Beam Management in Operational 5G mmWave Networks

Yufei Feng, Jin Wei, Phuc Dinh, Moinak Ghoshal, Dimitrios Koutsonikolas
Institute for the Wireless Internet of Things, Northeastern University, USA

ABSTRACT

Due to the directional nature of mmWave signal propagation, beam management plays a critical role in the performance of 5G mmWave deployments. However, the details of beam management in commercial deployments and its performance in real-world scenarios remain largely unknown. In this paper, we fill this gap by performing a comparative measurement study of the beam management procedure of two major US operators in Boston, MA. We study a number of beamforming parameters including beamwidth, number of beams, beam switching delay, and their impact on performance, and we explore the interplay between beam management and rate adaptation. We also investigate for first time Rx beam management on the UE side. Finally, we study the beam tracking performance and the quality of the selected beams for the two operators.

CCS CONCEPTS

• **Networks** → **Mobile networks**; **Network measurement**; **Network performance analysis**;

ACM Reference Format:
Yufei Feng, Jin Wei, Phuc Dinh, Moinak Ghoshal, Dimitrios Koutsonikolas. 2023. Beam Management in Operational 5G mmWave Networks. In *The 7th ACM Workshop on Millimeter-Wave and Terahertz Networks and Sensing Systems (mmNets '23), October 6, 2023, Madrid, Spain.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3615360.3625092

1 INTRODUCTION

5G mmWave is being rapidly deployed by major mobile operators, especially in urban settings. To cope with the high propagation loss in the mmWave frequency bands, mmWave systems use directional beams on both the transmitter (Tx) and the receiver (Rx) end. However, directionality introduces new challenges – vulnerability to blockage and beam misalignment due to mobility. Consequently, *beam management* plays a crucial role in the performance of 5G mmWave deployments.

5G mmWave beam management has been extensively studied in the literature, e.g., [5–8, 10, 11]. However, all these works are theoretical, and their evaluation methodologies rely on simulations. In practice, beam management depends on a large number of factors, including but not limited to operating band, number of Tx/Rx beams, beamwidth, path loss between Tx and Rx, and interference management. Additionally, unlike WiFi networks, cellular networks are "black boxes" from the user's point of view; users have no direct insight into the operations performed on either the BS or UE side. Further, the implementation details of beam management, e.g., how many beams are used, when to trigger beam adaptation, etc. are left to operators and equipment vendors leading to potential performance differences in different deployments. Hence, understanding the details of beam management and its impact on performance in commercial 5G mmWave deployments is important, as it can provide valuable insights into the performance of 5G mmWave networks, and enable more realistic simulation studies.

Nonetheless, the details of beam management in commercial deployments and its performance in real-world scenarios remain largely unknown. The only experimental study of beam management to our best knowledge is the work by Narayanan *et al.* [9]. In that work, the authors studied beam management of two major US operators in Chicago, IL. They found that the two operators use different beam management parameters (number of beams, beamwidth, beam switch time, etc.), and these differences result in different signal propagation characteristics and coverage for the two operators.

In this paper, we extend the work in [9] by studying the beam management procedures of the two major US 5G mmWave operators (Verizon and AT&T) in Boston, MA. Our study reveals that the two operators in Boston use very different beam management parameters compared to those reported by Narayanan *et al.* in Chicago, suggesting that extensive measurement studies in different cities are necessary to fully understand the details of beam management in commercial deployments. In contrast to the work in [9],

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
mmNets '23, October 6, 2023, Madrid, Spain
© 2023 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.
ACM ISBN 979-8-4007-0338-6/23/10…$15.00
https://doi.org/10.1145/3615360.3625092

which focused on coverage, our work focuses on the impact of beam management on performance and we explore the interplay between beam management and rate adaptation – the primary mechanism employed by all types of wireless networks to adapt to channel fluctuations. Additionally, we investigate for first time Rx beam management on the UE side. Finally, we study the beam tracking performance for the two operators and the quality of the selected beams.

2 BACKGROUND ON 5G MMWAVE BEAM MANAGEMENT

5G mmWave beam management, defined as a set of L1/L2 procedures, is designed to establish and maintain optimal transmission conditions between the Tx and Rx beams [3]. Although beam management applies to both uplink and downlink transmissions, we focus here on components essential to downlink beamforming. These components, executed at the 5G mmWave BS (gNB) or the UE, include:

Beam Sweeping: In 5G NR, the gNB periodically transmits Synchronization Signal Blocks (SSBs), typically every 20 ms (SSB periodicity) in multiple directions. Each SSB, broadcast in a specific direction, is tagged with a unique SSB index. Throughout the paper, the terms SSB Index and Beam Index are used interchangeably. On the Rx end, the UE uses the transmitted SSBs for initial synchronization with the cell and for initial beam establishment.

Beam Measurements and Beam Determination: In downlink communication, beam measurements are carried out at the UE. In the idle state, UEs conduct signal strength measurements based on the periodic SSBs transmitted from the gNBs. In the connected state, the UE measurements rely on SSBs for beam switching and Channel State Information-Reference Signals (CSI-RS) for beam refinement.[1] The UE uses these measurements to identify the best Tx SSB index and also keeps track of a set of N candidate Tx beams ranked according to their received signal strengths. For each of these SSB indices, the UE also selects one or more optimal Rx beams. The number of Rx beams depends is device-specific. For instance, the SM-G990U1 Galaxy S21 phone we use in our experiments uses the Qualcomm Snapdragon 888 SoC with X60 LTE/5G module integrated that can support two spatial beams [2].

Beam Reporting: The UE reports its signal strength measurements of the candidate beams to the gNB to enable multiple beam management procedures, including beam selection, beam tracking, and beam refinement. The reporting is triggered when certain conditions on the serving beam signal strength are met [1]. This enables the gNB to maintain updated signal strength information of different SSB Indices under channel dynamics for beam switching.

5G mmWave beam management employs two types of beams – SSB beams and CSI-RS beams. SSB beams are wider and play a vital role in initial beam establishment and beam tracking under mobility. On the other hand, CSI-RS beams, being narrower and more directional, are employed in the beam refinement phase following the establishment of SSB beams. In this work, we focus on SSB beams and leave the study of the beam refinement phase as future work.

3 METHODOLOGY

5G operators. We performed all the measurements in downtown Boston using Verizon and AT&T's 5G mmWave services. In Boston, Verizon's 5G mmWave service works in the 28 GHz frequency band (n261) using 48 SSB indices (Tx beams) per BS. On the other hand, AT&T works in the 39 GHz frequency band (n260) and uses 24 Tx beams per BS. This observation suggests that Verizon uses narrower beams compared to AT&T. Note that these numbers are very different from the numbers reported in [9] for Chicago (13 for Verizon, 56 for AT&T) suggesting that the same operator uses different hardware and/or beam management strategies in different cities. The BSes for both operators are mounted on top of traffic lights or lamp posts.

The 5G mmWave service of both operators utilizes carrier aggregation (CA), allowing the UE to simultaneously use up to 8 cells or component carriers (CCs) in the downlink direction and up to 2 cells in the uplink direction. Beam management is performed separately for each cell. For simplicity, we focus on the performance and beam management of the primary cell (PCell), which is used for initial access and control signaling along with data transmission. Under ideal conditions, the MAC throughput of each cell is approximately equal to $1/n$ of the total MAC throughput, where n is the total number of cells used for data transmission.

5G devices and cloud servers. We used a Samsung S21 phone as the UE, which supports the 5G mmWave bands n260/261 and 8-CC (8x100 MHz)/2-CC downlink/uplink CA achieving up to 3.5 Gbps and 350 Mbps in the downlink and uplink directions, respectively. We used an Accuver XCAL-Solo [4] device to collect PHY-layer KPIs. Finally, we used a Google cloud server located in Washington, DC for all the throughput measurements. The server's ingress/egress network bandwidth was 16 Gbps+, ensuring that the server's network capacity does not become the bottleneck.

Experiments. We used iperf3 to generate backlogged UDP downlink traffic and logged the total and cell-wise MAC downlink throughput, MCS, RSRP, RSRQ, SSB indices, and RX beam indices with XCAL. The throughput and MCS were

---

[1] In uplink communication, beam measurements are performed at the gNB and sounding reference signals (SRS) are used instead of CSI-RS.

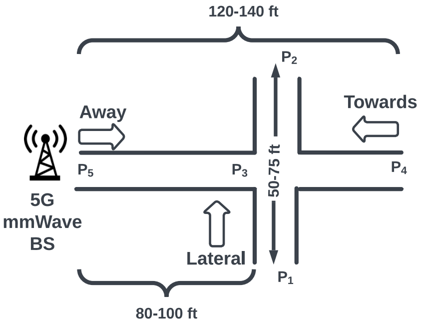

Figure 1: Trajectories for mobility experiments.

logged at a granularity of 100 ms. For the beam and signal strength related metrics, although their minimum logging granularity on XCAL is 100 ms, we observed that the logging interval in practice varies from 200 ms to a few s. In particular, when the UE was using the AT&T service, XCAL was able to log Rx beam information only once every few s.

To study beam management under mobility and self-blockage, we walked towards, away from, and laterally to a 5G mmWave BS (Fig. 1) at the typical walking speed (~3 ft/s), and took measurements along the way for 15-40 s. Given the deployment of 5G mmWave BSes on traffic lights and lamp posts, we believe these three mobility patterns represent realistic user mobility scenarios. In the case of walking towards (away from) the BS, we started at P4 (P5) and stopped at P5 (P4). For lateral motion, we started at P1, walked until reaching the BS at P3, and continued walking until P2. While performing the measurements, we sometimes noticed handovers to a different 5G mmWave BS. We carefully removed such traces and report results from cases where there were no handovers. We collected 10 traces for each trajectory-operator combination.

4 RESULTS

Figs. 2 and 3 show representative timelines of the PCell MAC layer throughput, PCell MCS, SSB indices, and Rx beams with Verizon and AT&T, respectively, for the three mobility scenarios. In all three scenarios, throughput is affected by mobility and exhibits fluctuations over time. Notably, the backward trajectory (Figs. 2b, 3b), impacted by a combination of mobility and self-blockage, yields the lowest MCS and average MAC throughput for both operators. In all three mobility scenarios, we observe that, as the UE moves, both the Tx and Rx beams change over time as the BS and UE try to maintain beam alignment. Nonetheless, there are fundamental differences in the beam alignment procedure between the two operators, which we discuss below.

4.1 Impact of beamwidth

Illustrating example. Figs. 2 and 3 clearly demonstrate the impact of different beamwidths employed by the two operators on beam management. Since Verizon uses narrower beams than AT&T (§3), it triggers beam adaptation more often in all three scenarios. As an example, in the lateral trajectory, the Verizon BS (Fig. 2c) keeps tracking the user motion as the user moves from P1 to P2 and switches to a new beam every few seconds, using a total of 7 beams (SSB indices); in contrast, the AT&T BS (Fig. 3c) uses only two beams for the whole trajectory and performs a single beam switching roughly in the middle of the trajectory, when the user's body starts incurring self-blockage as the user passes in front of the BS and continues walking towards P2.

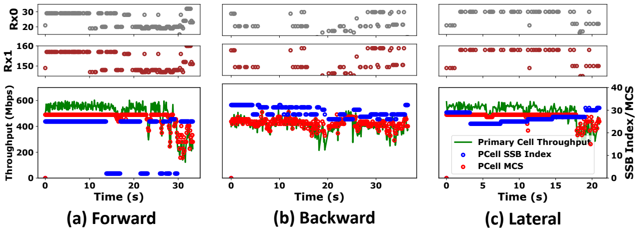

(a) Forward (b) Backward (c) Lateral

Figure 2: Verizon: PCell MAC throughput, MCS, SSB index, and Rx beams over time for different trajectories.

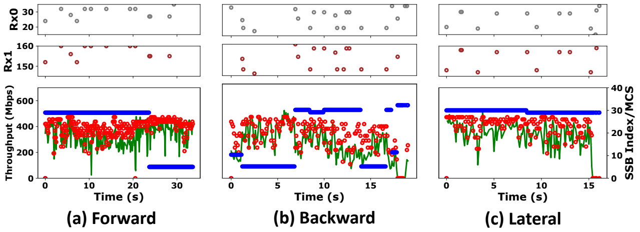

(a) Forward (b) Backward (c) Lateral

Figure 3: AT&T: PCell MAC throughput, MCS, SSB index, and Rx beams over time for different trajectories.

The impact of different beamwidths is further highlighted when observing the end of the lateral trajectory and the forward trajectory, where the two operators experience different degrees of throughput degradation. For the lateral case (Figs. 2c, 3c), the user incurs severe self-blockage towards the end of the trajectory, leading to a more severe throughput reduction for Verizon compared to AT&T due to the use of narrow beams. In the case of forward trajectory, one would expect the BS to maintain the same beam for the whole duration of the experiment, as the UE always faces the BS. Nonetheless, we observe that, as the user approaches the BS, the Verizon BS keeps switching beams starting at about 15 s (Fig. 2a) trying to maintain beam alignment on the vertical plane with the UE. Further, when the user approaches the BS, none of Verizon's narrow beams can be steered steeply downwards to serve the UE; as a result, throughput exhibits a drop of about 200 Mbps in the last 5 s. In contrast, AT&T's wider beams can maintain coverage requiring only one beam switch at 23 s and sustain a higher throughput at the end of the trajectory (Fig. 3a).

Overall statistics. Figs. 4a (left) and 4b (left) show the average number of unique beams used by each operator and the average number of beam changes, respectively, over all

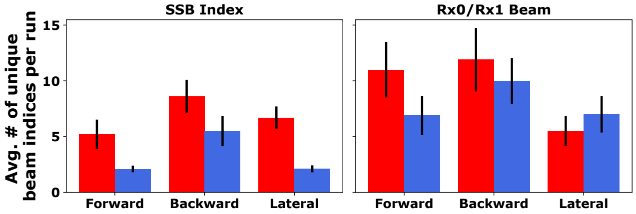

(a) Unique beams statistics.

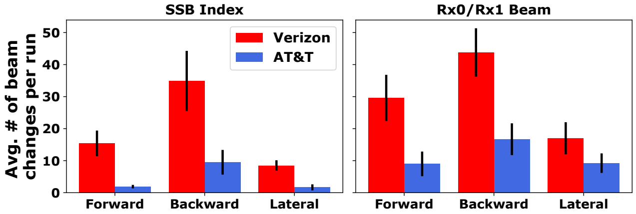

(b) Beam change statistics.

Figure 4: Tx and Rx beam statistics during mobility.

the experiments for each type of motion. Verizon utilizes a larger number of unique beams than AT&T in all three mobility patterns (5.2/8.6/6.7 vs. 2.1/5.5/2.1 in forward, backward, and lateral motion, respectively) and triggers a larger number of beam changes (15.4/34.9/8.5 vs. 1.9/9.5/1.7 in forward, backward, and lateral motion, respectively).

Figs. 5a (left) and 5b (left) plot the CDFs of the Tx beam coherence time, i.e., the time during which the BS maintains the same beam. Once again, we observe that, thanks to the wider beams, AT&T has a much longer beam coherence time than Verizon (median 16.7/1.3/4.2 s vs. 0.95/0.5/3 s, 75-th percentile 20.9/2.6/6.7 s vs. 1.8/1.1/4 s in forward, backward, and lateral motion, respectively). Both operators have the shortest beam coherence time under backward motion, when severe self-blockage triggers frequent beam switches. AT&T has the longest beam coherence time under forward motion; in the most extreme case, we noticed the AT&T BS maintained the same beam for 25 s. On the other hand, this is not always true for Verizon; the coherence time has a longer tail under forward motion, but in general it exhibits lower values compared to lateral motion, due to the difficulty of narrow beams to point downwards, as we explained previously.

4.2 Rx Beams

The beamwidth of the Tx beams also affects beam management on the UE side. A narrower Tx beamwidth will generally result in more beam misalignments as the UE moves and in a higher number of Rx beam switches. This is shown in the example timelines in Figs. 2, 3, where we observe a larger number of Rx beam switches for Verizon than for AT&T. Similarly, Figs. 4a (right) and 4b (right) show that the UE uses a smaller number of unique Rx beams with AT&T than with Verizon in forward and backward motion (6.9/10 vs. 11/11.9) and triggers fewer Rx beam switches in all 3 types of motion (9/16.7/9.2 vs. 29.6/43.8/17.0). The only exception is under lateral motion, where the UE uses a slightly higher number of unique Rx beams with AT&T (7 vs. 5.5) but still triggers a beam switch less often compared to Verizon. Figs. 2, 3 also show that the two Rx beams always change simultaneously; we confirmed this behavior in all our experiments.

Figs. 5a (middle, right) and 5b (middle, right) plot the CDFs of the beam coherence time for the two Rx beams with Verizon and AT&T, respectively. Similar to the Tx beam case, Rx beam coherence times are longer with AT&T (median 2/0.89/0.9 s vs. 0.5/0.5/0.5 s, 75-th percentile 4.18/1.45/2.05 s vs. 1.29/0.89/1.07 s in forward, backward, and lateral motion, respectively). Also, similar to the Tx beam case, the Rx beam coherence time is the shortest under backward motion with both operators, and exhibits its longest values under forward motion with AT&T. On the other hand, the Rx beam coherence time is similar under forward and lateral motion with Verizon. Interestingly, Fig. 5 shows that for a given trajectory and operator, the Rx beam coherence times are much shorter than the Tx Beam coherence time, i.e., the UE switches beams much more often that the BS.

4.3 SINR and Throughput

The choice of Tx beamwidth poses a design tradeoff. Wider beams are more robust to blockage, can cover a longer mobility range, and reduce the overhead of beam switching at the cost of reduced directivity, and hence, lower SINR in LoS conditions. Fig. 6a, which plots the CDFs of SINR with the two operators under each type of motion, shows that narrow beams indeed help Verizon maintain significantly higher SINR than AT&T under all three types of motion; the gap in the median case is 2.13 dB/7.57 dB/11.93 dB under forward, backward, and lateral motion, respectively. The higher SINR allows the Verizon BS to support higher MCS (Fig. 6b) and eventually achieve higher throughput than the AT&T BS (Fig. 6c), in spite of the much more frequent beam switches. In the median case, Verizon's Pcell throughput is higher than AT&T's throughput by 198.43/163.17/223.81 Mbps under forward, backward, lateral motion, respectively.

Another factor that contributes to performance is the beam switch delay. Fig. 6d shows that Verizon switches beams much faster than AT&T (0.32 vs. 1.12 s in the median case, 2.24 s vs. 7.84 in the worst case). The shorter beam switch delays counterbalance the overhead of frequent beam switches resulting in higher performance.

4.4 Interaction with MCS

Figs. 2 and 3 also demonstrate how beam management interacts with rate adaptation, the primary mechanism used by all wireless networks to deal with changes in link quality. Verizon heavily relies on beam management to deal with link

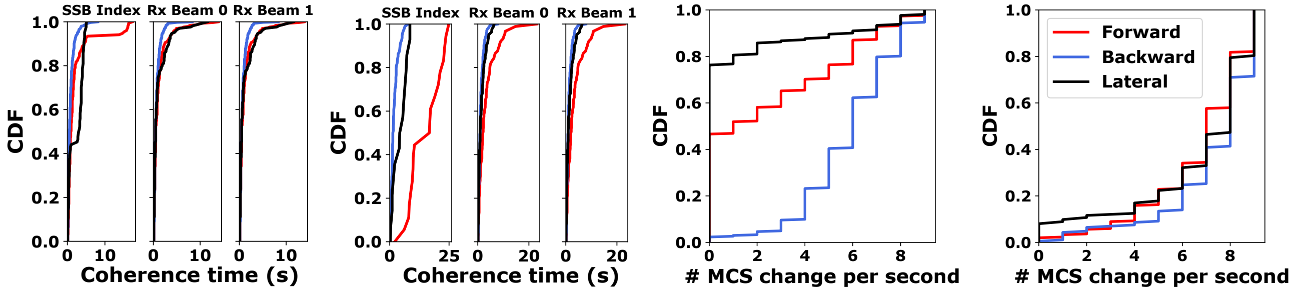

(a) Verizon. (b) AT&T. (c) Verizon. (d) AT&T.

Figure 5: CDFs of coherence time of SSB Index and number of MCS changes per s.

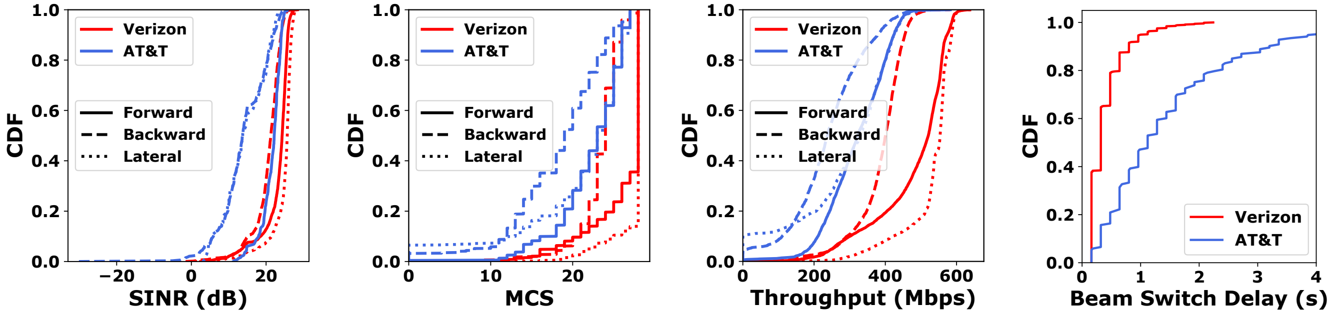

(a) SINR. (b) MCS. (c) Pcell throughput. (d) Beam switching delay.

Figure 6: CDFs of MCS, SINR, PCell Throughput, and beam switching delay.

Table 1: % of time when the selected beam is different from the best beam.

	Verizon	AT&T
Forward	5.0%	0.1%
Backward	18%	20%
Lateral	0.8%	0.6%

Table 2: % of time when that selected beam is not the best-RSRQ beam.

	Verizon	AT&T
Forward	35%	24%
Backward	57%	30%
Lateral	33%	15%

quality due to mobility, ensuring high SINR in general with the exception of the end of the towards and lateral motions, as we described in S4.1. As a result, rate adaptation is triggered infrequently in Figs. 2a-2c. In contrast, AT&T cannot ensure high SINR most of the time using wide beams and relies on frequent rate adaptation to adjust to the varying channel conditions during motion (Figs. 3a-3c). This is also shown in Figs. 5c, 5d, which plot the CDFs of the number of MCS changes per s with the two operators over all the experiments. We observe that AT&T (Fig. 5d) triggers a large number of MCS changes (6-8 per s in the median case) under all three types of motion. In contrast, the number of MCS changes is much lower with Verizon (Fig. 5c), especially in the case of forward and lateral motion, where MCS remains stable 45% of the time and 75% of the time, respectively.

4.5 Beam Tracking Performance

We evaluate beam tracking performance by observing how often the BS chooses a beam other than the best beam. In Table 1, we observe that both AT&T and Verizon select the best beam more than 99% of the time in the case of lateral motion. The same is true for AT&T under forward motion. In contrast, Verizon selects a suboptimal beam 5% of the time under forward motion, suggesting that beam tracking is harder on the vertical plane for narrow beams. This is also reinforced by looking at Fig. 2a, where the BS keeps switching back and forth mostly between two beams, indicating some uncertainty in the beam tracking mechanism for determining the best beam. In the case of backward motion, both operators perform poorly, selecting a beam other than the best beam roughly 20% percent of the time. This is expected, since, in the presence of body blockage, the UE often has to rely on non-LoS paths, the signal strength of which may change frequently and arbitrarily with mobility, thereby increasing the probability of choosing wrong beams.

Note that, following 3GPP notation, the term "best beam" indicates the highest-RSRP beam measured at the UE. Accordingly, there are two reasons due to which the BS may not choose the best beam. First, the UE's reporting is triggered only when certain conditions are met, e.g, RSRP drops below a threshold [1]; thus, the gNB does not always have the most up-to-date beam information. Second, the BS does not always follow the UE's recommendations even when updated information is reported, since it has to consider multiple perspectives including network load, thresholding, and beam switch delays. For example, the BS may decide to continue with the same beam if the RSRP difference between the serving beam and the best beam is below a threshold.

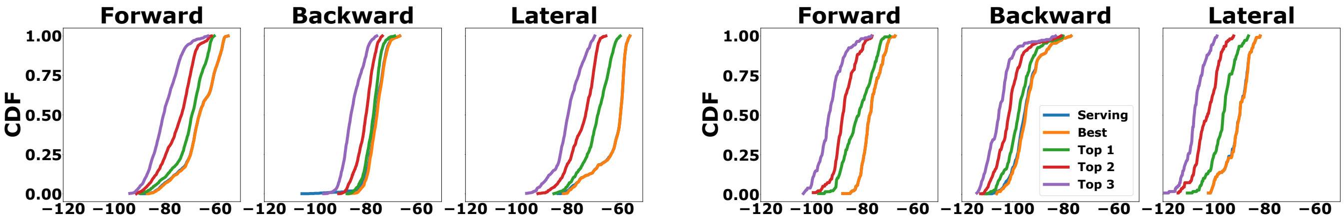

(a) Verizon. **Figure 7: CDFs of BRSRP of different beams.** (b) AT&T.

Fig. 7 plots the CDFs of the beam RSRP (BRSRP) for the best, serving, and top 3 candidate beams for the two operators under each type of motion. We observe that the difference in RSRP between the serving beam and the best beam is almost always negligible for both operators. When the serving beam is not the best beam, the median BRSRP difference between the two is 1.4, 2.2, 1.5 dB for Verizon and 0.4, 1.8, 1.3 dB for AT&T under forward, backward, and lateral motion, respectively, emphasizing the marginal impact of not selecting the best beam. Additionally, the RSRP of the serving beam is significantly higher than the RSRP of the other candidate beams. Together, these two observations demonstrate the ability of both operators to almost always maintain high-quality, near-optimal beams with only a few exceptions in the case of Verizon under backward motion (notice the long tail for the serving beam in Fig. 7a (middle)).

In Table 2, we explore another dimension of beam tracking performance by evaluating how frequently the selected beam has the highest RSRQ. While RSRP quantifies the received signal strength of the beam, RSRQ accounts for both the strength of the received signal as well as the level of interference and noise. As such, beams with the higher RSRQ may be more desirable. We observe from Table 2 that Verizon selects beams with suboptimal RSRQ more than 50% of the time under backward motion and more than 30% of the time under forward and lateral motion. Although AT&T demonstrates better performance, it still selects a lower-RSRQ beam at least 15% of the time. This raises an interesting question about whether RSRP should be the sole reporting metric for beam tracking in 5G mmWave. The common assumption is that in signal propagation environments with high received signal strength or minimal noise and interference, beams with the best RSRP also have the best RSRQ. However, our experimental results suggest otherwise.

5 CONCLUSION

We conducted a measurement study of the beam management procedures in operational 5G-mmWave networks. Our study shows that the two major US cellular operators employ very different beam management parameters and these differences have a direct impact on performance. Our measurements reveal that narrower beams employed by Verizon yield higher SINR, resulting in higher and more stable MCS and higher throughput compared to wider beams employed by AT&T, which heavily relies on rate adaptation to cope with changes in link quality due to mobility. The overhead of frequent beam switching in Verizon due to narrow beams is counterbalanced by the shorter beam switching delay compared to AT&T. We also found that both operators select near-optimal beams in terms of RSRP in diverse scenarios. However, the current RSRP-based beam selection strategy is ineffective in choosing the highest RSRQ beam 15-57% of the time in different scenarios. Our findings are a first step towards a comprehensive understanding of 5G NR beam management procedures in diverse deployment scenarios.

REFERENCES

[1] Online. *Requirements for support of radio resource management.* Technical Report 38.133 version 15.6.0 Release 15. 3GPP.

[2] Online. *Snapdragon X60 5G Modem-RF System.* https://www.qualcomm.com/products/technology/modems/snapdragon-x60-5g-modem

[3] Online. *Study on New Radio (NR) access technology: Physical layer aspects.* Technical Report 38.912 version 14.0.0 Release 14. 3GPP.

[4] Online. *XCAL Solo.* https://accuver.com/sub/products/view.php?idx=11;.

[5] Phuc Dinh, Tri Minh Nguyen, Sanaa Sharafeddine, and Chadi Assi. 2019. Joint Location and Beamforming Design for Cooperative UAVs With Limited Storage Capacity. *IEEE Transactions on Communications* 67, 11 (2019), 8112–8123.

[6] Shuangfeng Han, Chih-lin I, Zhikun Xu, and Corbett Rowell. 2015. Large-scale antenna systems with hybrid analog and digital beamforming for millimeter wave 5G. *IEEE Communications Magazine* 53, 1 (2015), 186–194.

[7] Marwan Krunz, Irmak Aykin, Sopan Sarkar, and Berk Akgun. 2023. Online Reinforcement Learning for Beam Tracking and Rate Adaptation in Millimeter-wave Systems. *IEEE Transactions on Mobile Computing* (2023).

[8] Cen Lin, Geoffrey Ye Li, and Li Wang. 2017. Subarray-Based Coordinated Beamforming Training for mmWave and Sub-THz Communications. *IEEE Journal on Selected Areas in Communications* 35, 9 (2017), 2115–2126.

[9] Arvind Narayanan, Muhammad Iqbal Rochman, Ahmad Hassan, Bariq S. Firmansyah, Vanlin Sathya, Monisha Ghosh, Feng Qian, and Zhi-Li Zhang. 2022. A Comparative Measurement Study of Commercial 5G mmWave Deployments. In *Proc. of IEEE INFOCOM.*

[10] Hossein Soleimani, Raul Parada, Stefano Tomasin, and Michele Zorzi. 2019. Fast Initial Access for mmWave 5G Systems with Hybrid Beamforming Using Online Statistics Learning. *IEEE Communications Magazine* 57, 9 (2019), 132–137.

[11] Bingpeng Zhou, An Liu, and Vincent Lau. 2019. Successive Localization and Beamforming in 5G mmWave MIMO Communication Systems. *IEEE Transactions on Signal Processing* 67, 6 (2019), 1620–1635.