

A Comparative Measurement Study of Commercial 5G mmWave Deployments

Arvind Narayanan^{†*}, Muhammad Iqbal Rochman^{‡*}, Ahmad Hassan[†], Bariq S. Firmansyah[§], Vanlin Sathya[¶],
Monisha Ghosh[‡], Feng Qian[†] and Zhi-Li Zhang[†]

[†]University of Minnesota, [‡]University of Chicago, [§]Institut Teknologi Bandung, [¶]Celona, Inc.

Abstract—5G-NR is beginning to be widely deployed in the mmWave frequencies in urban areas in the US and around the world. Due to the directional nature of mmWave signal propagation, improving performance of such deployments heavily relies on beam management and deployment configurations. We perform detailed measurements of mmWave 5G deployments by two major commercial 5G operators in the US in two diverse environments: an open field with a baseball park (BP) and a downtown urban canyon region (DT), using smartphone-based tools that collect detailed measurements across several layers (PHY, MAC and up) such as beam-specific metrics like signal strength, beam switch times, and throughput per beam. Our measurement analysis shows that the parameters of the two deployments differ in a number of aspects: number of beams used, number of channels aggregated, and density of deployments, which reflect on the throughput performance. Our measurement-driven propagation analysis demonstrates that narrower beams experience a lower path-loss exponent than wider beams, which combined with up to eight frequency channels aggregated on up to eight beams can deliver a peak throughput of 1.2 Gbps at distances greater than 100m.

I. INTRODUCTION

5G New Radio (NR) has been specified by 3GPP [1], [2] to operate in low-band (<1 GHz), mid-band (1 to 6 GHz) and high-band or mmWave (>24 GHz). While majority of 5G-NR’s commercial deployments are in the low- and mid-band frequency range, recent deployments in mmWave have also been rapidly increasing, especially in urban areas. 5G NR in the mmWave can harness much higher bandwidths (up to 800 MHz) compared to mid-band and hence offers the potential for greatly increased throughput. However, in order to compensate for the omni-directional and frequency dependent path-loss yet at the same time provide robustness and coverage, mmWave systems use directional beams at both – transmitter (Tx) and receiver (Rx) – ends. While a number of theoretical and experimental studies [3]–[5] have been conducted for the different elements of a 5G mmWave cellular system (e.g., path loss measurements, beam management, etc.), it has not been feasible until recently to verify and understand the performance of commercially deployed 5G mmWave systems. It is important to understand its performance in the real-world, since mmWave propagation can be extremely variable due to a plethora of factors including but not limited to the position

* These authors contributed equally to this paper.

Corresponding authors: arvind@cs.umn.edu, muhiqbalcr@uchicago.edu

This research was in part supported by NSF under Grants CNS-1618836, CNS-2128489, CNS-1814322, CNS-1836772, CNS-1901103, CNS-1915122, CCF-1903880 and a Cisco University Research Grant.

TABLE I: Dataset statistics collected from real world deployments with mmWave 5G coverage at Chicago Loop area.

5G BEAMS Dataset Statistics	
Cumulative distance	73.76 km+ (walk), 69.1 km (drive)
Cumulative time of traces	1260 minutes+
# of commercial operators	2
# of unique 5G PCIs	OpX: 265, OpY: 105
Total area covered	2.3 km ²
mmWave 5G-NR bands	n260/39 GHz and n261/28 GHz

of the body when a smartphone or other user equipment (UE) is held in the hand, orientation with respect to the base station (BS), obstructions due to foliage, vehicles and buildings *etc.* In other words, real-world mmWave systems exhibit a propagation environment that is very different from theoretical models and propagation analysis based on limited channel sounding experiments [6]–[9] which do not take these factors into account. Furthermore, the implementation details of beam-management such as how many beams are used, carrier aggregation across beams *etc.* are left to operators and equipment manufacturers/vendors leading to potential performance differences in deployments. In order to study their effects on end-user performance, conducting in-situ measurements and analysis are called for.

Our objective in this paper is to contribute towards the understanding of real-world 5G mmWave deployment. We performed detailed measurements of two major 5G operators with different 5G mmWave deployment parameters in two representative environments – an urban canyon and an open field – in Chicago, a major U.S. city. We summarize the key contributions as well as the findings of this measurement study.

• **mmWave Deployment Parameters (§V-A).** For both the operators, we reveal the configuration of several key parameters related to deployment and beam management such as 5G-NR bands, channel bandwidth, maximum number of channel aggregation, sub-carrier spacing, PCI assignment, number of Tx beams, *etc.* We also highlight key differences in their strategies and use them to reason our findings.

• **mmWave City-Wide Coverage Analysis (§V-B & §V-C).** We conducted a city-wide outdoor drive-test coverage analysis of mmWave. Our survey finds that within an area of 2.23 km², even with a very dense deployment of mmWave 5G BSs (e.g., 34+ BSs with each BS having 1 to 3 antenna panels), operators can achieve a coverage¹ or no more than 35%. Our

¹Having coverage does not necessarily translate to perceiving good network performance or throughput.

analysis is one of the first to conduct such a city-wide study to reveal the challenge of providing seamless mmWave 5G coverage in a large urban area. We also conducted controlled experiments in select regions within the city to reveal beam coverage maps in diverse environments.

- **Path loss in Different Frequencies & Environments (§VI).**

We demonstrate a signal-strength based method to characterize propagation loss on receive signal power. We show that the path loss exponent (PLE) is relatively higher in NLoS environment than in LoS. We also quantify and show that 28 GHz signals has relatively lower PLE compared to 39 GHz.

- **Beam Selection Analysis (§VII-A).**

We also empirically evaluate the beam selection mechanisms adopted by the operators. In open fields with not much signal reflection, operators deploy mmWave BSs with overlapping footprints to compensate the inability to establish NLoS path. In such areas, operators generally are able to select the best beam as the serving beam. However, in an urban canyon which has a mix of both NLoS and LoS path to BS, operators on an average select beams with degraded quality ($\sim 3.6\text{dBm}$). This is due to the highly dynamic environment coupled with the known sensitivity of mmWave signal propagation.

- **Throughput Performance. (§VIII).** Using the differences in the deployment parameters of the two operators, we proceed to understand its impact on throughput performance. We observe that higher channel aggregation and wider Tx beam contributed to an increase in median throughput. However, the throughput gain is reduced at a distance compared to narrower Tx beam which performs consistently at all distances. We also conducted congestion experiments with multiple UEs connected to the same beam and observe a persistent pattern of uneven throughput performance between the UEs.

- **5G BEAMS: An Open and Real 5G mmWave Dataset (Tables I & II, §IV).** A lack of wide-scale deployment of mmWave 5G, high-cost of tools/license to access low-layer information, need to conduct on-field experiments are just few of the challenges faced by the research community to conduct 5G research on real deployments. We believe the research opportunities provided by 5G BEAMS dataset is beyond the scope of the topics studied in this paper and have therefore released it to the public. The URL of our dataset is:

<https://5gbeams.umn.edu>

The paper is organized as follows: Section II provides a quick primer of beam management in 5G mmWave. Section III discusses related work. Section IV describes the measurement tools and methodology used in this study. Section V provides an overview of the two operator deployments as well as provides insights about their coverage. In Section VI, we present details about our measurement-driven analysis and findings of signal propagation. In Section VII, we conduct empirical analysis of the beam selection strategy observed in the data collected from the two 5G operators and provide key insights. Section VIII presents throughput analyses and results with multiple devices, and conclude in Section IX.

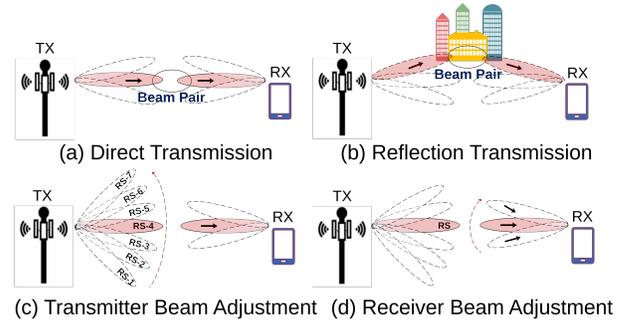


Fig. 1: Beam Management in 5G NR mmWave

II. BACKGROUND ON 5G MMWAVE MANAGEMENT.

5G NR mmWave performance depends on efficient beam selection and beam management, both in stable Line-of-Sight (LoS) and dynamic Non-Line-of-Sight (NLoS) environments. The primary goal of beam-management is to choose the best Tx-Rx beam pair that maximizes the Reference Signal Received Power (RSRP) at the receiver. As depicted in Fig. 1, the best beam pair may not necessarily correspond to Tx and Rx beams that are physically pointing towards each other. In NLoS environments, there can be static obstructions such as buildings, and moving obstructions such as vehicles and pedestrians. Thus, the direct path between the Tx and Rx may be blocked and a reflected path may provide better connectivity (see Fig. 1b); the use of narrow Tx beams increases the likelihood of such occurrences [10], [11].

In general, beam management is divided into three phases. **(1) Initial Beam Establishment.** In this phase a device will acquire broadcast synchronization signals from a BS and proceed to synchronize with the BS using random access for the initial connection. **(2) Beam Adjustment.** This phase is important for maintaining connection to the BS. If the environment changes due to motion, orientation, or vehicles passing by, beam choices might have to be re-evaluated. Beam adjustment happens in two ways: either Tx beam at the BS is varied while Rx beam is fixed (Fig. 1c), or Tx beam is fixed while Rx beam at UE is varied (Fig. 1d). In both, RSRP is measured to determine beam quality. NLoS environments will require adjustments to both Tx and Rx beams more often than LoS environments. **(3) Beam Recovery.** This is required if the current beam pair connection breaks without sufficient time for the regular beam adjustment to adapt. When narrower Tx beams are used, this is more likely due to higher gains [10], [11]. 3GPP specifies certain procedures to handle such beam-failure events [12], [13]: *Beam Failure Detection*, where the UE detects a beam connection failure; *Candidate Beam Identification*, where the UE tries to identify a new beam and connect to it; and *Recovery Request Transmission*, where the UE sends a beam recovery request to the BS. After beam failure, connectivity can be often re-established by means of a new beam pair within the current BS.

Hence, beam adjustment plays a crucial role in mmWave performance. This function depends on multiple factors and deployment parameters, such as operating band, number of

frequency channels, number of Tx and Rx beams, beamwidth, path loss between Tx and Rx, interference management (based on RSRP/RSRQ), and congestion. The choice of different beam configuration parameters will affect the network throughput, latency, and range. In this work, we perform detailed measurements of mmWave in 5G deployments and study their beam management (adjustment).

III. RELATED WORK

We discuss two categories of mmWave research that are relevant to our work: (i) research using theoretical methods and controlled experimental setups, and (ii) measurements conducted using commercial deployments.

A. Theoretical and Controlled Experimental Research

Wireless systems in the mmWave band have been an active area of research for a number of years. Most existing literature discusses the feasibility [14], design [15], and deployment challenges [16]. There are a number of contributions that perform theoretical studies, modeling and simulations on beam management [10]–[12] and beam selection algorithm [13], [17], [18]. Authors in [14] provide a comprehensive overview of emerging 5G mmWave propagation characteristics, including the free-space path loss, material penetration loss, rain and foliage induced attenuation, atmospheric induced attenuation, and other propagation factors. In [12], [13], the authors provide an overview of 5G standardization approaches to beam management procedures for different network architectures (standalone and non-standalone) and signal transmission directions (downlink or uplink). The authors showed that there exist trade-offs between better detection accuracy, improved reactivity and reduced overhead in beam management. However, most of the analysis and theoretical modeling in the literature do not adequately answer all the questions that need to be addressed in real-time deployments, such as: what should the practical inter-distance between two mmWave BSs be, what role does path loss play in beam selection mechanism, what is the trade-off between number of Tx and Rx beams, do more antennas imply higher throughput, do more beams lead to more inter/intra-beam handover or latency?

B. Research on Commercial mmWave Deployments

Since mmWave deployments, especially commercial cellular ones, are fairly recent, the literature on results obtained from studying actual deployments is fairly limited. Authors in [19] explored mmWave usage beyond serving the end-users and demonstrated four novel use cases: 28 GHz as a backhaul point-to-point link, 60 GHz unlicensed access with edge computing, mmWave mesh networks for cost-effective backhauling of small-cell BSs in dense urban scenarios, and automated driving enabled by mmWave-based Vehicular-to-Vehicular (V2V) and Vehicular-to-Everything (V2X). In our previous work [20], we captured the network performance of 5G’s very first commercial mmWave deployments. Further, [21] using commercial 5G services studied the power consumption characteristics as well as application QoE on

smartphones. Finally, in [22], we seek to use user-side factors to characterize and predict the application-level throughput of 5G mmWave transmission at the client device using machine learning techniques. However, due to their inability to access the control plane messages and chipset logs, none of these studies had visibility into the lower-layer information thus were unable to provide insights on the beam management aspects of commercial mmWave-based 5G deployments.

With the access to tools that provide visibility into both the upper and rich lower-layer messages, coupled with a systematically devised measurement methodology involving diverse locations, we collect real-world data of 2 mmWave 5G operators. We analyze mmWave channels and propagation, deployment parameters, beam management, beam selection, beam-to-beam handover, *etc.* and the impact of these on end-user’s network performance.

IV. MEASUREMENT TOOLS AND METHODOLOGY

A. 5G Operators and Locations

In this measurement study, we pick the city of Chicago where two 5G operators: OpX (**Verizon**) and OpY (**AT&T**) have deployed mmWave-based 5G service (in non-standalone or NSA mode) commercially. To better understand the impact of the surrounding environs on beam management and signal propagation, we surveyed the area of Chicago Loop and carefully pick two regions with diverse environmental characteristics (1) BP– the Upper Hutchinson Field **Baseball Park** (*near E Balbo Dr & S Columbus Dr*) representing an open field space, and (2) DT– **DownTown** Chicago (*W Adams Blvd & S LaSalle St to W Jackson Blvd & S State St*) representing an urban canyon surrounded by tall buildings on both sides of the road with high pedestrian and vehicular traffic. More details about the measurement methodology, the data collected as well as an overview about the operator’s deployment is presented later (see §IV-C, §V-A and Table III for details).

B. 5G Smartphones and Measurement Tools

The 5G mmWave network is designed to support ultra-high throughput. For instance, recent studies have shown that commercial mmWave operators can support a downlink throughput of ~ 2 to 3 Gbps [20], [21]. To ensure that the end user’s smartphone device does not become a bottleneck in supporting such high bandwidth, we use $3\times$ state-of-the-art Samsung Galaxy S21 Ultra 5G (S21) smartphones as (SM-G998U1) the user equipment (UE). This model is equipped with the Qualcomm Snapdragon 888 (SM8350) chipset with X60 modem [23] to handle 5G in the low, mid, and mmWave bands. On the mmWave bands, it is capable of receiving up to 8 Tx beams using 2 Rx beams, utilizing up to 8×100 MHz wide channels.

In order to understand the beam management and signal propagation characteristics of commercially deployed mmWave 5G networks, access to *PHY*, *MAC*, *RRC* layer messages (received or sent by the UE) is critical. However, Android APIs do not provide such information. Accessing lower-layer information requires access to Qualcomm *Diag*

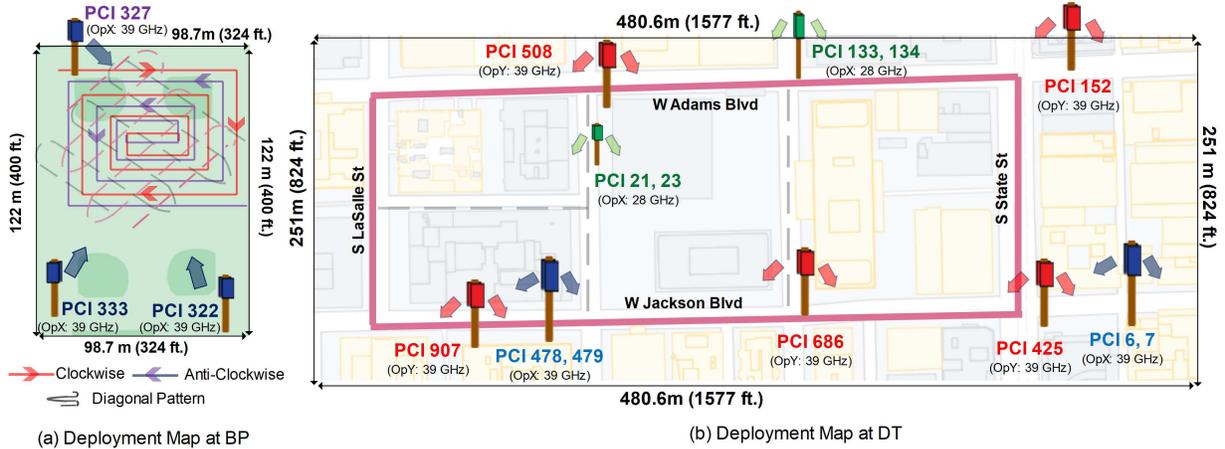


Fig. 2: Map of 5G mmWave BS deployments and walking route/trajectory.

TABLE II: Fields captured in the 5G BEAMS dataset.

Field	Description
Latitude, Longitude	UE's geographic coordinates and estimated accuracy from the Android API
PCell PCI	Primary cell PCI for LTE/NR cell
SCell[x] PCI	Secondary cell PCI for LTE/NR cell [$x = 1 \sim 7$]
RSRP/RSRQ*	Signal strength values for LTE and NR PCell/SCell
Pathloss	Path Loss b/w Tx and Rx for NR PCell/SCell
UL, DL NR-ARFCN	Absolute radio-frequency channel number used in uplink and downlink for NR PCell/SCell
PDSCH/PUSCH Throughput	UL/DL physical throughput for LTE and NR PCell/SCell
Beam SSB Idx*	SSB (Secondary Synchronization Block) Tx/Rx beam index for NR Cell
Best Beam Idx	Tx beam index of dominant beam (highest RSRP) on serving cell
Best Beam State	Status of whether serving beam has the best RSRP over all possible beams (serving + neighbor)
Beam Switch Delay	Delay time when switching between beams on the same or different PCI

* these fields are also captured for neighbor (non-serving) cells and per beam for mmWave NR

(or the diagnostic interface), which needs special licenses and tools. We therefore rely on a professional tool called *Accuver XCAL* which has access to Qualcomm Diag. This tool runs on a laptop and can simultaneously collect the lower (and higher) layer information from up to 4 smartphones concurrently. These smartphones are tethered to the laptop running XCAL using USB cables. Table II provides a summary of a subset of fields captured by the XCAL measurement tool.

C. Data Collection Methodology

We focus our measurement campaign on two regions. The first region is BP, a baseball park with large open fields spanning an area of approximately 17,170m². Fig. 2a depicts this park where OpX has deployed 3 mmWave BSs. Each BS was equipped with 3 directional mmWave transceivers. In order to understand the coverage of OpX within the baseball park area, we constructed two patterns of walking trajectory (see Fig. 2a): (1) a rectangular spiral pattern and (2) a zig-zag pattern which respectively took ~ 27 mins (~ 2.2 km long) and ~ 55 mins (~ 4.6 km) to complete a single route. We repeatedly

walked pattern (1) in a clockwise and anti-clockwise directions for 3 times each, and pattern (2) in two opposing diagonals (*i.e.*, NW \leftrightarrow SE and NE \leftrightarrow SW) for 2 times each. The second region is DT, a section of downtown Chicago region that is surrounded by tall buildings, restaurants, tourist hotspots, *etc.*, with high pedestrian as well as vehicular traffic. Both OpX and OpY have fairly dense 5G mmWave deployments in this area. We pick a 970m walking route in this region that passes through the coverage of both operators. We completed 9 walking loops of the route in an anti-clockwise direction. Fig. 2b depicts the walking route in DT as well as the location of the BSs. These two regions are particularly useful for this study from two perspectives. First, it allows us to compare the two and understand the impact of the environment characteristics on beam management and signal propagation. Second, the DT region allows us to compare the same between the two operators who have different deployment parameters as described in the next section.

In both regions, we use X-CAL to passively collect all the lower-layer information, and run two types of active experiments: (1) *Ping* – measures the round trip latency every second with the target set to a Google DNS server (8.8.8.8), and (2) *HTTP* – download a large YUV data blob over HTTPS [24] (and repeat if the download is complete). For understanding beam management and coverage, the *Ping*-based measurements helped us ensure the 4G and 5G radios always remain in the RRC_CONNECTED state, thus avoiding any fallback to 4G due to data inactivity. *HTTP*-based measurement is used to understand the implications of beam management and configuration over network performance (*e.g.*, downlink throughput). Table I provides a statistical summary of our collected dataset over the full campaign at Chicago. In this work, we particularly only focused on data collected on foot.

V. OVERVIEW OF OPERATORS AND DEPLOYMENTS

A. Deployment Parameters used by OpX and OpY

Table III summarizes several parameters observed in the data collected from our coverage analysis of Chicago city for

TABLE III: mmWave Deployment Parameters (as of June'21).

Parameter	OpX	OpY
Radio Make	Ericsson	Samsung
Radio Model	AIR 5121/6701	HT5H01-60A
# of Antenna Panels	2 to 3 per BS	2 per BS
PCI Assignment	1 per panel	1 per BS
Max. Ch. Agg. (CA)	4 or 8 channels	8 channels
Max. # of Tx Beams	13 per PCI	56 per PCI
5G Deployment Model	NSA	NSA
5G-NR Band	n261, n260	n260
LTE Anchor Band	Band 2, 5, 66	Band 2 & 66
Ch. Width	100 MHz	100 MHz
Sub-Carrier Spacing	120 kHz	120 kHz

both OpX and OpY. The key differences between the two operators include: (i) *PCI assignment*: OpX has a unique PCI for every directional panel (e.g., if a single BS has 3 panels, we observe three unique PCIs) whereas OpY has one per BS; (ii) *number of Tx beam indices*: OpX uses fewer beam indices (13 per PCI or 26 for a BS with 2 panels) compared to OpY (56 per PCI/BS)². This observation suggests OpX uses wider beams than OpY; (iii) *5G-NR band*: OpX uses both 28 GHz and 39 GHz in DT and only 39 GHz in BP while OpY uses only 39 GHz in DT. All BSs of both operators in DT use carrier aggregation (CA) to aggregate a maximum of 8 mmWave channels (1 primary and up to 7 secondary channels), each 100 MHz wide. We also find that depending upon the location (or radio model and/or band), OpX might either aggregate a maximum of 4 or 8 channels. OpY was observed to support up to 8 aggregated channels. With majority of our HTTP-based experiments (that saturated the downlink capacity) focused in the DT and BP regions, we observe that OpX aggregated up to 4 channels in 28 GHz and up to 8 channels in 39 GHz. OpY was observed to support up to 8 aggregated channels.

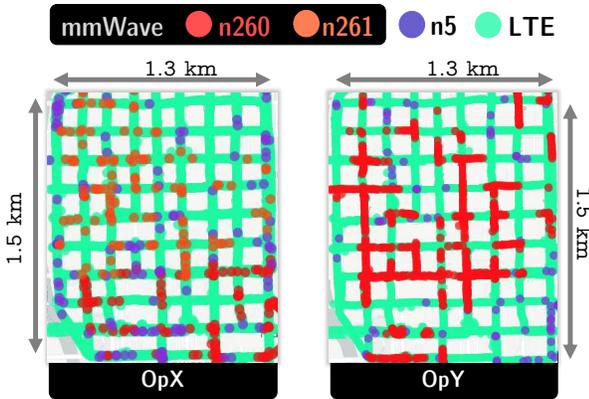


Fig. 3: Outdoor mmWave coverage around Chicago Loop area.
B. City-Wide mmWave Coverage Analysis

In order to understand the outdoor coverage of mmWave, we drove at low-speeds (<25km/hr) through the Chicago Loop

²We observe a maximum beam index of 29 and 63, per PCI for OpX and OpY, respectively. While we observe only a subset of the beam indices in our walking experiments, the unseen beam indices might either be deactivated or can be observed at other regions (e.g., at higher altitude).

area and its surroundings with two UEs (each equipped with a carrier's sim) mounted on the car's windshield. Fig. 3 shows the NR-band mapping. Our survey finds that OpX is able to provide mmWave connectivity to the UE for ~41% of the time during the entire drive session, while OpY provides ~33%. Of course, connectivity to mmWave also depends on how densely the operators have deployed their mmWave base stations as well as the signal propagation characteristics. While driving, to the best of our ability, we visually identify the mmWave BSs deployed by OpX and OpY and record its geographic locations. Our survey finds that within an area of 2.23 km², a minimum of 34 and 19 BSs were deployed (via visual survey) by OpX and OpY, respectively. The BSs were often mounted over poles. Note, each BS further has several antenna panels pointing at different directions. We found anywhere from 1 to 3 panels per BS. Our study exemplifies that even with such a dense deployment of mmWave base stations, operators can achieve no better coverage footprint of 35%, in absolute terms, 24.1 km out of 69.1 km. However, this is purely from connectivity perspective. Signal strength characteristics might widely vary which is described in later sections.

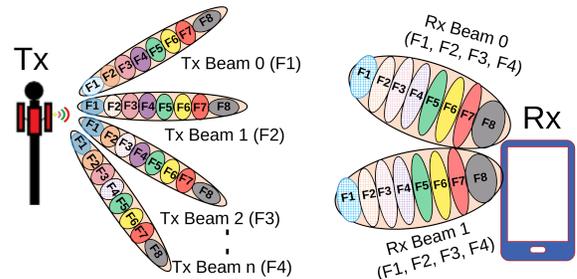


Fig. 4: Carrier Aggregation (CA) in mmWave.

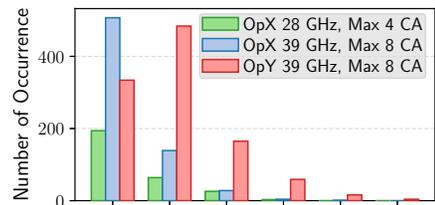
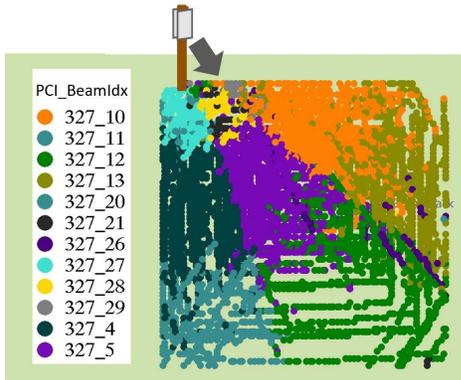
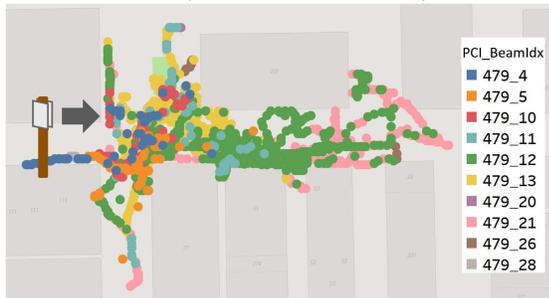


Fig. 5: Number of beams used when CA >= 2.

Carrier aggregation (CA) in frequency can occur over a single beam or multiple beams. As shown in Fig. 4, each Tx beam can transmit on up to eight 100 MHz wide channels. In the deployments we measured, a particular frequency was transmitted only on one Tx beam, i.e., there were no simultaneous data transmissions on the same frequency from multiple beams. On the Rx side, the phone has two beams, each of which can receive over all frequency channels. Thus, this is not a pure implementation of Multiple-Input-Multiple-Output (MIMO) where BS transmit the same signal over multiple frequencies, but an aggregation of different signals over space-frequency with receive combining. OpX and OpY exhibit differences in how they combine CA with beams as shown in Fig. 5. When more than 2 channels are aggregated, OpX uses a single beam most often in both 28 GHz and 39 GHz,



(a) LoS: Serving Beams of PCI 327, OpX at BP.



(b) NLoS+LoS: Serving Beams of PCI 479, OpX at DT.

Fig. 6: Coverage by different serving beams of a PCI (OpX). whereas OpY uses two beams. This could be due to our earlier observation where we found OpX uses wider beams than OpY.

C. Beam Coverage in LoS and NLoS Conditions

Line of Sight (LoS) Coverage. In an open field such as BP, mmWave beams have little to no scope to reflect off surfaces to establish non-line of sight (NLoS) path to the UE. Not surprisingly, OpX has densely deployed 3 mmWave BSs to cover the same geographic spot. For instance, as shown in Fig. 2a, the central region at BP falls in the coverage footprint of three different panels or PCIs (322, 327 and 333). We find that the UE typically gets connected to a PCI with LoS. We therefore use this location to understand and quantify the beam coverage. Fig. 6a shows OpX’s PCI 327’s beam coverage. This figure clearly shows each unique serving beam’s demarcated geographical footprint under LoS conditions. Upon further investigation, we find that each unique beam associated with PCI 327 at BP location covers an area approximately between 1,800-4,000 m^2 .

NLoS Coverage. We conduct similar analysis in the DT area that is surrounded by tall buildings. An urban canyon like this provides plenty of opportunity for the beams to reflect and establish NLoS path to the UE. Fig. 6b shows the coverage of OpX’s PCI 479. Unlike the clear footprints observed under LoS conditions in BP, in NLoS, we find each beam’s coverage can be highly dynamic and sparse. No doubt, coverage with NLoS is largely dictated by two key factors: (i) the mmWave signal reflection characteristics provided by the surrounding environment, and (ii) the UE’s contextual factors such as geolocation, moving speed, direction, *etc.*

To summarize, in this section we presented an overview of the deployment parameters and coverage of real-world mmWave 5G deployments by two carriers. We also presented the beam coverage in both LoS and NLoS conditions. Next, we will dig deeper to understand the path propagation characteristics of mmWave.

VI. PROPAGATION ANALYSIS

Propagation measurements in mmWave have been conducted in multiple environments by a number of researchers, e.g. [7]–[9]. Most of these studies were carried out in a precisely controlled manner using high-fidelity channel sounding equipment that enables not only path loss measurements but also channel impulse responses when wideband channel sounding signals are used. These and many of similar measurements have formed the cornerstone of mmWave system development, including 5G. However, there is a dearth of measurement data on propagation in real-world environments using such as ones collected using commercial off-the-shelf (COTS) hand-held smartphones. Factors such as body-loss, hand obstructions on the receive antenna, foliage and building blockage have been considered in isolation but not in combination with real-world deployments and constraints. In this analysis, we use RSRP values recorded from the S21 smartphones running simultaneous *Ping* workload. Although the RSRP value may not be calibrated between the phones, they all uses the same modem chipset (see §IV-B for details). Thus, RSRP values recorded by the modem are assumed correct within the smartphone model (*i.e.*, the same value will lead to the same behavior for all S21 smartphones).

Primarily imposed by the both the tool and the UE/smartphone, there are two main limitations in the measurements available to us for analyzing propagation: (i) it is unclear as to how the “path loss” measurement obtained from the tool is being computed, since the transmit power could vary with the use of power control. We have observed that the RSRP of the primary channel is always higher than the secondary channels, indicating a higher transmit power. Furthermore, the combined RSRP from the two receive beams on the phone is used to compute the path loss, not the RSRP on each individual receive beam, and (ii) we compute distances based on the GPS coordinates available from the phone, which have an inherent inaccuracy exacerbated by tall buildings in the DT location. With these constraints, it is impossible to “fit” the path-loss measurements to any of the well-known path-loss models [6]. Instead, since the RSRP calculation is well-defined, we focus on the RSRP measurements on a beam-pair level to perform relative comparisons of RSRP using the approach in [9] where a floating intercept model is used:

$$RSRP[dB] = \alpha + 10\beta \log_{10}(d) + X_{\sigma} \quad (1)$$

where d is the distance in meters, α is the intercept in dB, β is the slope, and X_{σ} is a zero mean Gaussian random variable with a standard deviation σ in dB. It should be noted here that β **should not** be considered as the path-loss exponent (PLE) since the intercept α is **not** the reference power at the reference distance of 1 m that is commonly assumed for

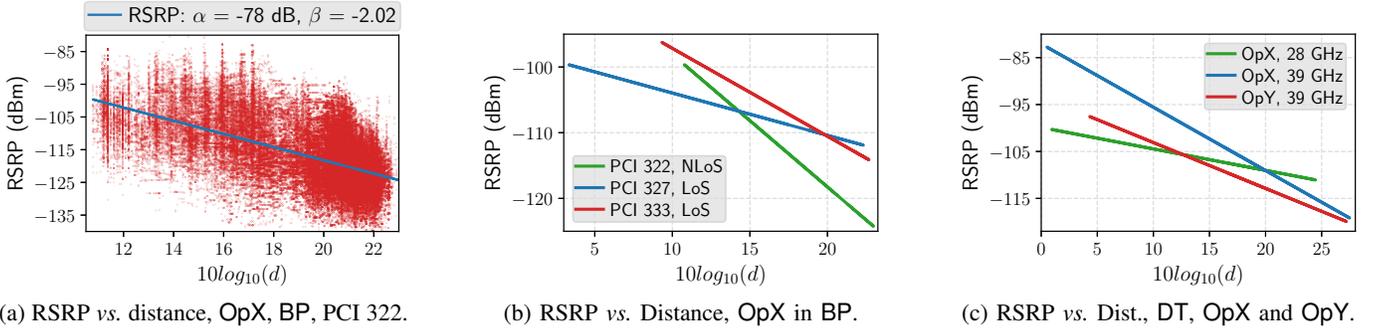


Fig. 7: Line fitting of RSRP vs. distance.

mmWave propagation. Instead, α includes all contributions due to frequency dependence, Tx and Rx antenna gains, clutter, body loss, foliage, *etc.* However, β can be used to make relative comparisons as will be described later. The RSRP analysis in this section is based on observations made over the primary channel from the data collected using the *Ping*-based measurements. We fit the linear model described above to the RSRP for every Tx-Rx pair, where each Rx beam is considered separately.

A. RSRP vs. Distance for OpX in BP

As shown in Fig. 2a OpX has deployed 3 BSs in the BP location, with PCIs 322, 327 and 333 providing coverage footprint to the inside of the baseball field. PCI 322 is partially obstructed by foliage (NLoS) while the other two PCIs are less obstructed. Fig. 7a shows the RSRP vs distance performance of PCI 322 where the scatter plot of all individual Tx-Rx beam pairs (not just the best beam pair) is shown, along with the best linear fit. Fig. 7b shows the best-fit line computed similarly for all three PCIs in that location. As mentioned above, we can use the relative difference in the slopes, β , of these three PCIs in the same area to conclude that the obstructed PCI, PCI 322, has a higher PLE than the other two PCIs in the area.

B. RSRP vs. Distance for OpX and OpY in DT

We combine the RSRPs for each Tx-Rx pair through the entire DT area for OpX in 28 GHz and 39 GHz and OpY in 39 GHz over all deployed PCIs. Fig. 7c shows the performance. OpX at 28 GHz has a lower slope (smaller PLE) compared to 39 GHz, due to the frequency difference, while OpX at 39 GHz exhibits a slightly higher slope (larger PLE) compared to OpY which could be due to the use of wider Tx beams leading to less power received at the same distance. However, overall there is not a significant difference at 39 GHz between the two operators since the deployment environment is basically the same.

VII. BEAM SELECTION ANALYSIS

A. LoS vs. NLoS: Beam Selection

To better understand the impact of environmental features (*e.g.*, open-space vs. urban canyon), Fig 8 compares the RSRP of the serving beam between BP (LoS) and DT (NLoS+LoS). Overall, we find the RSRP at BP (which, except PCI 322,

mainly propagates via LoS to UE) is higher by 3 to 4 dBm when compared to DT which is a mix of LoS and NLoS.

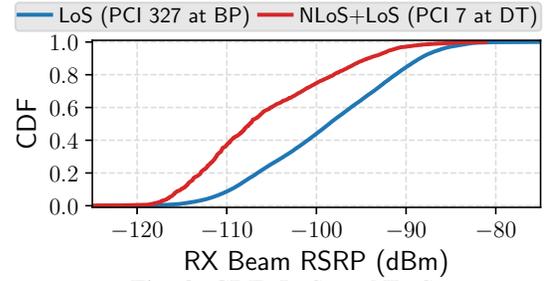


Fig. 8: CDF: LoS vs. NLoS

When selecting the serving beam especially under situations where multiple PCIs (or beams) can cover the same geographic region or when the UE is on the move, operators have to track the UE's location and perform beam switching. To better understand the beam selection strategy used by the operator, we select several metrics (*e.g.*, RSRP, RSRQ, CSI, *etc.*) of the serving beam and compare it against that of the *neighboring beams* (up to 3, which can be from same or different PCI) as seen by the UE. We find that in general, operators use RSRP to make beam selection. We therefore use RSRP for further analysis on evaluating the beam selection strategies deployed by both the operators. We also compare the *Serving Beam's* RSRP with the *Best Beam*³ as reported in the Qualcomm chipset's ML1 Searcher Measurement log messages.

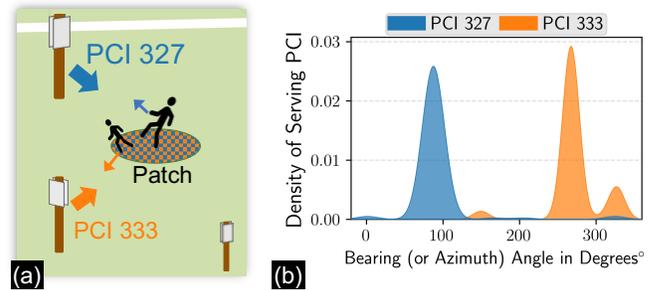


Fig. 9: OpX: UE @ BP prefers PCI (and Beam) with LoS.

As discussed earlier, BP represents an open field (with high density of people during events) providing less opportunity

³Details on how Qualcomm decides which beam is the best is not fully known. Our correlation analysis suggests this to be chosen from the beam with the highest instantaneous RSRP measured by the chipset.

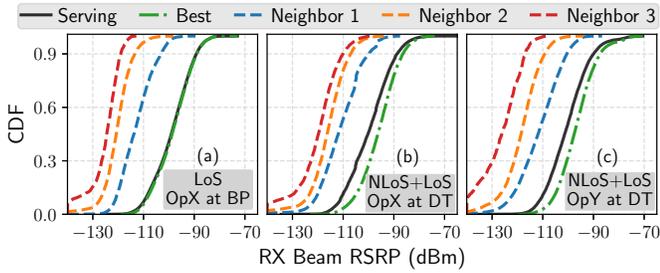


Fig. 10: CDF of OpX RSRP: serving beam vs. best beam vs. top $\{1, 2, 3\}$ neighbor.

TABLE IV: Beam switching statistics at DT.

Operator	Intra-BS Beam Switch			Inter-BS Beam Switch		
	Switch. interval	Delay (s)		Switch. interval	Delay (s)	
		μ	σ		μ	σ
OpX-28 GHz	6.99 s	0.16	0.09	43.39 s	4.72	2.35
OpX-39 GHz	70.18 s	0.35	1.66	N/A	N/A	N/A
OpY-39 GHz	1.29 s	0.2	1.12	85.03 s	24.03	73.14

for establishing NLoS paths between the BS and UE. Not surprisingly, in BP we find OpX has deployed multiple PCIs with overlapping coverage footprint and depending upon the UE’s moving direction, the UE gets connected to the PCI with LoS. For instance, the patch illustrated in Fig. 9a falls under the footprint of all the three PCIs. The bearing (or azimuth) angles which represents the direction of UE’s mobility shows distinct density distributions when connected to PCI 327 versus PCI 333. In terms of how well an operator performs in selecting the best beam as the serving beam, Fig. 10a shows that in BP (or under LoS conditions), the *Best* beam clearly match the *Serving* beam. On an average, the selected beam is also 14.4 dBm higher than the *Neighbor Top 1* beam. This suggests, operators show the ability to in general select the best beam under LoS. Note, selecting the best beam does not always result in better coverage especially in open space settings with limited to no scope of signal reflection.

On the other hand, Fig. 10b (and Fig. 10c) show the RSRP at the DT (*i.e.*, LoS + NLoS environment), on an average there is a degradation of 3.6 dB of the *Serving* beam’s RSRP when compared to that of the *Best* beam. Nonetheless, our study highlights and quantifies the challenges faced by operators which could have several implications on network and application performance. Clearly, differences in the environmental features has an impact of signal reflection and propagation. Such impact is known but challenging to quantify especially in-the-wild. we believe our initial analysis on beam selection as well as the dataset will establish a baseline for the research community as well as to track longitudinal insights.

B. Intra-BS and Inter-BS Beam Switching Statistics

We perform an analysis on OpX and OpY data collected at DT using the *Http* workload which capable of activating more beams. Table IV shows the statistics of beam switching within different beams in the same BS (intra-BS) and beam switching between different BSs (inter-BS) at DT. The phones are side-by-side during the measurement walk, thus enabling comparison of beam switching data from different phones.

OpY uses a larger number of narrower beams compared to OpX and hence we see that the switching interval is smaller, *i.e.* beams are switching more often within the BS. Since DT is a mix of NLoS and LoS, this is expected. The delay incurred by switching beams are varying between all operators. However, given that OpY has less switching interval and considerably low delay, it may demonstrated a more efficient beam management algorithm. For OpX in 39 GHz, we observe high beam switching interval and delay compared to 28 GHz, indicating that the 39 GHz BS holds on to a beam more due to its high delay. Over all operators, the 39 GHz beams shows a higher intra-BS beam switching delay compared to 28 GHz.

The deployment map in Fig. 2b shows that OpX’s 28 GHz BSs are deployed closer together than OpY’s, which results in a lower inter-BS switch interval for OpX compared to OpY. Meanwhile, OpX’s 39 GHz BSs are deployed 340 m between each other, thus the UE never switch between them. It is interesting that the inter-BS switch delay for OpY is much higher: it may due to a less optimized switching algorithm.

VIII. PERFORMANCE ANALYSIS

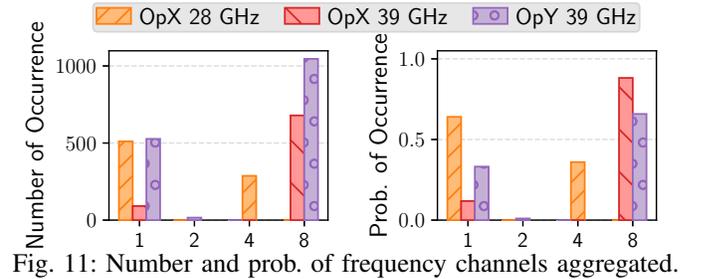


Fig. 11: Number and prob. of frequency channels aggregated.

A. Throughput Comparison at DT

Fig. 12 shows the throughput CDFs of OpX and OpY at 39 GHz and OpX at 28 GHz in DT for **all** data and separated by $d < 100m$ and $d > 100m$, with d as the distance between BS and phone. There are a number of interesting observations we can draw from these results:

1) OpX’s **maximum throughput over all data** is lower than OpY’s in both bands. This is due to 2 reasons: (i) Fig. 11 shows that OpX aggregated a maximum of 4 channels at 28 GHz compared to 8 by both operators at 39 GHz, and (ii) we see from the deployment map in Fig. 2b that OpX’s 39 GHz BSs are much farther apart than OpY’s. In fact, it is rather curious that OpX has deployed their 28 GHz BSs closer together than the 39 GHz BSs: given the theoretical difference in path loss at these 2 frequencies (corroborated in Fig. 7c), it should have been the other way around.

2) OpX’s **median throughput over all data** at 39 GHz is higher than OpY’s. Fig. 11 show that this is due to OpX at 39 GHz aggregates 8 channels at 88% of the time compared to OpY at 62%. Also, OpX’s median throughput is dominated by the data for $d < 100m$, *i.e.*, the increased carrier aggregation outweighs the reduced RSRP due to distance.

3) By comparing OpY’s throughput between $d < 100m$ and $d > 100m$, we see that the throughput distribution is relatively

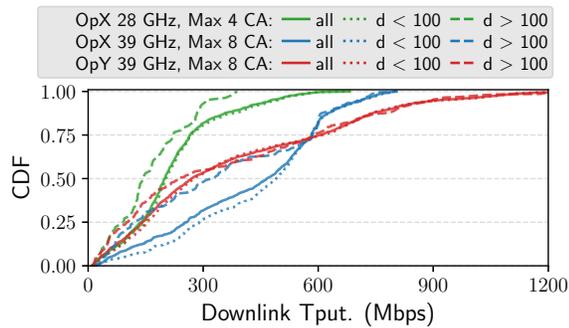


Fig. 12: Throughput comparison for different distances in DT.

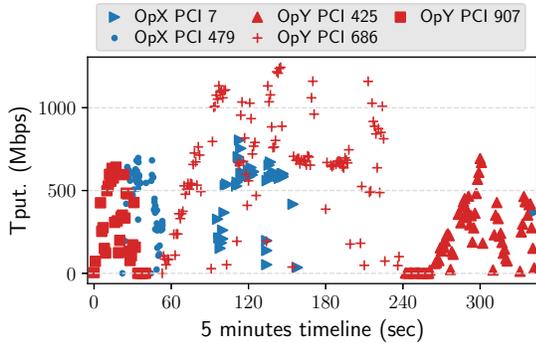


Fig. 13: Throughput vs. time at DT.

unaffected by distance. This is due to the average distance between their BSs being less than OpX at 39 GHz.

4) Fig. 13 shows throughput from a 5 minutes section of the measurements, which starts from LaSalle & Jackson and ends at State & Jackson. Clearly, the denser deployment of OpY (3 BSs) leads to more uniform and higher throughput over that region, with the best throughput achieved by OpY’s PCI 686 due to the LoS environment surrounding the BS.

B. Congestion Experiment at BP

We exploit the stable LoS condition of BP for a multi-UE congestion experiment over static PCIs and beams. Three S21 smartphones are used (UE 1, 2, & 3) to initiate an iPerf3 [25] session to public cloud servers in Illinois and Minnesota. All sessions establish 8 TCP connections in parallel to saturate the downlink capacity. All UEs are handheld stationary at a distance of 50m from BS. Each phone started iPerf3 session one after the other, each with an interval of 1 minute. We performed 3 tests while varying the start order of UE’s transmission: UE 1→UE 2→UE 3, UE 2→UE 3→UE 1, and UE 3→UE 1→UE 2, *i.e.*, a total of 9 runs.

Over the repetitions, we observe a similar trend: The first UE to start the transmission will have its throughput dropped lower than other UEs after a period of time. We pick a representative session and show the results in in Fig. 14. This figure shows the total DL throughput over all channels, the number of CA, primary channel RSRP, and linear sum of RSRP over all channels (summed in the *milliwatt* domain). We omitted RSRQ values, as there is no significant change in primary and secondary channels’ RSRQ. All UEs used the maximum number of RB with modulations of 64-QAM, 16-QAM, and QPSK used 86%, 13%, and 1% of all time,

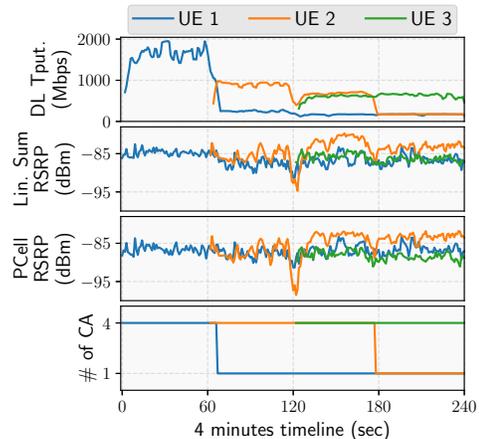


Fig. 14: Congestion Experiments: throughput, linear sum RSRP, primary RSRP, # of CA over time.

respectively. UE 1 starts with the highest throughput (~ 1.6 Gbps) with 4 CA. As the second UE becomes active at the 1-minute mark, UE 1’s throughput drops to 260 Mbps (1 CA). In the RSRP domain, there is no change in primary RSRP. However, we observe a reduction in the linear sum RSRP, indicating a lower overall channel condition. For the next minute (*i.e.*, 120s to 180s), UE 2 started with an average throughput of 850 Mbps followed with small decrease as UE 3 becomes active. UE 2’s throughput further decreased to 165 Mbps at the 3 minute mark due to reducing to 1 CA.

However, the black-box nature of real-deployment measurement still prevents us from further exploration since we lack detailed BS information. With limited visibility only from the UE’s perspective, we can only attribute the throughput degradation to BS not serving more channel. However, our research reveals opportunity lie in understanding how operators/vendors deploy mechanisms related to deciding whether to use channel aggregation and at what level, and evaluate its effectiveness.

IX. CONCLUSIONS AND FUTURE WORK

We presented what we believe is the first measurement study of real-world 5G NR mmWave cellular networks with emphasis on understanding mmWave beam management. Using state-of-the-art measurement tools that collect both application layer as well as rich lower-layer information, and by carefully designing systematic measurement methodologies, we conduct field experiments in Chicago city involving two major 5G operators. We reveal the different deployment parameters used by the operators, provide interesting insights on the mmWave coverage as well as conduct a comparative study to reveal the relative importance of the configuration parameters on beam management, signal propagation characteristics and network performance. We hope that these insights from in-situ measurements will benefit the broader researcher community’s efforts in understanding and improving mmWave performance in diverse deployment scenarios, far beyond the scope of the topics studied in this paper. The authors have provided public access to their data at <https://5gbeams.umn.edu>.

REFERENCES

- [1] 3rd Generation Partnership Project, “3GPP Release-15,” 2016, <http://www.3gpp.org/release-15/>.
- [2] —, “3GPP Release-16,” 2021, <http://www.3gpp.org/release-16/>.
- [3] S. Rangan, T. S. Rappaport, and E. Erkip, “Millimeter-wave cellular wireless networks: Potentials and challenges,” *Proceedings of the IEEE*, vol. 102, no. 3, pp. 366–385, 2014.
- [4] R. W. Heath, N. González-Prelcic, S. Rangan, W. Roh, and A. M. Sayeed, “An Overview of Signal Processing Techniques for Millimeter Wave MIMO Systems,” *IEEE Journal of Selected Topics in Signal Processing*, 2016.
- [5] M. Shafi, A. F. Molisch, P. J. Smith, T. Haustein, P. Zhu, P. De Silva, F. Tufvesson, A. Benjebbour, and G. Wunder, “5G: A tutorial overview of standards, trials, challenges, deployment, and practice,” *IEEE journal on selected areas in communications*, vol. 35, no. 6, pp. 1201–1221, 2017.
- [6] T. S. Rappaport, Y. Xing, G. R. MacCartney, A. F. Molisch, E. Mellios, and J. Zhang, “Overview of millimeter wave communications for Fifth-Generation (5G) Wireless Networks—with a focus on propagation models,” *IEEE Transactions on Antennas and Propagation*, vol. 65, no. 12, pp. 6213–6230, 2017.
- [7] C. Gentile, P. B. Papazian, N. Golmie, K. A. Remley, P. Vouras, J. Senic, J. Wang, D. Caudill, C. Lai, R. Sun, and J. Chuang, “Millimeter-wave channel measurement and modeling: A NIST perspective,” *IEEE Communications Magazine*, vol. 56, no. 12, pp. 30–37, 2018.
- [8] S. Sun and T. S. Rappaport, “Wideband mmWave channels: Implications for design and implementation of adaptive beam antennas,” in *2014 IEEE MTT-S International Microwave Symposium (IMS2014)*, 2014.
- [9] T. S. Rappaport, G. R. MacCartney, M. K. Samimi, and S. Sun, “Wideband Millimeter-Wave Propagation Measurements and Channel Models for Future Wireless Communication System Design,” *IEEE Transactions on Communications*, vol. 63, no. 9, pp. 3029–3056, 2015.
- [10] M. N. Akbar, S. Atique, M. Saquib, and M. Ali, “Capacity enhancement of indoor 5g mmwave communication by beam steering and narrowing,” in *2018 10th International Conference on Electrical and Computer Engineering (ICECE)*. IEEE, 2018, pp. 85–88.
- [11] W. Attaoui, K. Bouraqia, E. Sabir, M. Benjillali, and R. Elazouzi, “Beam alignment game for self-organized mmwave-empowered 5g initial access,” in *2019 15th International Wireless Communications & Mobile Computing Conference (IWCMC)*. IEEE, 2019, pp. 2050–2057.
- [12] M. Giordani, M. Polese, A. Roy, D. Castor, and M. Zorzi, “A tutorial on beam management for 3GPP NR at mmWave frequencies,” *IEEE Communications Surveys & Tutorials*, vol. 21, no. 1, pp. 173–196, 2018.
- [13] Y.-N. R. Li, B. Gao, X. Zhang, and K. Huang, “Beam management in millimeter-wave communications for 5G and beyond,” *IEEE Access*, vol. 8, 2020.
- [14] A. N. Uwaechia and N. M. Mahyuddin, “A comprehensive survey on millimeter wave communications for Fifth-Generation Wireless Networks: Feasibility and challenges,” *IEEE Access*, vol. 8, 2020.
- [15] I. K. Jain, R. Subbaraman, and D. Bharadia, “Two Beams Are Better than One: Towards Reliable and High Throughput MmWave Links,” in *ACM SIGCOMM 2021 Conference*, 2021, p. 488–502.
- [16] Y. Niu, Y. Li, D. Jin, L. Su, and A. V. Vasilakos, “A survey of millimeter wave communications (mmWave) for 5G: opportunities and challenges,” *Wireless networks*, vol. 21, no. 8, pp. 2657–2676, 2015.
- [17] Y. Wang, A. Klautau, M. Ribero, A. C. Soong, and R. W. Heath, “Mmwave vehicular beam selection with situational awareness using machine learning,” *IEEE Access*, vol. 7, pp. 87 479–87 493, 2019.
- [18] F. Fernandes, C. Rom, J. Harrebek, and G. Berardinelli, “Beam management in mmwave 5g nr: an intra-cell mobility study,” in *2021 IEEE 93rd Vehicular Technology Conference (VTC2021-Spring)*. IEEE, 2021, pp. 1–7.
- [19] K. Sakaguchi, T. Haustein, S. Barbarossa, E. C. Strinati, A. Clemente, G. Destino, A. Pärssinen, I. Kim, H. Chung, J. Kim *et al.*, “Where, when, and how mmWave is used in 5G and beyond,” *IEICE Transactions on Electronics*, vol. 100, no. 10, pp. 790–808, 2017.
- [20] A. Narayanan, E. Ramadan, J. Carpenter, Q. Liu, Y. Liu, F. Qian, and Z.-L. Zhang, “A First Look at Commercial 5G Performance on Smartphones,” in *The Web Conf. (WWW '20)*, 2020.
- [21] A. Narayanan, X. Zhang, R. Zhu, A. Hassan, S. Jin, X. Zhu, X. Zhang, D. Rybkin, Z. Yang, Z. M. Mao, F. Qian, and Z.-L. Zhang, “A Variegated Look at 5G in the Wild: Performance, Power, and QoE Implications,” in *ACM SIGCOMM 2021 Conference*, 2021.
- [22] A. Narayanan, E. Ramadan, R. Mehta, X. Hu, Q. Liu, R. A. K. Fezeu, U. K. Dayalan, S. Verma, P. Ji, T. Li, F. Qian, and Z.-L. Zhang, “Lumos5G: Mapping and Predicting Commercial mmWave 5G Throughput,” in *ACM IMC '20*, 2020.
- [23] (2021) Snapdragon X60 5G modem-RF system. [Online]. Available: <https://www.qualcomm.com/products/snapdragon-x60-5g-modem>
- [24] (2021) Derf’s test media collection. [Online]. Available: <https://media.xiph.org/video/derf>
- [25] (2021) iPerf3 - the TCP, UDP and SCTP network bandwidth measurement tool. [Online]. Available: <https://iperf.fr/>