed a possible inefficiency of channel selection: the BW-scaled RSRP of B14 is the highest thus it may lead to a better choice for the primary channel. However, the difference in RSRP is negligible, its RSRQ is around 3 dB lower, and it only provides 10 MHz bandwidth (for comparison, B2 has 15 MHz). Additionally, while the propagation of B14 is better due to the low-band frequency, it may also leads to more neighboring cell interference when the same channel is used on neighboring cells. Fig. 2e and 2f show the CDF of BW-scaled primary RSRP and RSRQ for

T-Mobile, respectively. There are only two choices for primary channel bands, with B66 (E66811, 92% of data) as the majority. This choice seems justified from the BW-scaled RSRP CDF and the RSRQ CDF. For Verizon, Fig. 2g shows the CDF of BW-scaled primary RSRP with B66 (E66356, 93% of data) is selected more often than B2 and B13 that have higher RSRP. Additionally, Fig. 2h shows a lower RSRQ for B66 compared to B2 and B5. Again, the choice of B66 as primary channel is optimal since the alternative B13 shows slightly lower RSRQ than B66, while B2 and B5 has 5 and 10 MHz bandwidths, respectively, which is lower than 20 MHz available in B66. The above data indicate that each operator's primary channel choice is based primarily on optimizing RSRP, RSRQ, and bandwidth.

Lastly, Fig. 2i shows the CDF of the primary channel bandwidth. Verizon has the highest available bandwidth for its primary channel, followed by T-Mobile and AT&T. While these primary channel bandwidth, RSRP, and RSRQ analysis can be an indicator of throughput performance, it is missing a key component, *i.e.,* carrier aggregation, which is not available on SigCap. However, it provides insight into the deployment quality: the higher the primary bandwidth and RSRP, the more likely that the operator will have good coverage and throughput. This is corroborated by the throughput analysis in the next sub-section.

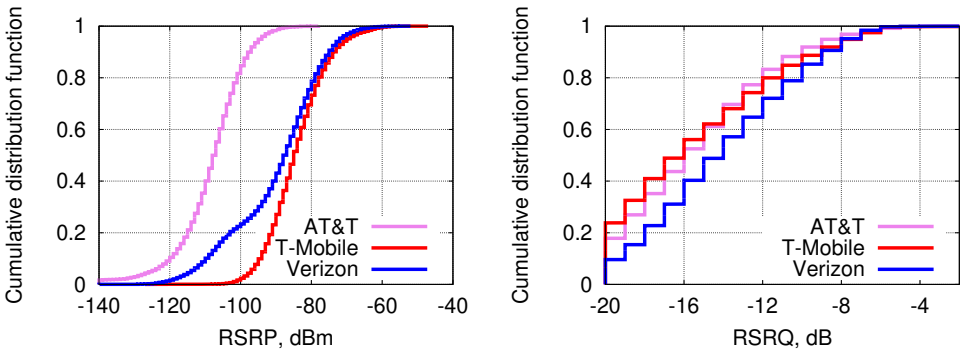

(a) All 4G Channels' RSRP (b) All 4G Channels' RSRQ

Figure 3: AT&T, T-Mobile, and Verizon in Hutchinson Field: CDF of 4G RSRP & RSRQ

To compare coverage performance, Fig. 3a and 3b show the CDF of RSRP and RSRQ for all 4G Licensed carriers (*i.e.,* primary, secondary, neighboring) in the area. Based on the CDFs, T-Mobile has the best 4G licensed coverage, followed closely by Verizon. However, the RSRQ CDF shows a better overall channel quality on Verizon compared to T-Mobile. On the other hand, AT&T's RSRP and RSRQ values indicate inferior coverage, partly due to the fact that the cells are mostly deployed outside Hutchinson Field.

Fig. 4a and 4b show the 5G-RSRP and 5G-RSRQ CDF, which are the RSRP and RSRQ measured on the 5G channel's synchronization signal block. Unlike previous analysis, the 5G-RSRP cannot be scaled

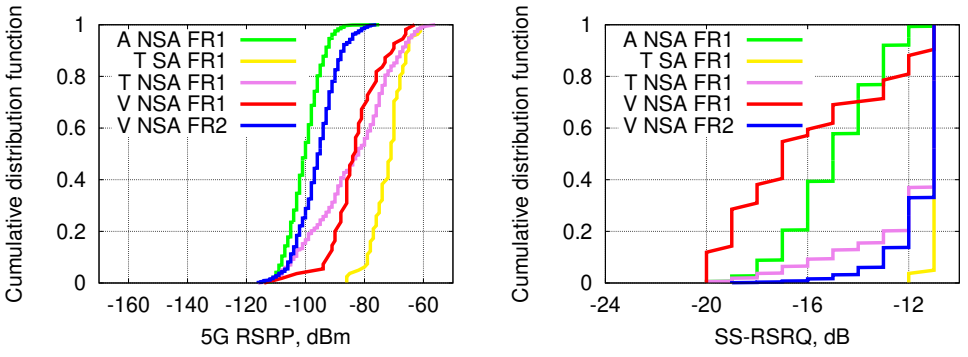

(a) 5G RSRP (b) 5G RSRQ

Figure 4: AT&T (A), T-Mobile (T), and Verizon (V) in Hutchinson Field: CDF of 5G RSRP & RSRQ

with bandwidth since the SigCap does not provide this information. Overall, the 5G-RSRP of the FR1 bands is higher than FR2 due to the difference in operating frequency and its propagation. First, since 5G-RSRP data of T-Mobile NSA in FR1 are a combination of low-band (n71) and mid-band (n41), its CDF does not follow Gaussian distribution. In SA mode, T-Mobile shows the highest 5G-RSRP since the device can only connects to the low-band n71 in SA mode. Our devices were rarely connected to Verizon 5G in FR1, possibly due to the lower 5G-RSRQ values as shown by Fig. 4b. We further utilized NSG to block the dive from connecting to 5G mmWave. However, the device would not connect to 5G FR1, it rather connects to 4G+LAA/CBRS, perhaps because the latter configuration provided higher bandwidth due to carrier aggregation. Finally, AT&T FR1 displays a very low 5G-RSRP and 5G-RSRQ which indicating inferior 5G coverage in the area.

While LAA and CBRS information was collected, we do not include them in the comparisons since there is a substantial difference in transmit power compared to the licensed channels; the U-NII-3 spectrum used by LAA, allows a maximum of 30 dBm transmit power, while CBRS allows a maximum of 47 dBm in outdoor deployments.

Using NSG data, we extracted the average RB allocation per device as an indicator of network load, as shown in Fig. 5a. There are slightly fewer RBs allocated on Verizon's licensed carrier compared to the other operators, indicating a higher load or a load balancing scheme by allocating more RBs on the secondary LAA/CBRS/5G carriers. However, the difference is insignificant, and we can conclude that the network load is similar for all operators during the measurements.

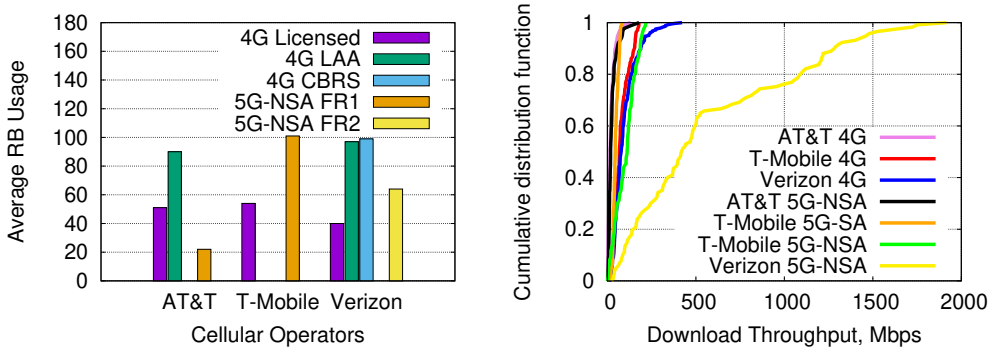

(a) RB Usage

(b) CDF of Downlink Throughput

Figure 5: AT&T, T-Mobile, and Verizon in Hutchinson Field: RB Usage and Downlink Throughput

4.2.2. Throughput and Latency Performance, using FCC ST

On each download, upload, and latency test, FCC ST recorded the type of cellular technology used at the beginning and end of the test (*i.e.,* 4G, 5G). We sorted the data based on the cellular technology used and removed data where the technology switched between 4G and 5G during the test. Additionally, we added SA and NSA labels on the 5G data based on the 5G mode forced on the device: AT&T and Verizon always use NSA mode, while T-Mobile uses NSA by default but can be forced to use SA mode.

Fig. 5b shows the downlink throughput CDF of AT&T, T-Mobile, and Verizon in 4G and 5G. AT&T had the worst 4G and 5G throughput in Hutchinson Field, further confirming the coverage and bandwidth analysis from prior section. Verizon 5G mmWave had the best throughput: the maximum throughput achieved was 1.92 Gbps, which is constrained by Pixel 5's support of a maximum of four aggregated mmWave channels [5]. Most of the FCC ST data for Verizon 5G was captured using mmWave since there was a sparse deployment of 5G in FR1.

[5]Other 5G phones may have higher maximum downlink throughput due to greater mmWave aggregation capability, *e.g.*, Samsung Galaxy S21 Plus supports a maximum of eight aggregated mmWave channels.

The next best downlink throughput performance is achieved closely by Verizon 4G and T-Mobile in 4G and 5G-NSA. Both Verizon and T-Mobile achieved a very similar performance in 4G, which correlates to the similarity of their 4G RSRP, RSRQ, and primary bandwidth distribution. However, Verizon delivered the highest 4G throughput of 421 Mbps due to LAA/CBRS usage, which is better than the highest 5G throughput in FR1 of 219 Mbps, achieved by T-Mobile 5G-NSA. Due to device limitations, only a maximum of one secondary 5G FR1 carrier can be aggregated. Thus the 80 MHz is available on T-Mobile's 5G Band n41 is comparable to its 4G counterparts, leading to diminished throughput increase between 5G-NSA and 4G. Similarly, T-Mobile 5G-SA offered low throughput due to the single 5G channel usage without any aggregation. The average downlink throughput recorded in the Hutchinson Field region for all operators are as follows: (i) AT&T: 20.7 Mbps and 27.1 Mbps in 4G and 5G-NSA, respectively; (ii) T-Mobile: 77.2 Mbps, 46.2 Mbps, and 101.3 Mbps in 4G, 5G-SA, and 5G-NSA, respectively and (iii) Verizon: 95.8 Mbps and 574.4 Mbps in 4G and 5G-NSA, respectively. Verizon achieved the best downnlink throughput performance due to its usage of mmWave.

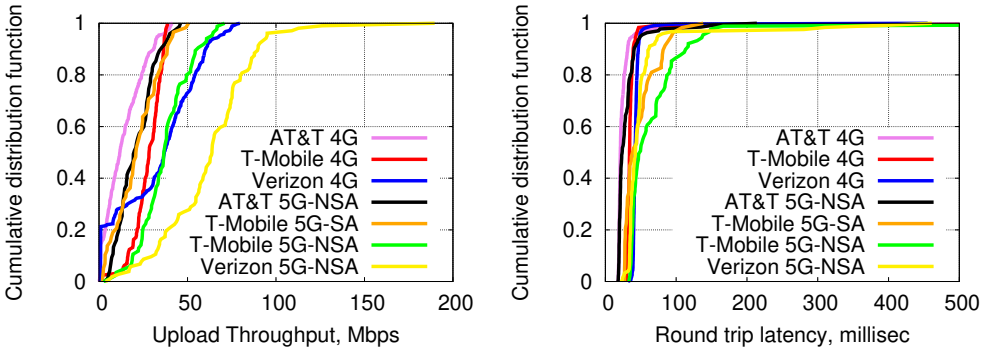

(a) CDF of Uplink Throughput

(b) CDF of Round Trip Latency

Figure 6: AT&T, T-Mobile, and Verizon in Hutchinson Field: Uplink Throughput and Latency

Fig. 6a shows the CDF of uplink throughput of the three operators over 4G and 5G. In uplink transmission, we only observe the primary channel utilization, regardless of how many channels aggregated. Thus, 5G NSA performs the better overall compared to 5G SA and 4G since both primary channel of 4G and 5G is utilized. First, AT&T 4G performs the worst due to lack of coverage. Interestingly, AT&T 5G NSA and T-Mobile SA performs similarly on uplink throughput: the prior mostly aggregates 10 MHz B66 and 5 MHz n5, while the latter only utilized 20 MHz n71. Verizon 5G-NSA predictably performs the best due to utilizing a mmWave channel with 100 MHz bandwidth. The following are the average uplink throughput values for all operators: (i) AT&T: 13.9 Mbps and 20.9 Mbps in 4G and 5G-NSA, respectively; (ii) T-Mobile: 27.3 Mbps, 21.5 Mbps, and 37.2 Mbps in 4G, 5G-SA, and 5G-NSA, respectively and (iii) Verizon: 33.1 Mbps and 62.3 Mbps in 4G and 5G-NSA, respectively.

Lastly, Fig. 6b similarly shows the CDF of the round trip idle latency of the three operators. The median values are: 30.5 ms and 30.7 for AT&T 4G and 5G-NSA, respectively; 44.1 ms, 48.4 ms, and 74.8 ms for T-Mobile 4G, 5G-SA and 5G-NSA, respectively; 44.1 ms and 54.4 ms for Verizon 4G and 5G-NSA, respectively. Overall, the poor latency performance in 5G-NSA shows the problem of non-optimal deployment of 5G-NSA [3, 29], causing additional overheads due to dual connectivity. It should be noted that while the latency measurement is end-to-end, the effects of back-haul are the same for all the operators since all the latency tests were conducted via the same two servers located in Chicago. We also did not notice any significant difference in throughput and latency between tests conducted over the two servers.

It is clear that 5G mmWave provides a significantly improved throughput performance, but the latency performance could be improved. The dense deployment in Hutchinson Field has provided a very good 5G mmWave coverage even with the directional nature of mmWave transmissions: there are 6 BSs over 0.1 km^2, with average distance of 140 m. However, the directional nature also results in a higher variance of 5G mmWave throughput as seen in Fig. 5b.

5. 5G mmWave Measurements in Miami

Next, we present a focused analysis of 5G mmWave measurements conducted in Miami. Our primary objective is to address the impact of various scenarios on 5G mmWave throughput performance, focusing on three key factors: varying distance under different environments, human body blocking, and the presence of multiple devices. While the prior Chicago analysis encompassed a broader scope, assessing coverage and performance in a more general sense, the Miami study focused on the intricate details of 5G mmWave technology in targeted scenarios. Together, these two analyses contribute to a comprehensive understanding of 5G network performance, shedding light on both the overall coverage capabilities and the nuances of specific technologies and usage scenarios.

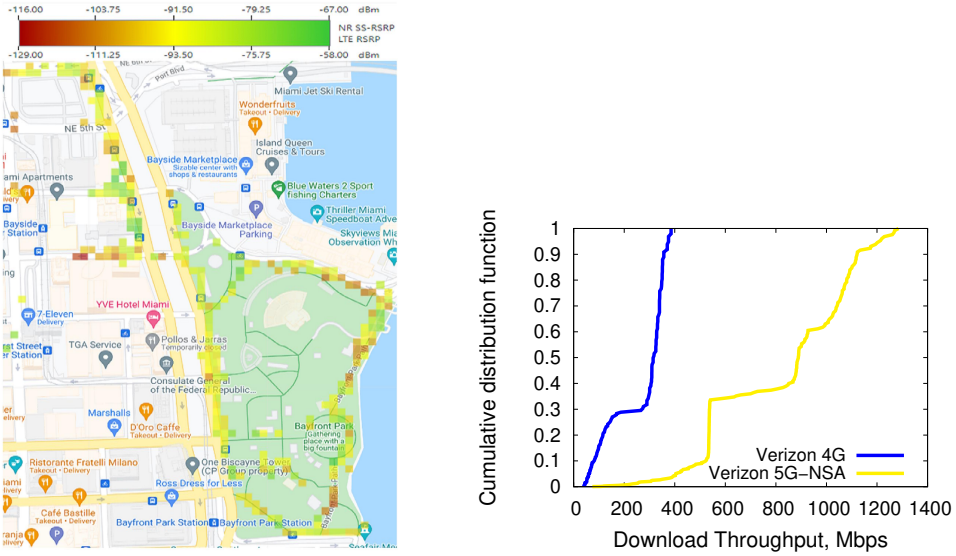

(a) 5G Coverage map in Miami downtown

(b) Throughput of 4G+LAA Vs 5G

Figure 7: Overview of 5G mmWave Deployment at Miami, Florida and Throughput Comparison with 4G+LAA

5.1. Methodology and Deployment Overview

We utilized two Pixel 5 phones as summarized in Table 3, all equipped with SigCap, FCC ST, and NSG. Fig. 7a shows the area of measurements in the park and city streets within the downtown area. We measured the Verizon network while walking between January and June 2021. Verizon has a diverse deployment in downtown Miami with a mix of 4G, 4G+LAA, and 5G mmWave, as was previously summarized in Table 2. We did not detect CBRS in the measurement area at the time[6]. On the other hand, we observed mmWave operating in Band n261 (28 GHz) with a bandwidth of 400 MHz (*i.e.,* 4 × 100 MHz). The selection of mmWave frequency is interesting since it contrasted the n260 (39 GHz) deployment at Hutchinson Field, *i.e.,* a shorter wavelength, possibly chosen due to the dense deployment.

Fig. 7a shows the *coverage map* of 5G deployment in downtown Miami, we focus on 7 locations labelled M1-7. 4G+LAA is also widely deployed in the same area. Fig. 7b the throughput distribution (as collected by FCC ST) of 5G compared to 4G+LAA on all locations. The 5G throughput gain is in the range of 4× to 14×, compared to 4G. Also, we observe a similar 5G throughput distribution between Miami and Hutchinson Field. However, 4G throughput at Miami is higher due to high occurrence of LAA aggregation.

[6]According to the FCC database, Verizon has 30 MHz PAL license in the Miami-Dade county area. However, we did not observe CBRS deployment at the time of measurements.

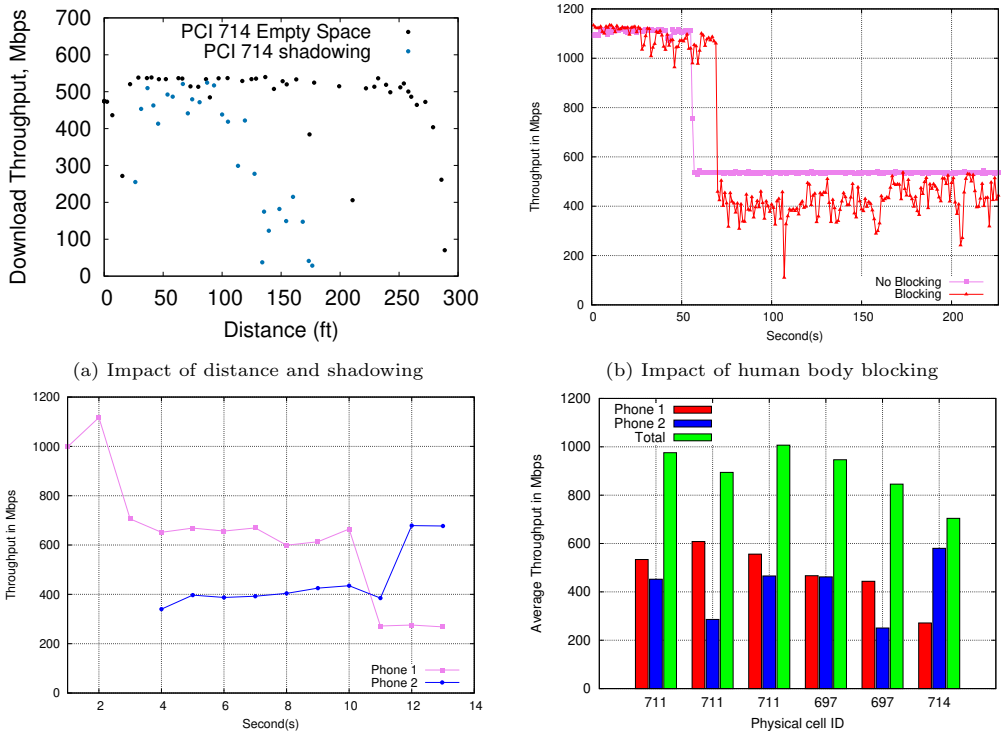

(a) Impact of distance and shadowing

(b) Impact of human body blocking

(c) Throughput between 2 phones

(d) 2 phones throughput on multiple PCI

Figure 8: 5G mmWave-focused measurements at Miami, Florida

5.2. 5G mmWave-focused Measurements in Miami

Next, we selected a measurement area at the Bayfront Park near a mmWave BS. We initialized HTTP download of a 10 GB dataset file from [30] and utilized NSG to capture PHY throughput metrics. Fig. 8a shows the impact of *distance* on the 5G mmWave coverage. In an empty space (*i.e.,* line-of-sight environment), we observe a maximum throughput around 550 Mbps up to 250 ft (76 m). At 300 ft (91 m), the throughput drops to 60 Mbps. Furthermore, having trees (*i.e.*, shadowing effect) reduces the coverage range down to 125 ft (38 m), *i.e.,* a 50% drop in coverage. Next, we conducted experiments to quantify the impact of human body blocking. We define 2 different trials: one trial had the user's body blocking the phone, while the other did not. We controlled both trials by standing a fixed distance to the tower with no other obstructions, thus the phones were connected to the same PCI 714 and the same beam number throughout the trial. Fig. 8b shows the results of the human body blocking experiment, we observe a lower throughput and a higher variance on the blocked phone. The average degradation due to human body blockage is about 20%. At the 50 s and 70 s marks, we observe a reduction of throughput on both mark due to temperature throttling. This problem will be further expounded on §7, but in summary: due to an increase in device temperature initiated by the mmWave transmission, the phone throttled its connection by reducing from 4 channel aggregation to 1 channel. Further, the blocked phone experienced the throttling later compared to the non-blocked phone, possibly due to the lower throughput reflects on the lower CPU/modem load, thus the phone took longer to heat up. Another possible factor is the ambient temperature: even when the experiments were conducted at the same day, there can be a small variation of ambient temperature between the experiments. Unfortunately, since thermal factors was not our focus at the time, we did not take note on the ambient temperature.

Finally, we focus on the impact of having two simultaneously served phones. In this experiment, we use two Google Pixel 5 phones, were held within arms-length of one another near a cell tower, thus both phones are connected to the same PCI. Phone-1 first starts its downlink transmission, then followed by phone-2

four seconds later. Fig. 8c depicts the throughput achieved by each of the two phones over time. We observe phone-1 starts with a high throughput, indicating that all RBs are allocated to phone-1. Once phone-2 starts, the throughput of phone-1 drops given that the total resource blocks are now shared between the two phones. Fig. 8d shows the repeated experiments over different PCIs. Over the total of 6 trials on each PCI, we observe comparable the throughput values between the two phones

6. Outdoor-to-Indoor Performance Analysis of a Commercial Deployment of 5G mmWave

As the previous measurement campaigns were focused on outdoor measurements of 5G, there is a lingering research question regarding the performance of 5G indoors while the BSs are outdoors. To address this, we conducted indoor measurements at a University of Chicago dormitory building which conveniently has a Verizon mmWave BS across the street. We first present an indoor performance analysis of 5G NR among the three US operators (*i.e.,* AT&T, T-Mobile, Verizon). Then we present a focused analysis of Verizon's 5G mmWave performance (in terms of downlink throughput, uplink throughput, and latency) under varying window opening gap sizes.

6.1. Methodology and Deployment Overview

Fig. 9a shows the indoor measurements site: Woodlawn Residential Commons, a 7 storied dormitory building at the University of Chicago. The building is located at 1156 E 61st St, conveniently beside (∼25 m) a Verizon mmWave BS deployment. The building has completed its construction in 2020, with an unknown type of glass used in the windows. However, we believe that the windows are most likely Low-E glass given the very recent construction of the building. We conducted indoor measurements in various rooms in July 2021, with special permission from the university.

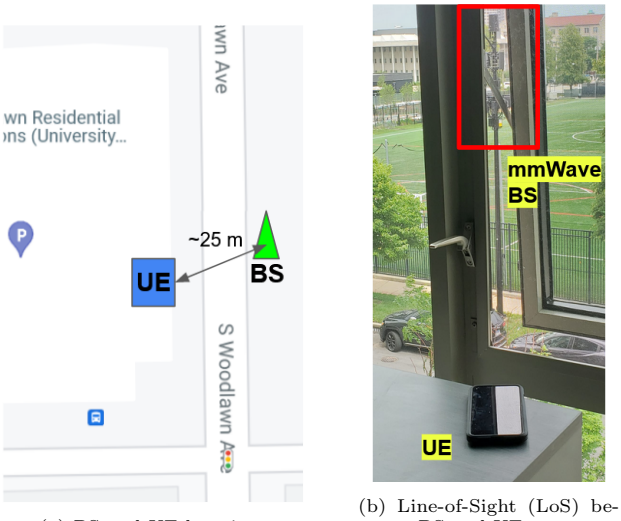

(a) BS and UE location.

(b) Line-of-Sight (LoS) between BS and UE.

Figure 9: Outdoor-to-indoor (OtI) measurement location.

For the UEs, we utilized up to three Google Pixel 5 phones, equipped with AT&T, T-Mobile, and Verizon SIMs with unlimited data plans. As with our prior works, we utilized multiple SIMs to avoid data cap throttling. Each phones passively ran SigCap and NSG on Google Pixel phones to collect detailed signal parameters. Conversely, the FCC ST app is utilized to actively run download, upload, and latency test on every minute.

Table 6: Indoor Cellular Reception at UChicago Dormitory. SA: standalone, NSA: non-standalone

Operator	**5G NR Mode**	**5G NR Band (Max. Bandwidth)**	**4G LTE Band (Max. bandwidth)**
AT&T	NSA	n5, 850 MHz (5 MHz)	2 (15 MHz), 12 (10 MHz), 14 (10 MHz), 17 (10 MHz), 30 (10 MHz), 66 (10 MHz)
T-Mobile	NSA	n41, 2.5 GHz (100 MHz), n71, 600 MHz (20 MHz)	2 (15 MHz), 66 (15 MHz)
Verizon	NSA	n5, 850 MHz (10 MHz), **n261, 28 GHz (400 MHz)**	2 (5 MHz), 13 (10 MHz), 66 (20 MHz)

6.2. Comparison of 5G NR performance among different bands and operators

We surveyed 2nd - 7th floors of the building (the accessible floors for measurement) by placing three Pixel 5 phones on a cart. Each phone is equipped with AT&T, T-Mobile, and Verizon SIMs. We then turn the passive and active measurements, and wheeled the cart through corridors and rooms in the east side of the building which faces Verizon mmWave BS. All windows were shut during this survey. Table 6 shows the summary of indoor reception in the dormitory: there are a varying number of different 4G LTE and 5G NR channels, in low-, and mid-bands. First, all 5G deployments are NSA, and we were not able to "force" SA mode on T-Mobile by disabling LTE connection using NSG, possibly due to lack of SA deployment nearby. We were able to receive low-band n5 deployments by AT&T and Verizon indoors. On the other hand, T-Mobile has a more varying 5G deployment: a low-band n71 and and a mid-band n41 can be received indoors. For high-band/mmWave, we observed Verizon 5G mmWave reception indoors in 28 GHz with a maximum bandwidth of 400 MHz using 4-channel aggregation (CA). However, the reception was very limited and with poor RSRP. The mmWave signals are not available at all on the 7th floor possibly due to the downwards orientation of mmWave panels. While on the 2nd - 6th floors, there were only a handful of rooms that were LoS to the BS that could receive 5G mmWave signals when the window was open. This is likely due to the Low-E glass used in the windows. For AT&T and T-Mobile channels, we are uncertain about the exact location of their outdoor BSs. However, we confirmed from the PCIs that all indoor reception of Verizon's NR and LTE channels were being transmitted from the BS on the pole right outside the building. This particular Verizon BS also shows an outdoor availability of Bands 46 (LAA) and 48 (CBRS). However, we also did not detect both bands indoors, most likely due to the lower transmitted power allowed in these bands.

Since all 5G deployments are NSA, each 5G connection consists of LTE and NR channels. For AT&T and T-Mobile, the choice of LTE and NR channels are predictable based on the RSRP. However for Verizon, but there was a difference in the LTE primary channel used depending on the NR band. When NR Band n5 (low-band, bandwidth 10 MHz) was used, the LTE primary on the DL was always Band 66 with a bandwidth of 20 MHz, whereas when NR Band n261 (mmWave) was being used, the LTE primary carrier on the DL was either band 66 (20 MHz) or band 13 (10 MHz). This difference in choice of LTE primary channel has an effect on overall DL and UL throughput as shown on the later analysis.

DL Throughput comparison: From Figs. 10a and 10d we see that NR clearly delivers significant DL throughput improvements over LTE, especially for T-Mobile and Verizon. AT&T NR performance is limited due to the low-band only deployment using only 5 MHz of bandwidth, compared to T-Mobile's 100 MHz at 2.5 GHz and Verizon's 10 MHz. Since Verizon 5G mmwave was not received during these tests, the DL throughput is solely via aggregation of LTE and low-band NR. With no mmWave reception, T-Mobile NR DL throughput is superior to Verizon's, even though the Verizon BS is very close to the building. Once again, this survey demonstrates the severe limitation of indoor 5G mmWave reception.

UL Throughput comparison: From Figs. 10b and 10e we see that here too NR clearly delivers significant UL throughput improvements over LTE, for all operators. There is a clear advantage of Verizon UL, most likely due to the aggregation with the 20 MHz band 66 LTE carrier and the proximity of the location to the BS enabling higher modulation-coding settings. For example, the 80 Mbps throughput is due to 65 Mbps

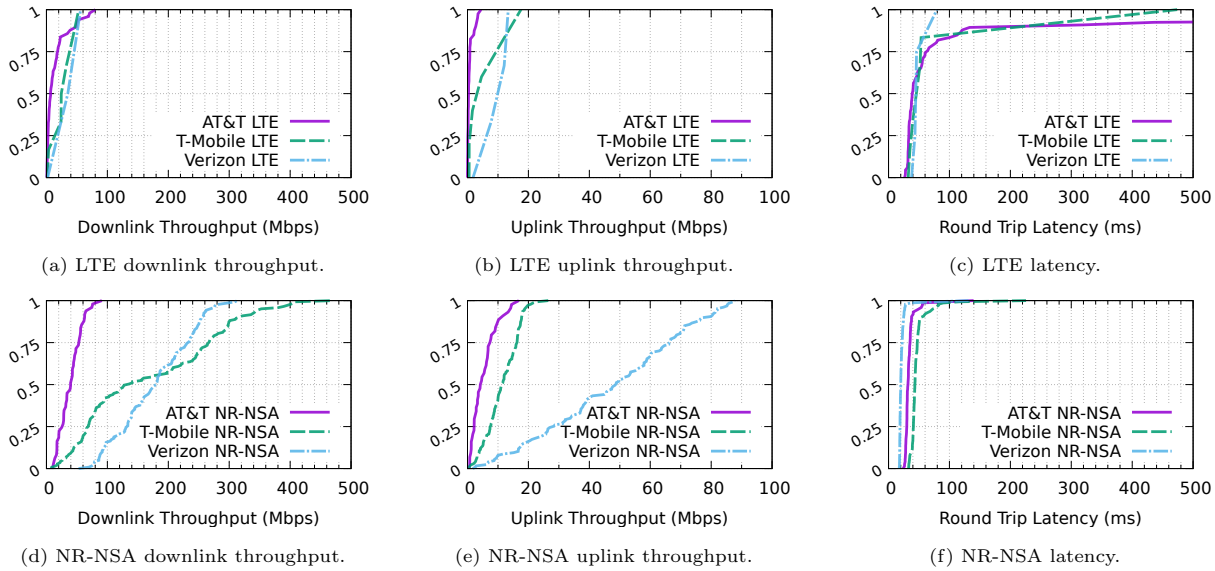

(a) LTE downlink throughput. (b) LTE uplink throughput. (c) LTE latency.

(d) NR-NSA downlink throughput. (e) NR-NSA uplink throughput. (f) NR-NSA latency.

Figure 10: Indoor survey of AT&T, T-Mobile, and Verizon in terms of throughput and latency performance.

over band 66 and only 15 Mbps over NR band n5.

Latency comparison: From Figs. 10c and 10f we see that there is not an appreciable reduction in latency with NR, though overall Verizon latency with NR is the lowest. However previous results already noted that NR latency was lower in the low-band compared to mmWave, and these results only include low-band NR. It should also be noted that since most of these measurements were over the NSA mode of NR, the latency could be higher due to the dual connectivity, channel aggregation and the use of the 4G core network. As SA with the new 5G core begins to be deployed, we anticipate that the latency results will improve.

6.3. Performance of 5G mmWave as a function of window opening gap size

We first did a preliminary measurement where we open the window of each room facing Verizon 5G mmWave BS on all floors, and we observe the best 5G mmWave performance in room E206 on the 2nd floor. Therefore, we utilized the room for additional experiments to quantify performance as a function of the window opening gap size. The UE was placed on top a desk with line-of-sight (LoS) to the mmWave BS as shown in Fig. 9b. We then vary the width of the window opening as shown in Fig. 11, where Gap 1 is the widest gap and Gap 4 is a fully closed window. For each gap scenario, passive and active measurements were taken over 15 minutes.

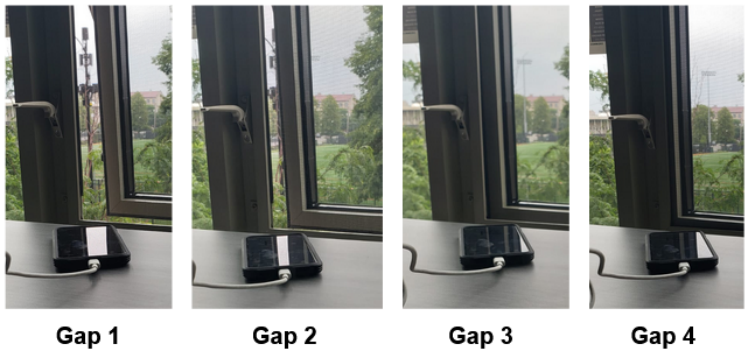

Figure 11: Window opening gap diagram.

Table 7: NR Reception on Different Gaps

Gap #	NR Channels	Avg. RSRP	Avg. RSRQ
Gap 1	4 × n261 (400 MHz)	-89.52 dBm	-11 dB
Gap 2	4 × n261 (400 MHz)	-98.98 dBm	-11 dB
Gap 3	1 × n5 (10 MHz)	-74.34 dBm	-11 dB
Gap 4	1 × n5 (10 MHz)	-75.60 dBm	-11 dB

First, we present coverage performance on each gap scenario, in terms of RSRP and RSRP as shown in Table 7. The phone was connected 5G mmWave for Gaps 1 and 2, then it was handed-over to 5G low-band for Gaps 3 and 4. Thus, the overall lower RSRP values in Gaps 1 and 2 are due to the shorter wavelength. While the average RSRQ stays constant on all Gaps, we observe reduction of average RSRP from Gap 1 to Gap 2, and from Gap 3 to Gap 4. The decrease of RSRP by 10 dB from Gap 1 to Gap 2 illustrates the vulnerability of mmWave connections when encountering obstructions such as Low-e glass. While on Gaps 3 and 4, there is not much difference in the RSRP since Band n5 at 850 MHz propagates very well indoors and is less dependent on the window gap size.

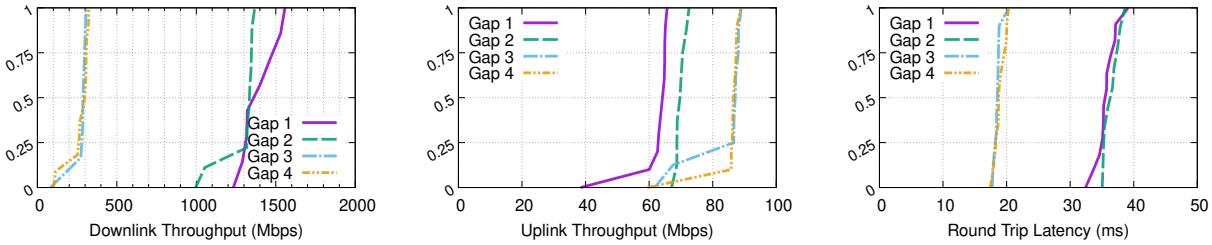

(a) Downlink throughput. (b) Uplink throughput. (c) Round trip latency.

Figure 12: Throughput and latency performance as a function of window opening size

Fig. 12 show the throughput and latency performance on the varying window gap scenarios. For downlink throughput as shown by Fig. 12a, we observe a better performance on Gaps 1 and 2. This can be explained due to the high bandwidth of the 5G mmWave connection. Contrarily, Fig. 12b and 12c shows the best UL throughput and latency is achieved by Gaps 3 and 4, possibly due to the LTE channel aggregation of the NSA mode. These results shows that a true Gbps downlink throughput over 5G mmWave can only be delivered indoors with unobstructed LoS. On the other hand, the mmWave connection has worse uplink throughput and latency performance compared low-band NR.

While our work on the performance of OtI mmWave is comparable to a prior work by [21], the methodologies employed are very different, leading to different and contradictory conclusions. First, the throughput results reported in [21] are based on predictions from signal strength measurements using a specific channel sounder that utilizes a continuous-wave tone at 28 GHz as the sounding signal, rotating horn antennas on the receiver and omni-directional transmit antennas: very different from actual operating conditions of 5G mmWave. On the other hand, we specifically conducted our measurements on deployed 5G mmWave systems using consumer handsets which captures real-world conditions such as beam-management using phased-arrays at both BS and UE, wide-bandwidth operation (400 MHz), and handset limitations. Additionally, our performance metrics are measurements of throughput and latency over all of the network stack, thus it includes overheads due to the MAC, transport, and network layers. Contrarily, we believe the prior work lack in accounting the effects of the intermediate layers in their prediction. These major differences in measurement methodologies and environment has lead to the contradictory results: the prior works [21] shows a prediction of 500 Mbps in a building with high-loss glass (*i.e.,* Low-e) windows for 90% of users located up to 49 m away from the BS, while our results demonstrate that in a building with Low-E glass windows located about 25m from a 5G mmWave BS, there is no 5G mmWave connectivity at all through closed windows and limited connectivity in a few locations with the window open. Further, the measure-

ments over different bands and operators demonstrate that the low and mid-band 5G NR can still provide DL throughput of up to 400 Mbps even when 5G mmWave is unavailable in the building.

7. Impact of Device Thermal Performance on 5G mmWave Communication Systems

In all of our prior 5G mmWave measurements, we noticed a persistent pattern of throttling occurring after a minute of initiating download traffic. This pattern consist of a switch from 5G mmWave connection with all mmWave antenna used (4 or 8 antennas depending on the phone model), to only one mmWave antenna used, then a full handover to 4G. This introduced a problem when conducting a 5G mmWave-focused measurement since we were only interested on 5G mmWave performance in that scenario. In a prior work, we have identified Android's hardware temperature API which describes CPU, GPU, and skin temperature, with its respective "throttling" threshold [9]. We have previously shown that the increase in skin temperature (over its throttling threshold) correlates with the throughput reduction in 5G mmWave network. However, this temperature API is not available in some phones, thus we aim to show the correlation between device thermal and 5G mmWave throughput performance by capturing the device temperature using IR camera.

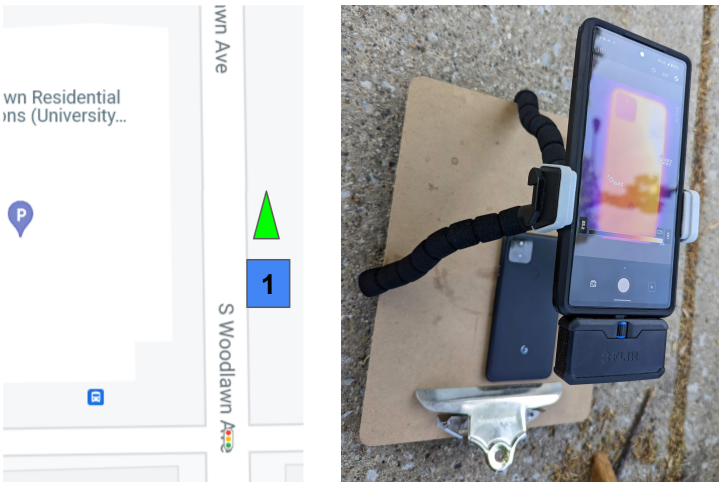

(a) Measurements Location (b) IR Thermal Experiment Setup

Figure 13: IR Experiment Setup and Measurements Location in Chicago.

7.1. Measurement Tools and Methodology

The experiments were conducted in Chicago near a Verizon mmWave beside the Woodlawn Commons Dormitory (as shown in Fig. 13a) in May 2022, with ambient temperature of 30° C. Fig. 13b depicts the mmWave UE, *i.e.,* a Samsung S21 Plus phone without a phone case, mounted ∼17.5 cm below a FLIR One Pro LT IR camera, which capture a thermal image every one second. We utilized the Qualipoc measurement tool[26] which is a root-based measurement tool very similar to NSG, but with easier data extraction. The S21 Plus phone is running Android 12 on a Verizon network with an unlimited data plan[7]. The Verizon 5G mmWave network at the location utilizes Band n261 at 28 GHz. Fig. 13a shows Location 1 where the phone is connected to the 5G mmWave BS (green triangle) with a good signal condition (∼90 dBm RSRP). On each of the 10 minutes experiment run, the phone either initiates a downlink or upload traffic while the Qualipoc app captures PHY layer data (e.g., per channel SS-RSRP, SS-RSRQ, and throughput) and HTTP/application layer throughput. The downlink traffic is generated by initiating HTTP download of a 10 GB dataset file sourced from [30]. Conversely, the uplink traffic is generated by uploading a 1 GB bin file using the HTTP POST method to [31]. We ran downlink and uplink experiments each with two different phone orientation *w.r.t.* the BS, thus a total of 4 experiment runs.

[7]Subscribed Verizon plan indicates a throttling after 50 GBytes for 4G and 5G low/mid-band data, and no throttling for 5G mmWave data.

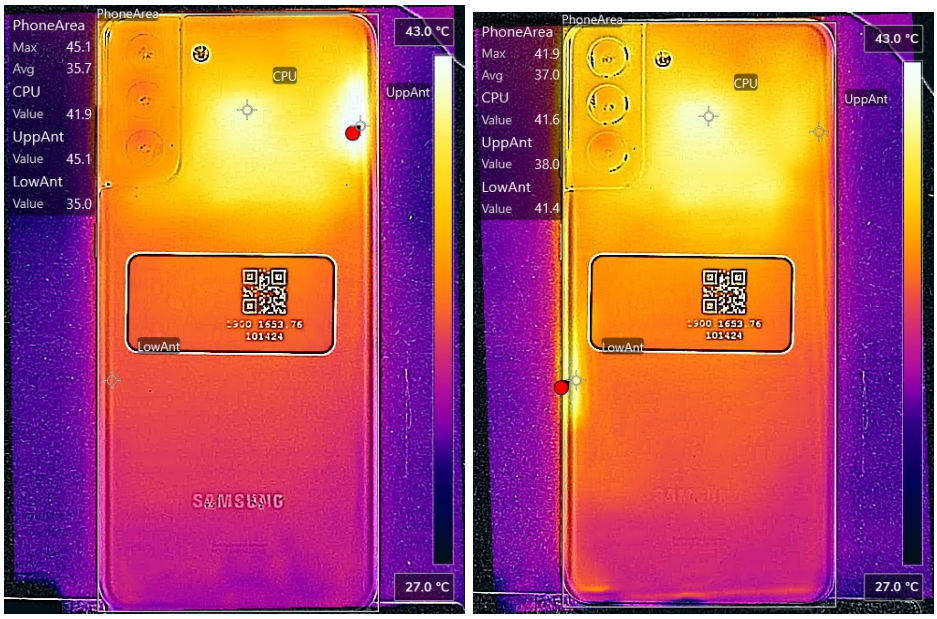

(a) IR thermal capture on O1 (BS is to the right of phone).

(b) IR thermal capture on O2 (BS is to the left of phone).

Figure 14: IR thermal captures in orientations O1 and O2.

7.2. Thermal Performance Investigation with IR Camera

To capture thermal performance of the phone, we define a "hot-spot", which is a spot in the captured data with high possibility of heating up during mmWave transmission. The hotspot is defined as an x-y coordinate relative to the phone frame. The relative coordinate is used to ensure that the data captured on different experiment runs are comparable. Fig. 14a and Fig. 14b show representative thermal images[8] of two orientations of the phone *w.r.t.* the BS: O1 and O2, respectively. For O1, the BS is located to the right of the phone, while for O2, the BS is located to the left of the phone. Fig. 14a shows a bright spot at the upper-right part where the upper mmWave antenna is located, while Fig. 14b shows a bright spot at the middle-left part where the lower mmWave antenna is located. There is also a considerably bright spot at the upper-center part, which may correspond to the CPU and modem area. Thus, we defined three hot-spots based on the observations: CPU and modem area (CPU), upper mmWave antenna (UppAnt), and lower mmWave antenna (LowAnt). On each of the thermal photo, we extracted the temperature values of each hot-spots. It should be noted that none of these hot-spots corresponds to Android API's CPU, GPU, or skin temperature value. The Android API captures thermal data from physical sensors and the exact location of the sensors are unknown to us.

Fig. 15a shows downlink throughput and temperature data over time for orientation O1. We observe an overall higher temperature on the UppAnt spot which indicates higher activity on the upper antenna. Conversely for O2, Fig. 15b shows an overall higher LowAnt temperature when the phone is connected to mmWave, as indicating higher activity on the lower antenna. These are expected since the upper antenna is the closest to the BS for O1, while the lower antenna is the closest for O2. In both orientations, the phone started with 4 mmWave channels then throttled to 1 channel and finally a complete handover to LTE. However, the timings of the throttling and handover are different: the throttling to 1 channel occurred at the 95 sec mark for O1 and 85 sec mark for O2, while the LTE handover occurred at the 154 sec mark for O1 and 174 sec mark for O2. Thus, O1 shows a worse overall throughput performance compared to O2. This is due to UppAnt spot heating up on O1 (because the upper antenna is active), and subsequently,

[8]Full video available at `https://youtu.be/LGTQoXCOoDU`

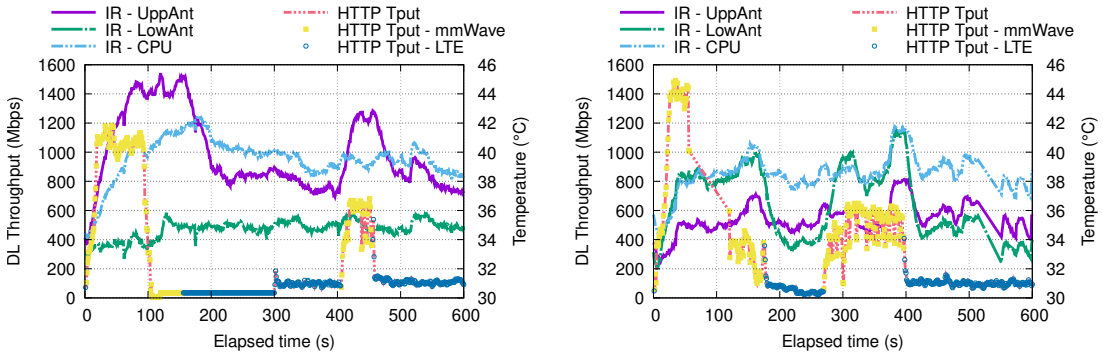

(a) Throughput & temperatures vs. time in orientation O1. (b) Throughput & temperatures vs. time in orientation O2.

Figure 15: Downlink throughput and temperatures vs. time.

the nearby CPU spots heats up at the same time. The heating from two spots introduces a constructive feedback that leads to a faster skin temperature increase on O1. On the other hand, the lower antenna which is active on O2, is far from the CPU spot. Fig. 15b even shows a quick temperature reduction on CPU and LowAnt spots after the phone switched to LTE, thus resulting in a faster switching back to mmWave at the 270 s mark. It should be noted that when the phone is reconnected to 5G mmWave on both O1 and O2, it is only connected to one mmWave channel, thus the lower throughput compared to the beginning of the experiment runs.

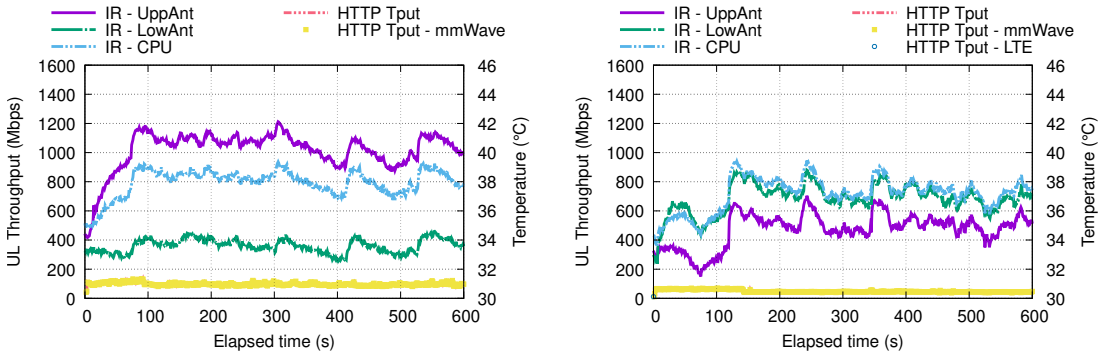

(a) Throughput & temperatures vs. time in orientation O1. (b) Throughput & temperatures vs. time in orientation O2.

Figure 16: Uplink throughput and temperatures vs. time.

Additionally, we performed UL traffic experiments using a similar setup and observed that the phone was connected to 1 mmWave channel during the entire 10-minute run, even as the temperature of the UppAnt spot increased to a maximum of 42° C. Fig. 16 shows the uplink throughput measurement with the corresponding with the temperature values over time. On both orientations, the 5G mmWave connection was never handed over to LTE. In uplink transmission, only a single primary channel is utilized for uplink transmission and available modulation is limited (only up to 64-QAM). Thus, the two factors may lead to a lower CPU/modem load, and consequently, a lower rise in skin temperature, even with the additional power consumption due to transmitting at the UE side. In other words, **transmitting** over 1 mmWave channel caused less rise in skin temperature compared to **receiving** over 4 mmWave channels.

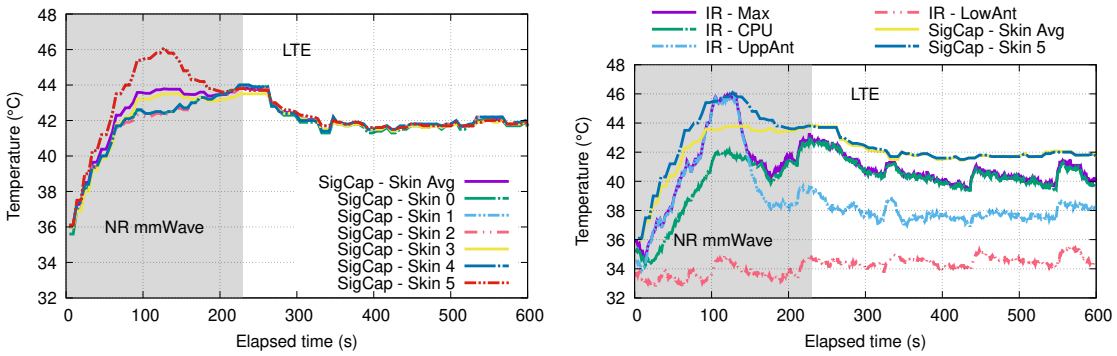

(a) SigCap skin temperature values.

(b) Comparison of IR and SigCap temperature values.

Figure 17: Correlation between IR and SigCap temperature values on S21 Ultra phone with O1 orientation.

7.3. Correlation between SigCap Skin Temperature and IR Spot Temperature

Due to inability of accessing device owner privilege, we could not capture the API's skin temperature value and correlate it to the spot temperatures on the prior experiments[9]. Thus we opted for an S21 Ultra phone similar to S21 Plus, but with device owner privilege availability. We ran an exactly similar experiment with S21 Ultra phone on O1 orientation and initiating downlink traffic. Fig. 17a shows six device skin temperature values captured by SigCap, which possibly correlates to six actual sensors in the device. Furthermore, we correlate the skin temperatures with the IR spot temperatures as shown by Fig. 17b. Specifically, the skin temperature #5 has high correlation to the UppAnt spot. However, as the LTE handover happened at the 230 s mark, it is unlikely that the single skin temperature sensor is triggering the throttle, as the value of skin temperature #5 has spiked before the handover. Rather, the average of all sensors' values may have triggered the handover since it is peaked at 54.4° C at the 230 s mark. This is a clear proof that the location of mmWave antenna and skin temperature sensors contributes to the thermal throttling in 5G mmWave transmission.

8. Conclusions and future work

The methodology developed in this paper, using a variety of apps on smartphones, is a quick, scalable, way of obtaining comprehensive information about complex cellular deployments that use a mix of frequency bands and technologies. At the present time, 5G deployments are evolving rapidly and such measurement campaigns enable researchers to uncover issues that can be further studied on experimental test-beds. It is clear that 5G performance will continue to improve, both in network deployment as well as device performance. Some of the research issues uncovered by the work presented in this paper, which we plan to address in future research, are:

(i) Operator's choice of primary channel is primarily determined by RSRP and RSRQ, as per our measurements. It is not clear however if this choice correlates with higher throughput, thus, we intend to explore learning algorithms to determine if there are better channel choices, given the large number of available channel aggregation options;

(ii) The latency performance of 5G is worse than 4G at present, due to the 5G-NSA deployments. However, even the limited SA data available on T-Mobile's mid-band 5G network does not exhibit improved latency. The latency analysis will be our focus in our future work, including latency under load (FCC ST only measures idle latency).

[9]In our S21 Plus phone, the Qualipoc app limits any elevated privilege. The device owner privilege would have been available on a S21 Plus phone without Qualipoc installed.

(iii) 4G with the LAA and CBRS bands aggregation can deliver throughput in the mid-band that is comparable or even higher than mid-band 5G: Verizon's maximum 4G throughput, using LAA and CBRS is 421 Mbps compared to T-Mobile's 5G throughput of 219 Mbps. However, there will be coexistence issues in LAA and synchronization issues in TDD deployment of CBRS, as both deployments continue to roll out.

(iv) While 5G mmWave has significantly higher data rate compared to 4G+LAA/CBRS, this higher performance cannot be guaranteed in all locations, due to distance limitations, body loss, and non-line-of-sight to the mmWave BS caused by foliage and other obstructions.

(v) One such obstruction is the low-e glass separating outdoor mmWave BS and indoor UE. Our outdoor-to-indoor experiments demonstrate that OtI performance over 5G mmWave is severely limited and is available in only very few locations with unobstructed LoS.

(vi) Another limitation of 5G mmWave is its capability to sustain throughput over a prolonged time. Our IR imaging experiments confirms that prolonged download traffic on 5G mmWave network heats up the UE's mmWave antennas and raises the phone's skin temperature, resulting in throughput throttling. Including our prior work on the impact of device skin temperature to 5G mmWave network [9], we can conclude that the temperature throttling event in 5G mmWave transmission depends on various factors: ambient temperature, CPU/modem usage (as a function of # of aggregated mmWave channels and transmission modulation), location of mmWave antenna, location of skin temperature sensors, and phone orientation *w.r.t.* BS.

Our work underscores the continued need for performing measurements and experiments on deployed networks with consumer devices to understand 5G NR performance in general and mmWave in particular, under real-world conditions and constraints. As theoretical predictions such as by [21] based on channel sounding alone can be overly optimistic. For future applications like AR/VR, having robust coverage for reliable connection is essential. Thus, reducing the variance of 5G mmWave throughput by addressing obstructions and thermal throttling problems will be a focus of our future work. Finally, our current measurement methodology can be scaled up and improved. First, crowd-sourcing can be enabled by combining passive and active measurements into a single app with intuitive design, while also implementing user incentives. We also recommend that these 5G information to be made available using the API for further detailed measurements: PCI, frequency, bandwidth, and beam index for serving and neighboring channels.

References

[1] 3rd Generation Partnership Project, 3GPP Release 15, `https://www.3gpp.org/specifications-technologies/releases/release-15`, accessed: May 2023 (2023).

[2] Statista, 5G core deployments worldwide from 2018 to 2023, by type, `https://www.statista.com/statistics/1330511/5g-core-deployments-worldwide-by-type/`, accessed: June 2023 (2023).

[3] A. Narayanan, E. Ramadan, R. Mehta, X. Hu, Q. Liu, R. A. Fezeu, U. K. Dayalan, S. Verma, P. Ji, T. Li, et al., Lumos5G: Mapping and predicting commercial mmWave 5G throughput, in: Proceedings of the ACM Internet Measurement Conference, 2020, pp. 176–193.

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

[5] M. I. Rochman, V. Sathya, N. Nunez, D. Fernandez, M. Ghosh, A. S. Ibrahim, W. Payne, A comparison study of cellular deployments in Chicago and Miami using apps on smartphones, in: Proceedings of the 15th ACM Workshop on Wireless Network Testbeds, Experimental Evaluation & CHaracterization (WiNTECH'21), 2022, p. 61–68.

[6] V. Sathya, M. I. Rochman, M. Ghosh, Measurement-based coexistence studies of LAA & Wi-Fi deployments in Chicago, IEEE Wireless Communications (2020).

[7] V. Sathya, M. I. Rochman, M. Ghosh, Hidden-nodes in coexisting LAA & Wi-Fi: a measurement study of real deployments, in: 2021 IEEE International Conference on Communications Workshops (ICC Workshops), IEEE, 2021, pp. 1–7.

[8] V. Sathya, M. I. Rochman, T. V. Pasca, M. Ghosh, Impact of hidden node problem in association and data transmission for laa wi-fi coexistence, Computer Communications 195 (2022) 187–206.

[9] M. I. Rochman, D. Fernandez, N. Nunez, V. Sathya, A. S. Ibrahim, M. Ghosh, W. Payne, Impact of device thermal performance on 5G mmWave communication systems, in: 2022 IEEE International Workshop Technical Committee on Communications Quality and Reliability (CQR), IEEE, 2022, pp. 1–6.

[10] Speedtest by Ookla, https://www.speedtest.net/, accessed: May 2023 (2023).
[11] FCC APPs, FCC Speed Test, https://play.google.com/store/apps/details?id=com.samknows.fcc, accessed: May 2023 (2023).
[12] A. Narayanan, X. Zhang, R. Zhu, A. Hassan, S. Jin, X. Zhu, X. Zhang, D. Rybkin, Z. Yang, Z. M. Mao, et al., A variegated look at 5g in the wild: performance, power, and qoe implications, in: Proceedings of the 2021 ACM SIGCOMM 2021 Conference, 2021, pp. 610–625.
[13] A. Narayanan, M. I. Rochman, A. Hassan, B. S. Firmansyah, V. Sathya, M. Ghosh, F. Qian, Z.-L. Zhang, A comparative measurement study of commercial 5G mmWave deployments, in: IEEE INFOCOM 2022-IEEE Conference on Computer Communications, IEEE, 2022, pp. 800–809.
[14] H. Lim, J. Lee, J. Lee, S. D. Sathyanarayana, J. Kim, A. Nguyen, K. T. Kim, Y. Im, M. Chiang, D. Grunwald, et al., An empirical study of 5g: Effect of edge on transport protocol and application performance, IEEE Transactions on Mobile Computing (2023).
[15] A. Hassan, A. Narayanan, A. Zhang, W. Ye, R. Zhu, S. Jin, J. Carpenter, Z. M. Mao, F. Qian, Z.-L. Zhang, Vivisecting mobility management in 5g cellular networks, in: Proceedings of the ACM SIGCOMM 2022 Conference, 2022, pp. 86–100.
[16] M. Giordani, M. Polese, A. Roy, D. Castor, M. Zorzi, A tutorial on beam management for 3GPP NR at mmWave frequencies, IEEE Communications Surveys & Tutorials 21 (1) (2018) 173–196.
[17] Y.-N. R. Li, B. Gao, X. Zhang, K. Huang, Beam management in millimeter-wave communications for 5G and beyond, IEEE Access 8 (2020) 13282–13293.
[18] V. Sathya, L. Zhang, M. Goyal, M. Yavuz, Warehouse deployment: A comparative measurement study of commercial wi-fi and cbrs systems, in: 2023 International Conference on Computing, Networking and Communications (ICNC), IEEE, 2023, pp. 242–248.
[19] M. I. Rochman, V. Sathya, M. Ghosh, B. Payne, M. Yavuz, A measurement study of the impact of adjacent channel interference between C-band and CBRS, arXiv preprint arXiv:2304.07690 (2023).
[20] Platforms for Advanced Wireless Research - PAWR, https://advancedwireless.org/, accessed: May 2023 (2023).
[21] M. Kohli, A. Adhikari, G. Avci, S. Brent, A. Dash, J. Moser, S. Hossain, I. Kadota, C. Garland, S. Mukherjee, et al., Outdoor-to-indoor 28 ghz wireless measurements in manhattan: Path loss, environmental effects, and 90% coverage, arXiv preprint arXiv:2205.09436 (2022).
[22] 3rd Generation Partnership Project, TS 38.331 5G Radio Resource Control (RRC) version 16.1.0 release 16 (2020).
[23] Federal Communications Commission, REPORT TO THE COMMITTEE ON COMMERCE, SCIENCE, AND TRANSPORTATION OF THE SENATE AND THE COMMITTEE ON ENERGY AND COMMERCE OF THE HOUSE OF REPRESENTATIVES, https://www.fcc.gov/sites/default/files/report-congress-usps-broadband-data-collection-feasibility-05242021.pdf, accessed: May 2023 (2021).
[24] M. I. Rochman, SigCap, https://people.cs.uchicago.edu/~muhiqbalcr/sigcap/, accessed: May 2023 (2023).
[25] QTRun Technologies, Network Signal Guru, https://play.google.com/store/apps/details?id=com.qtrun.QuickTest, accessed: May 2023 (2023).
[26] Rohde & Schwarz, QualiPoc Android - The premium handheld troubleshooter, https://www.rohde-schwarz.com/us/products/test-and-measurement/network-data-collection/qualipoc-android_63493-55430.html, accessed: May 2023 (2023).
[27] M. I. Rochman, NR cell map - Hutchinson Field, https://people.cs.uchicago.edu/~muhiqbalcr/grant-park-may-jun-2021/nr-heatmap.html, accessed: May 2023 (2023).
[28] X. Ying, M. M. Buddhikot, S. Roy, SAS-assisted coexistence-aware dynamic channel assignment in CBRS band, IEEE Transactions on Wireless Communications 17 (9) (2018) 6307–6320.
[29] D. Xu, A. Zhou, X. Zhang, G. Wang, X. Liu, C. An, Y. Shi, L. Liu, H. Ma, Understanding operational 5G: A first measurement study on its coverage, performance and energy consumption, in: Proceedings of the Annual conference of the ACM for computer communication, 2020, pp. 479–494.
[30] Derf's test media collection, https://media.xiph.org/video/derf/, accessed: May 2021 (2021).
[31] K. Reitz, Httpbin - a simple HTTP request & response service, https://httpbin.org/, accessed: September 2023 (2023).