

Vivisecting Beam Management in Operational 5G mmWave Networks

YUFEI FENG*, PHUC DINH*, MOINAK GHOSHAL, and EDUARDO BAENA, Northeastern University, USA

HAMZA BOUCHEBBAH, Kaizen Tech, Japan

DIMITRIOS KOUTSONIKOLAS, Northeastern University, USA

Beam management plays a fundamental role in 5G mmWave networks, due to the highly directional nature of mmWave signals. Despite its importance, beam management procedures and real-world performance in operational deployments remain underexplored. This paper addresses this gap by presenting the first in-depth empirical study of beam management procedures in commercial 5G mmWave networks. Our study is based on an extensive measurement campaign across two major US operators, covering six cities, different base station types, and varying mobility scenarios, including walking and driving. We evaluate key beam management parameters on both the base station and the user equipment side, assessing their impact on network performance. In addition, we examine the interaction between beam management and two other mechanisms critical to performance, rate adaptation and carrier aggregation. Finally, we quantify the beamforming overhead and analyze the effectiveness of beam tracking. Our findings provide novel insights into the real-world performance of beam management in 5G mmWave networks, offering guidance for further optimization.

CCS Concepts: • **Networks** → **Mobile networks**; **Network measurement**; **Network performance analysis**.

Additional Key Words and Phrases: 5G, Measurement, Beam Management, Network Performance, Beam Refinement, Dataset

ACM Reference Format:

Yufei Feng, Phuc Dinh, Moinak Ghoshal, Eduardo Baena, Hamza Bouchebbah, and Dimitrios Koutsonikolas. 2025. Vivisecting Beam Management in Operational 5G mmWave Networks. *Proc. ACM Netw.* 3, CoNEXT2, Article 12 (June 2025), 26 pages. <https://doi.org/10.1145/3730982>

1 Introduction

The wide-scale deployment of 5G over the past five years represents a significant advancement in mobile communications, providing enhanced speed, connectivity, and capacity. Central to this advancement is the use of millimeter wave (mmWave) frequencies, above 24 GHz. This spectrum offers extensive bandwidth and very high data rates, necessary for supporting the requirements of modern applications, such as ultra-high-definition video streaming, mobile AR/VR, and connected autonomous vehicles [28, 32, 35, 76]. However, the propagation characteristics of mmWave signals – limited range and high susceptibility to obstacles – present considerable challenges to the wide-scale deployment of 5G mmWave networks. Due to these challenges, in the early stages of 5G

*Both authors contributed equally to this research.

Authors' Contact Information: Yufei Feng, feng.yuf@northeastern.edu; Phuc Dinh, dinh.p@northeastern.edu; Moinak Ghoshal, ghoshal.m@northeastern.edu; Eduardo Baena, e.baena@northeastern.edu, Northeastern University, Boston, Massachusetts, USA; Hamza Bouchebbah, Kaizen Tech, Tokyo, Japan, hamza.bouchebbah@kaizen-tech.net; Dimitrios Koutsonikolas, Northeastern University, Boston, USA, d.koutsonikolas@northeastern.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 2834-5509/2025/6-ART12

<https://doi.org/10.1145/3730982>

Table 1. #beams per gNB/op. frequency for the two US carriers that offer 5G mmWave service in 11 US cities.

	Los Angeles	Las Vegas	Salt Lake City	Denver	Chicago	Indy	Cleveland	Boston	Atlanta	Miami	Charlotte
Verizon	36/28 GHz	36/28 GHz	36/28 GHz	36/28 GHz	36/28 GHz	36/28 GHz	144/39 GHz	144/224 GHz	36/224 GHz	36/28 GHz	36/28 GHz
AT&T	24/39 GHz	32/39 GHz	-	32/39 GHz	24/39 GHz	24/39 GHz	24/39 GHz	24/39 GHz	24/39 GHz	24/39 GHz	24/39 GHz

deployment, 5G mmWave services were extensively deployed only in two countries, US and Japan. Nonetheless, recently, there is a renewed interest in the use of mmWave frequencies, with several other countries in the early stages of commercialization, including Germany, Italy, Finland, and Spain in Europe, Taiwan, Hong Kong, and Singapore in Asia, and Australia [71].

5G mmWave networks rely on *beamforming* on both the transmitter (Tx) and the receiver (Rx) end to cope with the high propagation loss in the mmWave frequency bands. Beamforming involves the use of multiple antennas (antenna arrays) to direct signal transmission or reception to specific directions (beams), rather than broadcasting signals in all directions. 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.

In practice, beam management depends on a large number of factors, including but not limited to operating band, number of Tx/Rx beams, beamwidth, and interference management. Additionally, while conforming to 3GPP standards, operators and vendors have considerable freedom in implementing beam management logic, such as the number of active Tx/Rx beams, how often beam quality is reported, or whether to use beam refinement, yet the implications of these decisions are rarely visible to practitioners and researchers and remain poorly understood. For example, Table 1 lists the number of beams on a 5G mmWave base station (gNB) and the operating frequency for Verizon and AT&T, the two US operators that extensively use 5G mmWave services, in 11 major US cities, showing a very diverse landscape not only across operators, but also across cities for the same operator. Further, 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 gNB or the user equipment (UE) side. Consequently, understanding the details of beam management in commercial 5G mmWave deployments is very challenging, but, at the same time, extremely important, as it can provide valuable insights into the performance of 5G mmWave networks, and enable realistic simulation/emulation studies.

Nonetheless, the details of beam management in commercial deployments and its performance in real-world scenarios remain largely unknown. The only experimental studies of beam management to our best knowledge are those in [23, 51], but they are limited in both breadth and depth. Narayanan *et al.* [51] studied beam management of the two major US operators in Chicago, focusing on coverage. In our preliminary work [23], we studied beam management of the same two operators in Boston and found that they use very different beam management parameters compared to those reported in [51] for Chicago, highlighting the need for extensive measurement campaigns in different cities to fully understand the details of beam management in commercial deployments. Furthermore, these studies do not explore fundamental aspects of beam management, such as beam refinement, interaction of beam management with carrier aggregation, or beam reporting overhead.

This work fills this gap by conducting the *first* large-scale empirical study of beam management in operational 5G mmWave networks. Via an extensive measurement campaign spread out over a 16-month period (June'23-Sep.'24), spanning six US cities (Atlanta, Boston, Charlotte, Chicago, Las Vegas, Miami), two operators (Verizon, AT&T), and different types of mobility (walking, driving), using off-the-shelf smartphones and a commercial tool that captures lower layer KPIs and signaling messages, we collected a large cross-layer dataset of 5G mmWave beamforming-related KPIs and signaling events. Using the collected dataset, we conduct the first comprehensive characterization of beam management procedures in operational 5G mmWave networks and their impact on performance. The key contributions and findings of our study are summarized as follows:

- We analyze beam management on the gNB and UE side and study the impact of the number of available beams, traffic direction, and motion type (§4.1, §8). We find that beam management procedures in commercial 5G gNBs and UEs utilize only a small subset of their beams under typical mobility patterns and trigger beam switches rather infrequently – every 0.6-2.8 s on average. The number of available beams (which determines the beamwidth), vendor-specific 3D beamforming implementations, and motion type have a major impact on beam management, determining the rate of beam switches and the number of beams they utilize. We also find that uplink (UL) channels (from the UE to the gNB) are affected more by body blockage than downlink (DL) channels (from the gNB to the UE); as a result, UL traffic triggers more beam changes than DL traffic.
- We explore the impact of beam management on performance and the interplay between beam management and rate adaptation (RA) (§4.2). Our results show that narrow beams generally result in better link quality, support higher and more stable MCS, and provide higher throughput in spite of more frequent beam changes compared to wide beams. In contrast, gNBs featuring wide beams rely heavily on RA to deal with changes in link quality during motion.
- We analyze for the first time the overhead of beam management procedures (§5). Our results show that UEs send beam measurement reports every 160-310 ms on average, but many of these reports do not trigger beam changes on the gNBs, suggesting that there is scope for improvement with respect to the signaling overhead.
- We evaluate the quality of selected beams under motion (§6). Our results show that proprietary algorithms in commercial gNBs select near-optimal beams most of the time. Although the selected beams are not always the "best" beams (in terms of signal strength), their signal strength is within 1-2 dB from the signal strength of the best beam more than 90% of the time, highlighting a key insight of our study: despite efforts to design increasingly sophisticated beam switching algorithms, current commercial systems already perform this task quite effectively and further optimization of beam selection alone may yield diminishing returns.
- We study for the first time the interplay between beam management and carrier aggregation (CA) (§7). In contrast to IEEE 802.11ad mmWave systems that use a single beam over 2 GHz channels, we find that 5G mmWave commercial gNBs perform beam measurements and beam selection on a per carrier basis, even though this is not mandated by 3GPP. Our analysis suggests that this choice is often justified, showing high frequency selectivity even for adjacent carriers just 100 MHz apart.
- We study for the first time the beam refinement procedure in commercial 5G mmWave systems (§9). Our results show that beam refinement typically yields substantial signal strength gains, although it can sometimes result in signal degradation, especially under increased self-blockage or high mobility speeds. However, these gains come at the cost of significantly increased reporting overhead. Further, gNBs featuring a very large number of beams without beam refinement yield overall better throughput, as they still offer high SINR while avoiding the additional reporting overhead of the beam refinement process.
- Our measurements in walking and driving scenarios show that beam management performance is similar under both mobility patterns, but user performance is much lower under driving compared to walking (§8). Furthermore, our experiments with two different phone models show that the higher number of beams in the newer model does not always lead to better beam selection, but the newer phone model often yields higher throughput despite suboptimal beam selection (§10), indicating that other hardware and software improvements may decrease the impact of the beamforming gain on overall performance. Together, these observations suggest that future research efforts should not focus on improving beamforming performance alone but instead on the interplay between beam management and other MAC/PHY layer mechanisms that also affect throughput (e.g., MIMO, RA, or CA).

Our findings contribute to a deeper understanding of beam management in operational 5G mmWave networks. By providing practical insights into the performance and challenges of beam management in real-world deployments, our study aims to inform future network optimization and the development of more accurate simulation/emulation models. To assist towards this direction, we make our dataset publicly available [1].

2 A Primer on 5G mmWave Beam Management

Beam management in 5G mmWave networks consists of a set of Layer 1 and Layer 2 procedures designed to establish and maintain an optimal beam pair between the UE and the gNB. Although a large part of beam management algorithms is proprietary and vendor/device specific, 3GPP specifies the main components of beam management [35, 33] shown in Fig. 1.

Initial beam acquisition. Initial beam acquisition occurs when a UE attaches to a gNB. A 5G mmWave gNB periodically transmits Synchronization Signal Blocks (SSBs) in multiple directions through SSB beam sweeping. Each SSB is tagged with a unique SSB index corresponding to

Fig. 1. Beam management procedure.

a specific SSB beam. Throughout the paper, the terms SSB Index and Beam Index are used interchangeably. The UE aligns its own beam(s) to scan for incoming SSBs across different directions, measures the signal strength of the received SSBs, identifies the best gNB's beam (corresponding to the SSB with the highest signal strength), and reports the index of that beam to the gNB. At the same time, the UE adjusts its beam to optimize alignment with the selected gNB beam.

SSB Beam switching. Beam switching is crucial to maintaining gNB-UE beam alignment and high signal quality, especially under UE mobility. For DL transmissions, the gNB periodically transmits SSBs over the Physical Broadcast Channel (PBCH) in multiple directions. The UE receives these SSBs using multiple receive beams, measures the signal strength of different gNB-UE beam pairs, and reports the signal strength (SS-RSRP) of a subset of them to the gNB. If a (vendor-specific) beam switch triggering condition is met, the gNB informs the UE of its new SSB beam through a Transmission Configuration Indicator (TCI) state ID. The UE then switches to the corresponding best receive beam and sends a HARQ ACK to the gNB, finalizing the beam switching process. If beam re-orientation is not enabled, DL data traffic proceeds on the newly established beam pair after these steps. For UL transmissions, we found that both phone models we used in our experiments rely on beam correspondence: the UE uses its best receive beam and the gNB its best SSB beam, identified during the DL procedure, to transmit and receive data, respectively.

Beam re-orientation and refined beam switching. Beam re-orientation is an optional process designed to enhance communication between the UE and gNB after initial beam acquisition or switching to a new SSB beam and before data transmission. Refined beams, often referred to as CSI-RS beams, are narrower than SSB beams and improve communication quality. While beam re-orientation can theoretically occur on both the gNB and UE side, the phones we used in our experiments do not implement beam re-orientation. Therefore, we focus on the gNB's beam re-orientation process. In the DL case, the gNB transmits multiple Channel State Information Reference Signals (CSI-RS) on refined beams. The UE measures these signals and reports the best beam back to the gNB, enabling the gNB to select a refined beam within the coarser SSB beam. In the UL case, consistent with SSB beam switching, we found that gNBs rely on beam correspondence and reuse the refined beams selected during the DL procedure to receive data.

Beam management on different component carriers. CA is a key feature of 5G mmWave, simultaneously utilizing up to 8 mmWave cells or component carriers (CCs) in the DL direction and up to 4 cells in the UL direction for increased throughput. Each CC has a 100 MHz bandwidth, resulting in a total of up to 800/400 MHz of continuous spectrum in the DL/UL direction. 3GPP specifies two approaches to beam management across different CCs. The first approach simplifies operations and reduces control overhead, performing beam measurements and reporting only the primary CC (PCell), which is used for initial access and control signaling along with data transmission, and using the same beam for all secondary CCs (SCells). The second, more precise approach, performs separate beam measurements and reporting for each CC, ensuring optimal alignment but increasing control overhead.

Beam reporting interval. The UE constantly receives reference (SSB or CSI-RS) signals, measures their signal strength, and reports these measurements to the gNB for beam tracking. 3GPP specifies three reporting mechanisms: periodic, aperiodic, and semi-persistent. In periodic reporting, the gNB sets the reporting periodicity during UE admission, and the UE reports at fixed intervals via the Physical UL Control Channel (PUCCH). In aperiodic reporting, the gNB configures the UE to respond to sporadic trigger commands delivered through DCI messages on the Physical UL Control Channel (PDCCH); however, in this mode, the beam reports are sent over physical UL shared channel (PUSCH) instead of PUCCH. Semi-persistent reporting combines features of both periodic and aperiodic reporting, allowing the UE to transmit measurements at regular intervals while also accommodating sporadic trigger events from the gNB. Implementation details are vendor specific.

3 Methodology

3.1 5G Devices and Operators

3.1.1 5G Devices. We used Samsung S21 and S24 phones in different cities (See Table 3). Both models support the 5G mmWave bands n260/261 operating at 39/28 GHz, respectively, 8-CC CA in the DL direction and 2-CC/4-CC CA in the UL direction. Most of our results in this work are with S21 phones. We compare the beam management on the two UE models and the impact on performance in §10. We also used an Accuver XCAL Solo device, which taps into the Qualcomm diagnostic (Diag) interface of the smartphone and extracts lower layer KPIs and control-plane signaling messages necessary for a holistic understanding of the beam management procedures.

As noted in [51], the S21 phone uses two antennas, each of which can simultaneously transmit/receive over up to 8 CCs. Each antenna supports 36 beams. We found that the S24 phone uses a similar design but each antenna supports 43 beams. In other words, the phones always use two beams simultaneously. However, each beam in one antenna can only be combined with one specific beam in the other antenna, i.e., there are 36/43 possible beam pairs for S21/S24. In addition to the beam index of each antenna, XCAL also uses a third index to refer to a specific beam pair on the UE. For simplicity, we use the third index instead of showing results separately for each UE antenna.

3.1.2 5G Operators. We performed all the measurements in Table 2. Number of beams per panel/gNB in downtown areas of 6 cities using Verizon's and AT&T's 5G mmWave services. The two operators offer services in different bands (n260/n261, see Table 1), support different numbers of CCs (see Table 2), and have different numbers of beams per gNB covering all possible options (24, 32, 36, 144) in Table 1.

City	Beams per panel/gNB		Carrier agg. (DL/UL)	
	AT&T	Verizon	AT&T	Verizon
Atlanta	-	12/36	-	4CC/2CC
Boston	8/24	48/144	4CC/2CC	6CC/2CC
Charlotte	8/24	12/36	4CC/2CC	8CC/2CC
Chicago	8/24	12/36	8CC/4CC	8CC/4CC
Las Vegas	8/32	12/36	8CC/2CC	8CC/4CC
Miami	8/24	12/36	8CC/2CC	6CC/2CC

A typical 5G mmWave gNB structure consists of 3 antenna panels, each covering a 120° sector (see Figs. 22a, 22b in §C.2). The Verizon gNBs in all 6 cities of our study and AT&T gNBs in 5 cities use this structure. The only exception is some

AT&T gNBs in Las Vegas, consisting of 4 panels, each covering a sector (see Fig. 22c in §C.2). For Verizon gNBs, each panel is identified by a Physical Cell Identifier (PCI), whereas for AT&T gNBs, all panels share the same PCI. Verizon gNBs have 48 SSB beams per panel in Boston and 12 beams per panel in the other 5 cities, or 144/36 beams per gNB covering a circle. AT&T gNBs have 8 beams per panel in all 6 cities, or 32 beams per gNB in Las Vegas and 24 beams per gNB in the other 5 cities. Note that, unlike the UEs (§3.1.1), gNBs use one panel and hence one beam to communicate with a given UE.

3.2 Experiments

We used iperf3 to generate backlogged TCP DL and UL traffic from/to a Google cloud server located in Washington, DC for all the throughput measurements. The ingress/egress network bandwidth of the server was 16+ Gbps, much higher than the maximum 5G mmWave DL throughput (6 Gbps).

We conducted experiments under three mobility patterns: controlled walking, random walking, and driving, summarized in Table 3 for different cities. Our controlled walking experiments (Fig. 21 in §C.1) involve walking towards, away from, and laterally to a 5G mmWave BS, at the typical walking speed

Table 3. Summary of mobility patterns and UE models used in different cities and operators.

City	Verizon	AT&T
Atlanta	random walk (S21), drive (S21)	-
Boston	controlled walk (S21, S24), drive (S21)	controlled walk (S21)
Charlotte	controlled walk (S21)	controlled walk (S21)
Chicago	drive (S24)	drive (S24)
Las Vegas	drive (S24)	controlled walk (S21), drive (S24)
Miami	controlled walk (S21), drive (S21)	drive (S21)

(~3 ft/s). Given the deployment of 5G mmWave gNBs on traffic lights and lamp posts, we believe these three mobility patterns represent realistic pedestrian mobility scenarios. Most of our results are obtained via controlled walking experiments in front of Verizon gNBs in Boston, Charlotte, and Miami and AT&T gNBs in Boston, Charlotte, and Las Vegas. The selection of these cities covers all four possible numbers of beams per gNB (24, 32, 36, 144). 5G mmWave gNBs in Boston, Charlotte, and Miami are deployed across streets, and we were able to easily identify the gNB to which the UE was connected and complete the three trajectories in front of it. In Las Vegas, we conducted our controlled walking experiments in an empty parking lot. In contrast, in Atlanta, we conducted our experiments in the Centennial Olympic Park, where a large number of Verizon gNBs are deployed with overlapping coverage, making it extremely challenging to identify the gNB the UE was connected to at a given time. Hence, in Atlanta, we collected data under random walking, and we use these data in §8. Finally, we collected data under driving in cities with dense mmWave deployments along downtown streets: Atlanta, Boston, Chicago, Las Vegas, Miami for Verizon and Chicago, Las Vegas, Miami for AT&T. The driving trajectories are shown in Fig. 26 in §C.4. We use this dataset in §8, §9, §10.

During the walking experiments, we sometimes experienced handovers to a different 5G mmWave gNB. We removed such traces and report results from 10 traces (5 DL, 5 UL) for each trajectory-operator combination where there were no handovers. The duration of these traces varies from 15-40 s for different cities and operators, depending on the gNB location. Similarly, for the driving experiments, each spanning a few hours, we experienced multiple handovers to other 5G or LTE base stations. We removed those segments from the driving traces and only used the segments of the traces where the UE was connected to a 5G mmWave gNB.

4 Impact of Beamwidth

We begin by examining the impact of beamwidth on beam management procedures, using data from our controlled walking experiments (Table 3). We use the term *beam resolution* to refer to the number of beams per gNB. With perfectly-shaped beams of equal beamwidth, 24/32/36/144 beams would each have a beamwidth of $\pi/12$ to cover the 360° space. However, this assumption is not valid in practice. Beams generated by phased arrays in commercial gNBs have imperfect shapes with sidelobes, and also support 3D beamforming. As a result, each beam typically

(a) Unique beams observed. (b) Beam changes per second. (c) Beam coherence time.

Fig. 2. Beam management statistics DL, gNB side.

(a) Towards.

(b) Away.

(c) Lateral.

Fig. 3. Number of unique beams per run with different mobility patterns DL, gNB side.

has a different shape and beamwidth. Nonetheless, we verified that a very large beam resolution (144 beams) results in much narrower beams compared to the other three resolutions (see Fig. 2b). Consequently, in our analysis, we divide the figures into two regions separated by a dashed line: the left side represents gNBs with wide beams (24/32/36 beams per gNB), and the right side represents gNBs with narrow beams (144 beams per gNB).

We first present the overall beam management statistics for different beam resolutions in Sec. 4.1 and then analyze the impact of beamwidth on performance in Sec. 4.2. To decouple the impact of beam management on performance from the impact of CA, we focus on the performance and beam management of the PCell throughout the paper, and study beam management across different CCs in Sec. 7. Under backlogged traffic, the MAC throughput of each CC is approximately equal to $\frac{1}{N}$ of the total MAC throughput, where N is the total number of aggregated cells [20].

4.1 Overall Statistics

Fig. 2 presents the overall beam statistics on the gNB side for DL traffic and 4 different beam resolutions. We examine the average number of unique beams observed during a run, the average number of beam changes per s, and the beam coherence time, i.e., how long a certain beam is used before switching to another beam. The error bars in Fig. 2 and all the following figures denote the standard deviations.

We observe that a very high beam resolution (144 beams) results in a higher number of unique beams per run (8.3 vs. 3.4-4.3 in Fig. 2a), a higher number of beam changes per s (0.68 vs. 0.1-0.2 in Fig. 2b), and shorter coherence time (1.37 s vs. 2.5 s in Fig. 2c) compared to the other three (much) lower resolutions. This is expected as beam misalignment is more frequent during motion in the case of narrow beams. However, the three metrics do not change monotonically with respect to beam resolution for the three low beam resolutions. To analyze this, we present the number of unique beams observed for each mobility pattern in Fig. 3. For the case of 144 beams/gNB, we observe a consistently higher number of unique beams across all mobility types. In contrast, for 24, 32, and 36 beams/gNB, the number of unique beams observed varies for different mobility

(a) Unique beams observed. (b) Beam changes per second. (c) Beam coherence time.

Fig. 4. Beam management statistics UL, gNB side.

(a) Unique beams observed. (b) Beam changes per second. (c) Beam coherence time.

Fig. 5. Beam management statistics DL, UE side.

patterns. For example, gNBs with 32 beams/gNB use the lowest number of unique beams during "towards" mobility (only 2 beams), but the highest during "away" mobility. This suggests that vendors implement 3D beamforming differently, with some focusing more on vertical beams and others on horizontal beams, leading to the non-monotonic trends observed when the total number of active beams is not significantly different in Figs. 2a, 2b, 2c.

Fig. 4 shows beam management statistics on the gNB's side for UL traffic. Our first observation is that, even though the number of unique beams observed are similar for both UL and DL traffic (Figs. 4a vs. 2a), UL traffic triggers more beam changes per second, which results in shorter beam coherent times than DL traffic for 3 out of 4 beam resolutions (Figs. 4b vs. 2b and Figs. 4c vs. 2c). This observation suggests that channel reciprocity does not hold, and UL channels degrade more often than DL channels. In the case of UL traffic, any obstacle near the UE, such as the user's body, can directly obstruct the signal, especially if the user moves in ways that increase blockage (e.g., away from the gNB). In contrast, DL signals transmitted from the gNB travel a greater distance before encountering obstacles, allowing the beam to expand enough to (partly) go around the obstacle.

The beamwidth of the gNB beams also affects beam management on the UE side (Fig. 5). A narrower gNB beamwidth generally results in more beam misalignments as the UE moves and in a higher number of UE beam switches. This is shown in Figs. 5a, 5b, 5c, which plot the average number of unique beams, beam changes per second, and the coherence time on the UE side for DL traffic. Fig. 5b plots both the total number of beam changes per second on the UE side and the number of beam changes triggered by the UE itself, without a change on the gNB side.

Fig. 5a shows that the number of unique beams on the UE side also increases with the gNB beam resolution, from 9.4 (with 24 beams/gNB) to 12.2 beams (with 144 beams/gNB). These numbers are higher than those on the gNB side in Fig. 2a, as the UE performs beam switching both to adapt to beam changes on the gNB side and on its own (without a beam change on the gNB side). This is also shown in Figs. 5b, 5c; beam switching on the UE side occurs much more frequently than on the gNB side (Fig. 2b), and the average beam coherence time is less than 1 s. Figs. 5b also show that most beam changes on the UE side are triggered independently of the changes on the gNB side.

(a) SINR. (b) Pcell throughput. (c) MCS. (d) Rate of MCS changes.
 Fig. 6. Impact of beamwidth on performance and rate adaptation.

An interesting observation here is that the highest beam resolution (144 beams) on the gNB side results in the lowest number of beam changes on the UE side (Fig. 2b). The reason is that when the gNB features a smaller number of wider beams, the UE has to perform beamforming more frequently to maintain a high-quality link. This is also indicated by the number of UE beam changes without a gNB beam change, which accounts for 76-84% of the total beam changes in the case of 24, 32, and 36 gNB beams, compared to only 64% in the case of 144 gNB beams.

In summary, beam resolution, vendor-specific 3D beamforming implementation, and motion type have a major impact on beam management procedures, determining the rate of beam switches and the number of beams they utilize. Despite the differences across different gNB configurations, Figs. 2a, 4a show that the beam management procedure utilizes only a small subset of the available beams on the gNB under typical mobility patterns: fewer than 5, on average, out of 24/32/36 beams and fewer than 8 out of 144 beams. The UE utilizes a larger number of beams under the same mobility patterns (Fig. 5a) but still lower than 1/3 of the available beams on average. Similarly, we found that the UE reports signal strength measurements for only a small number of beam pairs during each run: 0.7-14% (1-19%) of the total beam pairs in the DL (UL) direction under typical mobility patterns (details in §C.3.1). Additionally, Figs. 2b, 2c, 4b, 4c, 5b, 5c show that beamforming is typically triggered infrequently, and both the gNB and UE maintain the same beam for long intervals: 1.2 s and 0.6-0.8 s on average, respectively. Nonetheless, the very large standard deviations, especially in the case of beam coherence time (Figs. 2c, 4c, 5c) show that there are extreme cases where the gNB and/or UE either maintain the same beam for several s or switch beams very fast, within a few ms.

4.2 Impact on Performance and Interplay with Rate Adaptation

The gNB's beam resolution poses a design trade-off. Wider beams are more robust to blockage, can cover a broader area, and can reduce the overhead of frequent beam switching during mobility, at the cost of reduced directivity, and hence, lower SINR in line-of-sight (LoS) conditions. Fig. 6a, which plots the SINR for different gNB beam resolutions, shows that the SINR with a beam resolution of 144 beams is about 6-14 dB higher in the median case compared to the other three resolutions. The higher SINR with a beam resolution of 144 beams allows the gNB to support higher MCS (Fig. 6c) and eventually achieve higher PCell throughput (Fig. 6b), in spite of the more frequent beam switches. In the median case, the PCell throughput with a beam resolution of 144 beams is 240/226/194 Mbps higher than the throughput with a beam resolution of 24/32/36 beams.

However, the SINR and throughput do not scale proportionally with the number of beams, due to several reasons. For instance, gNBs with 144 beams only double the throughput compared to those with 24 beams. First, gNBs with fewer beams often benefit from beam refinement that enhances link quality, whereas gNBs with many beams skip this step, as discussed in §9. Second, as explained in §4 and §B, commercial gNBs feature imperfect beams and 3D beamforming, and hence, the beamwidth (and beamforming gain on the horizontal plane) does not increase proportionally with the number of beams. Furthermore, proprietary vendor/operator policies, e.g., different RA or Tx power control algorithms, might also play a role.

We further explore how beam management interacts with RA, the primary mechanism used in wireless networks to deal with changes in link quality, in Fig. 6d, which shows the rate of MCS changes. gNBs with 144 beams rely heavily on beam management (Figs. 2a, 2b) to deal with link quality degradation due to mobility, ensuring high SINR (Fig. 6a). As a result, RA is triggered infrequently in this case, as shown in Fig. 6d. In contrast, gNBs with 24, 32, or 36 beams cannot ensure high SINR most of the time using wide beams (Fig. 6a) and rely on frequent MCS changes to adjust to the varying channel conditions during motion (Fig. 6d).

In summary, our analysis highlights a fundamental trade-off between beam management and RA: operators can either use narrower beams with more frequent beam switching or rely on RA with wider beams to cope with dynamic channel conditions. Our results show that the first choice results in better performance; narrow beams yield higher SINR, higher and more stable MCS, and higher throughput in spite of more frequent beam changes, compared to wide beams.

5 Beam Management Overhead

In this section, we study the beam management overhead using the beam reporting interval (i.e., how often the UE sends beam measurement reports to the gNB) as the metric of interest. We focus on the overhead for SSB beam management here and discuss the overhead of the beam refinement process in §9. We found that every measurement report sent by the UE reports signal strength for 3 beam pairs, in addition to the beam pair currently used, regardless of city, operator, and gNB type.

(a) DL (b) UL

Fig. 7. Beam reporting interval for DL and UL traffic.

Fig. 7 presents boxplots of the UE beam reporting interval for DL and UL traffic. First, we observe that the reporting interval is not constant, i.e., the UE uses the periodic reporting mechanism (§2). We further observe that the UE sends reports quite frequently every 160-310 ms in the median case (although with some outliers as high as 2 s), but many of these reports do not trigger a beam switch on the gNB (Figs. 2b, 4b, 5b). For example, with a beam resolution of 24 beams and DL traffic, the UE sends reports every 0.46 s on average, but the gNB performs a beam switch only every 2.8 s on average (Fig. 2b). DL traffic generally yields longer reporting intervals than UL traffic (160-310 ms vs. 160 ms in the median case), which is aligned with our observation in §4 that UL traffic triggers a higher number of beam changes per s (Figs. 2b vs. 4b). In addition, the beam resolution does not correlate with the reporting frequency, which suggests that the reporting frequency is mostly affected by environmental factors and/or vendor-specific configurations.

In summary, the UE sends measurement reports frequently, but many of these reports do not trigger a beam switch on the gNB, suggesting that there is scope for further optimization of the beam management process.

6 Beam Switching Analysis

In this section, we evaluate the effectiveness of the beam selection algorithms in operational 5G mmWave networks by examining how often the gNB selects the best beam. According to 3GPP, the term "best beam" indicates the beam with the highest RSRP measured at the UE. On the other hand, the serving beam is the beam chosen by the gNB and is not necessarily the best beam reported by the UE. The selection depends on the proprietary beam switching algorithms implemented by the vendor. For instance, if the signal strength of the current beam is high enough, beam switching may not be triggered even if a higher-RSRP beam is reported. Alternatively, the gNB may not always have the most up-to-date beam information, due to a delayed report from the UE.

(a) "Redundant" reports.

(b) "Useful" reports.

Fig. 8. RSRP of reported beam pairs and RSRP difference between the current beam pair and reported beam pair in "redundant" vs. "useful" reports.

Given that many beam measurement reports are ignored by the gNB, as we saw in §5, we begin our analysis by trying to distinguish between "useful" reports, i.e., reports that trigger a beam switch, and "redundant" reports, i.e., reports that are ignored by the gNB and thus, they could be omitted to reduce the signaling overhead. Since we have no visibility into the proprietary beam switching logic on the gNB side, we explore two simple heuristics: (i) making decisions based on the signal strength (RSRP) of the reported beam pairs and (ii) making decisions based on the RSRP difference between the reported beam pairs and the serving beam pair. When a report does not trigger a beam switch, we consider the highest RSRP included in that report. When a report triggers a beam switch, we consider the RSRP of the beam pair selected by the beam switching logic, even if it was not the highest RSRP in that report.

In Figs. 8a(left), 8b(left), we compare the reported RSRPs in reports that do not trigger a beam switch vs. those that trigger a beam switch, for different beam resolutions. We observe that, for each resolution, the RSRPs in reports that do not trigger a beam switch are generally higher than those in reports that trigger a beam switch, suggesting that the absolute reported RSRP is used as the beam switching criterion.

In Figs. 8a(right), 8b(right), we compare the RSRP difference (between the serving beam pair and the reported beam pair) in reports that do not trigger a beam switch vs. those that trigger a beam switch, for different beam resolutions. A negative RSRP difference indicates that the reported RSRP is higher than the RSRP of the serving beam pair. We observe that the RSRP is negative only 2-12% of the time in the cases that do not trigger a beam switch, but about 40% of the time in the cases that trigger a beam switch, indicating that the RSRP difference is probably used in the beam switch logic. Nonetheless, Fig. 8b(right) also shows that 60% of the "useful" reports trigger a beam switch to a new beam pair that has lower RSRP than the serving beam pair, suggesting a more complex beam switching logic that takes into account additional factors (e.g., history or hysteresis).

Given the complexity of the beam switching logic, we next study in Table 4 the fraction of time the serving beam is different from the best beam.

	# of beams	Towards	Away	Lateral
24	24	29%	22%	3%
32	32	7%	8%	2%
36	36	10%	24%	7%
144	144	5%	24%	5%

differs from the best beam for different beam resolutions and motion patterns. We observe that this fraction is quite small (3-7%) under lateral motion for all four beam resolutions. On the other hand, when the UE moves away from the gNB, a suboptimal beam is selected 22-24% of the time with 3 out of 4 resolutions. This is expected, since, in the presence of self-blockage, the UE often has to rely on non-LoS paths (via reflections); the signal strength over such paths may change frequently and arbitrarily with mobility, increasing the probability of selecting suboptimal beams. We also observe a non-negligible fraction of suboptimal beams for two resolutions when the UE

(a) 24 beams. (b) 32 beams. (c) 36 beams. (d) 144 beams.

Fig. 9. RSRP Difference between serving and best, top 1-3 beams for different # beams per panel.

walks towards the gNB; in particular, for 24 beams, a suboptimal beam is chosen 29% of the time, a fraction even higher than that of "walking away", suggesting again that beam tracking on the vertical plane may be challenging, depending on the implementation of 3D beamforming. Finally, we observe no correlation between the beam resolution and the effectiveness of beam tracking.

To evaluate the quality of the selected suboptimal beams, we plot in Fig. 9 the RSRP difference between the serving and the best beam, as well as between the serving beam and the top-3 candidate beams reported by the UE. We observe that the RSRP difference between the serving and best beam is almost always negligible in the case of 32, 36, and 144 beams (less than 1 dB 96-98% of the time, less than 2 dB 98-99% of the time). For the lowest resolution of 24 beams, these percentages are slightly higher, but still quite small (less than 1 dB 93% of the time, less than 2 dB 96% of the time). On the other hand, the RSRP of the serving beam is typically much higher than the RSRP of other candidate beams.

In summary, beam selection algorithms in operational 5G mmWave networks almost always select high-quality, near optimal beams, highlighting a key insight from our study: despite extensive efforts to design increasingly sophisticated beam-switching algorithms, today's 5G mmWave gNBs already perform this task quite effectively and attempts to further optimize the beam selection algorithms may yield diminishing returns.

7 Beam Management Across Different Carriers

As mentioned in §2, in the case of CA, the UE performs separate beam measurements and reporting for each CC with both operators in all 6 cities. This approach suggests high frequency selectivity, but appears counter-intuitive, given the small CC bandwidth (100 MHz) and total bandwidth (up to 800 MHz of contiguous spectrum in the DL direction, up to 200/400 MHz in the UL direction for S21/S24), and the fact that, in all our experiments, CA is performed over CCs from a single gNB.² If frequency selectivity is non-present across adjacent 100 MHz CCs, performing separate measurements for each CC can lead to an unnecessary increase of UE's power consumption and resource utilization. Hence, in this section, we analyze the efficiency of multi-carrier measurement and reporting.

Fig. 10 shows the fraction of time when all CCs use the same beam. We observe that this fraction is quite low for all beam resolutions and both traffic directions, ranging from 35% to at most 60%. This result indicates that the UE often measures different signal strength over different CCs and multi-carrier measurements may indeed be necessary. In particular, Fig. 10b (UL direction, where only two adjacent CCs are used) shows that frequency selectivity is present even for adjacent CCs just 100 MHz apart.

In spite of frequency selectivity, one could still use the PCell's best beam for all the SCells if the RSRP difference between the PCell's best beam and an SCell's best beam is negligible. In the

¹For comparison, the Wi-Fi mmWave standard (IEEE 802.11ad), operating at a much higher frequency (60 GHz), uses a single beam over much wider channels (2 GHz).

²CA can combine carriers from different gNBs, but we never observed any such case in our dataset.

(a) DL.

(b) UL.

Fig. 11. RSRP difference between an

Fig. 10. Fraction of time when all carriers share the same beam. SCell's and PCell's serving beams.

following, we explore this option. For each SCell, when its serving beam differs from the PCell's serving beam, we identify the top-3 highest-RSRP non-serving beams and calculate the fraction of time when the PCell's serving beam is among these top-3 candidates. The results in Table 5 show that the fraction of time when the Pcell's serving beam is the top-1 non-serving beam for an SCell (with different serving beam) is always less than 70% and can be as low as 53%. In other words, when an SCell selects a beam different from the PCell's beam, the beam chosen by the PCell is often not a good second choice. Additionally, the beam selected by the PCell is not even among the top-3 non-serving beams for an SCell for a non-negligible amount of time (up to 17%). We further analyze the RSRP difference between an SCell's serving beam and the PCell's serving beam (measured at the SCell) in Fig. 11. We observe that the median difference varies from 1.8 dB to 2.8 dB and in extremes cases it can be as high as 10-14.8 dB (more than ten-fold). These high RSRP differences suggest that it is often advantageous to perform beam selection per CC in spite of the increased overhead.

In summary, our results show high frequency selectivity in the 5G mmWave bands, even across adjacent carriers 100 MHz apart, and our analysis suggests that the choice made by operators to perform beam measurements and beam selection on a per carrier basis is often justified.

Table 5. Fraction of time (%) when the PCell's serving beam falls in an SCell's top-3 set.

# of beams	Top1 beam	Top2 beam	Top3 beam	Non-Top beam
24	53	23	10	14
32	69	14	8.5	8.5
36	66	20	9	5
144	55	19	9	17

This frequency selectivity can be attributed to both channel-induced selectivity and hardware-induced selectivity. A key example of hardware-induced selectivity is beam squint in phased arrays, where signals at different frequencies experience frequency-dependent phase shifts, causing the beam to steer differently. This effect is more pronounced at higher frequencies due to smaller wavelengths. Techniques such as true time delay beamforming [59] instead of phase-shifting can eliminate hardware-induced selectivity and reduce the need for beam selection on a per-CC basis, but they are more complex and costlier compared to phase-shifters used to implement commercial off-the-shelf antenna arrays.

8 Impact of Motion Type

In our analysis so far, we only considered walking experiments. In this section, we study the impact of increased mobility (driving) on beam management. We use our walking and driving data from Atlanta, Boston, and Miami with Verizon for this study.

In Fig. 12, we analyze the beam change rate under walking and driving mobility. Surprisingly, walking scenarios often exhibit a higher median rate of beam changes on both the gNB and UE side, for both traffic directions. We found that the reason for this counter-intuitive result is body blockage in walking scenarios, which causes significant signal strength degradation, prompting frequent beam switches as the gNB and UE search for better beam pairs, and in turn, the overall beam

(a) SINR. (b) MCS. (c) Throughput.

Fig. 13. Comparison of SINR, MCS, and Throughput under walking and driving mobility.

change rate. Fig. 24 in $\text{\textcircled{Y}C.3}$, where we removed "walking away" mobility data from our analysis, indeed shows lower beam change rates for walking mobility in this case. Further, Fig. 25 in $\text{\textcircled{Y}C.3}$, which plots the beam change rate separately for each city, shows that in Atlanta, where we only conducted random walking experiments, the beam change rate under walking is lower than under driving. In contrast, during our driving experiments, the phone was attached to the passenger's window and hence, it did not experience any body blockage.

We also found that, despite faster fluctuations of channel quality, the fraction of time during which there is a mismatch between the serving and the best beam is only slightly higher during driving than during walking (13.4% vs. 12.2%). However, the small differences in beam management performance between walking and driving cannot explain the large SINR (19.3 dB vs. 12.9 dB in the median case) and throughput (305 Mbps vs. 150 Mbps in the DL direction, 64 Mbps vs. 23 Mbps in the UL direction) drops under driving compared to walking in Fig. 13. This large performance drop is caused by other factors, e.g., longer UE-gNB distances during driving compared to walking and faster channel fluctuations, which also pose challenges to RA [30] and CA [46] in addition to beamforming.

In summary, our analysis reveals that walking mobility, despite lower speeds compared to driving, can result in higher beam switching rates due to body blockage, when walking away from the gNB. In contrast, driving mobility is characterized by rapid channel variations, but the number of mismatches between serving and best beams is only slightly higher than under walking, further strengthening our conclusion in $\text{\textcircled{Y}6}$ that beam management in today's 5G mmWave gNBs is quite optimized in terms of performance. On the other hand, the user experienced performance under driving (throughput) is significantly lower than under walking, suggesting that future research efforts should focus on optimizing other aspects of 5G, such as RA and CA.

9 Beam Refinement

Our analysis till now only considered SSB beams. In this section, we analyze the impact of beam refinement on performance from two perspectives: signal strength improvement and additional reporting overhead compared to using only SSB beams.

Recall from $\text{\textcircled{Y}2}$ that beam refinement is an optional feature. In our experiments, we found that beam refinement is typically used on the gNB side with two exceptions: Verizon in Boston, where the gNBs have a very large number (144) of narrow beams, and (surprisingly) AT&T in Las Vegas. The lack of beam refinement for AT&T in Las Vegas, which uses only 32 SSB beams, might explain the low performance and high reliance on RA in that city (Fig. 6). When beam refinement is enabled,

(a) DL.

(b) UL.

Fig. 14. Signal strength improvement with beam refinement.

Fig. 15. SSB and CSI-RS reporting overhead.

we found that two refined beams are used on top of a coarse SSB beam. Beam refinement is used regardless of the traffic direction, refining the Tx beams for DL and the Rx beams for UL traffic.

Fig. 14 shows the signal strength improvement with refined CSI-RS beams compared to wider SSB beams for both traffic directions and different types of mobility. In general, the use of refined beams leads to an improved RSRP. However, the benefits of refined beams vary for different types of mobility. The lateral and towards mobility patterns exhibit the highest gains, with median improvements of about 8-8.5 dB for both traffic directions. On the other hand, "walking away" and driving experience the lowest gain, with median improvements of 5-6 dB (4-4.5 dB) in the DL (UL) direction. Further, beam refinement can sometimes result in signal degradation, especially under "walking away" conditions and driving (6-9% of the time), and this degradation can be higher than 5 dB under extreme cases. The reduced performance when the UE walks away from the gNB is likely due to the fact that the narrower beams are more vulnerable to blockage, as they offer fewer multipath alternatives. On the other hand, driving is typically performed laterally to a gNB, but the higher speeds result in more frequent beam misalignments in the case of narrow beams.

Enabling beam refinement incurs additional reporting overhead. Fig. 15 plots the reporting interval in the case of SSB and CSI-RS beams. We observe that the median reporting interval for CSI-RS beams is 3-4x shorter than for SSB beams, indicating a significantly higher reporting overhead for refined beam management (note that when beam refinement is enabled, both SSB and CSI-RS beams are reported). Consequently, the trade-off between improved signal strength and increased signaling overhead must be carefully balanced to optimize system performance, considering that refined beams can be suboptimal in some cases, as shown in Fig. 14.

In summary, beam refinement typically yields substantial signal strength improvements, although it can sometimes result in signal degradation, especially under increased self-blockage or high mobility speeds. Additionally, the RSRP gains come at the cost of significantly increased reporting overhead. We also note that gNBs without beam refinement but with a very large number of SSB beams yield overall higher throughput (Fig. 6b), suggesting that this might be a better strategy, as it still ensures high SINR while it reduces substantially the reporting overhead.

10 Comparison of Different Phone Models

In the previous sections, experiments were conducted using the Samsung Galaxy S21, which integrates the Snapdragon X60 5G Modem-RF System. In this section, we discuss the differences between the S21 model and a newer model, the Samsung Galaxy S24, equipped with the Snapdragon X75 5G Modem-RF System. Re-

(a) gNB side.

(b) UE side.

call from 3GPP that the S24 model uses a larger number of beams compared to the S21 model (43 vs. 36). For a comparison under walking, we use data collected with both phones in Boston with Verizon at the same gNB. For a comparison under

(a) SINR. (b) MCS. (c) Throughput.
 Fig. 18. Comparison of MCS, SINR, and PCell's MAC throughput across different devices and mobility modes.

driving, we use data collected with S21 in Atlanta with Verizon and Miami with both operators, and data collected with S24 in Las Vegas with Verizon and Chicago with both operators, so that both datasets are obtained from gNBs with the same beam resolutions (24 and 36 beams). We focus on DL performance due to space limit, however, the findings are similar for both traffic directions.

Fig. 16 compares the beam change rate between the two phone models. S24 demonstrates consistently higher beam change rates than S21 across both walking and driving scenarios on both the gNB and UE sides. The higher beam change rates suggest the S24 employs a more aggressive beam management approach, with faster responsiveness to channel variations. Additionally, the larger number of (possibly) narrower beams available on S24 results in faster misalignments during motion, triggering beam changes more frequently.

However, the faster beam change rate with the S24 model does not help the gNB to consistently select better beams than with the S21 model, as shown in Table 6. Under driving, the gNB indeed selects the best beam 90% of the time with S24 vs. 85% of the time with S21. However, under walking, the result is reversed; the gNB selects the best beam 86% of the time with S21 but only 78% of the time with S24, suggesting that the larger number of beams on the latter cannot cope well with self-blockage when the user moves away from the gNB.

We also compare the beam refinement gain in Fig. 17. We only consider the driving scenario, as the walking experiments were conducted in Boston with Verizon, where the operator does not support beam refinement. The results show that beam refinement with the newer phone model provides lower RSRP gain than the older model (6 dB vs. 7.5 dB in the median case).

We conjecture that larger number of narrower beams on the S24 phone provide higher RSRP, rendering the beam refinement process less effective.

However, despite the lower quality of selected beams on the gNB under walking and the lower beam refinement gain under driving, the SINR with the S24 model is similar to the SINR of the S21 model (slightly higher during walking, slightly lower during driving, Fig. 18a). Similarly, the throughput with the S24 model is higher than with the S21 model under walking and exhibits a higher 75-th percentile (but lower 25-th percentile) under driving (Fig. 18c). In fact, the higher throughput comes in spite of the lower MCS (Fig. 18b), likely due to improvements in S24's 5G modem and signal processing.

In summary, the results in this section suggest that increasing the number of available beams on newer UE models does not necessarily yield better beam management performance. However, newer models come with a variety of additional hardware and software improvements, which may still yield better overall performance, making the beamforming gain less relevant.

Table 6. % of time the serving beam is different from the best beam for different UE models.

Mobility Mode	S21	S24
Walking	14	22
Driving	15	10

Fig. 17. Beam refinement gain with different UE models.

11 Related Work

Theoretical foundations of beamforming. A large number of theoretical contributions have been made in the topic of beamforming for 5G and beyond networks [16, 18, 21, 22, 24, 25, 36, 39, 40, 45, 48, 54, 56, 59, 63, 64, 66, 72, 78, 79]. The capabilities and potential of 5G networks, together with comprehensive overviews of beamforming, are extensively discussed [43, 45, 54, 56, 63, 66]. The works in [11, 15, 64, 72] study hybrid beamforming for mmWave systems, which combines the benefits of both analog and digital beamforming. However, operational 5G mmWave networks currently rely only on analog beamforming. The works in [6, 21, 22, 25, 40, 48, 78, 79] provide theoretical frameworks for mathematical optimization, analysis, and modeling of beamforming performance. All these works conduct simulation-based evaluation, which underscores the need for more empirical research to validate their findings in real-world environments.

Experimental works on mmWave beam management. A large number of research works have conducted experimental studies of beam management in [10, 60, 61, 69] or designed systems to improve beamforming in mmWave networks [26, 27, 34, 38, 41, 43, 50, 55, 65, 67, 68, 70, 77]. All these works, with the exception of [1, 42], focus on the 802.11ad/ay standards designed for indoor 60 GHz WLAN environments. 802.11ad/ay radios use different beam management techniques compared to 5G mmWave radios (e.g., they rely on omnidirectional reception that simplifies beam management at the cost of lower directionality gain) and indoor WLAN environments have very different propagation characteristics compared to outdoor cellular environments. Additionally, many of these works [10, 26, 27, 34, 38, 41, 43, 50, 55, 61, 67, 70, 77] have been evaluated on custom software-defined radio platforms. In contrast, our work performs the first large-scale empirical study of beam management in operational 5G mmWave networks.

5G measurements studies. Several measurement studies have analyzed various aspects of operational 5G networks, including coverage, performance, carrier aggregation, MIMO, and application QoE, over the past 5 years [7, 19, 23, 29, 31, 37, 44, 47, 49, 51, 52, 57, 58, 73, 75]. None of these works, with the exception of [23, 51] focuses on beam management. As we explained in §1, the works in [23, 51] perform small scale-studies limited in depth, in one city, with one smartphone model. In contrast, this work performs the first large scale study spanning 6 cities and two different smartphone models. Additionally, our work explores for the first time the performance of beam re-orientation, the beam switching overhead, and the interplay between beam selection and CA.

12 Discussion

12.1 Limitations of Our Study

Our study is subject to several practical constraints inherent to working with commercial cellular systems:

First, our experiments are conducted with limited visibility into the proprietary implementation details of commercial gNBs. Beam switching logic, re-orientation strategies, RA and CA algorithms are vendor-specific and not publicly documented. Our analysis therefore relies on UE-side observables and statistical patterns without direct access to gNB decision-making processes. As a result, we refrain from making definitive claims about the underlying algorithms or vendor-specific implementations.

Second, our control over system parameters is inherently asymmetric. While we carefully design UE-side experiments and control mobility patterns, we cannot configure gNB-side parameters, such as the number of supported beams, or disable specific features (e.g., beam switching, CA, or RA). The beam resolutions in our dataset—24, 32, 36, and 144 beams—are the ones available in commercial deployments; experiments with intermediate resolutions, such as 64 or 96 beams, are absent due to the lack of such configurations in commercial systems, not by design. Although inner

control over gNB-side parameters would definitely help establish more definitive performance trends, we believe our findings are still valuable to the research community.

Third, cellular networks and technologies are continuously evolving, with ongoing updates to gNB/UE software, firmware, and hardware, and deployment strategies. As such, our findings reflect general trends observed at the time of the study. While the core trade-offs and system interactions we highlight are likely to persist, specific metrics may change as the technology matures.

12.2 Implications and Future Research Directions

Our findings have several implications on future system design, research directions, and standardization efforts. Below, we outline three key takeaways that connect our empirical observations with broader engineering and methodological considerations.

First, our study shows that beam management is quite optimized in today's 5G mmWave networks in terms of performance, under both walking and driving mobility patterns, but user performance is much lower under driving compared to walking. Additionally, our experiments with different phone models show that further increasing the number of beams on the UE side does not always lead to better beam selection, but selecting suboptimal beams more often does not necessarily result in lower throughput. Together, these observations suggest that further attempts to optimize beam management may not be necessary, and future research efforts should focus instead on cross-layer interactions and the interplay between beam management and other MAC/PHY layer mechanisms (e.g., RA, CA, MIMO).

Second, we found that, although measurement reports are sent frequently, only a small fraction of them leads to beam switches at the gNB. This observation suggests that, although beam management is already optimized in terms of performance, there is still scope for improvement with respect to the signaling overhead. Future research efforts should focus on simplifying the reporting mechanism and reducing the frequency of beam measurement reports, e.g., by leveraging ML-based approaches [55], the UE's built-in sensors [3], or out-of-band information [62, 67] to predict blockage, changes in its moving direction, or the signal strength of neighboring beams.

Third, our experience highlights how limited visibility into gNB-side internals remains a fundamental barrier for rigorous research and replicable system analysis. Most beam management decisions are made based on proprietary vendor logic, which hinders efforts to fully understand or reproduce behavior across deployments. This motivates the importance of ongoing initiatives such as O-RAN, which aim to open up key interfaces and standardize telemetry. Broader adoption of such open frameworks will significantly enhance our ability to study, compare, and evolve beam management strategies in future cellular networks.

13 Conclusion

This work provides the first large-scale empirical evaluation of beam management in operational 5G mmWave networks. Our results show that beam management is quite optimized in terms of performance, selecting near-optimal beams most of the time, however, the reporting overhead can be high, especially when beam refinement is applied on top of coarse SSB beams. Our findings suggest that further attempts to optimize beam management performance alone may not be necessary. Instead, future research efforts should focus on reducing the beam management overhead or should consider beam management jointly with other components of the 5G protocol stack (e.g., signal processing, RA, CA, MIMO), which also have a major impact on performance.

Acknowledgments

We thank the anonymous reviewers and our shepherd Qing Wang for their helpful comments. This work was supported in part by NSF grant CNS-2340283.

References

- [1] Online. [CoNEXT '25] Vivisecting Beam Management in Operational 5G mmWave Networks. <https://github.com/NUWiNS/Beam-Management-in-Operational-5G-mmWave>.
- [2] Online. Requirements for support of radio resource management. Technical Report 38.133 version 15.6.0 Release 15. 3GPP.
- [3] 3GPP. 2018. TS38.306: NR; User Equipment (UE) radio access capabilities. 3rd Generation Partnership Project (3GPP), version 16.7.0.
- [4] 3GPP. 2020. TS38.213: NR; Physical layer procedures for control.
- [5] 3GPP. 2021. TS38.331: 5G NR: Radio Resource Control (RRC).
- [6] Accuver. 2024. XCAL Solo. <https://accuver.com/sub/products/view.php?idx=11>.
- [7] Shivang Aggarwal, Moinak Ghoshal, Piyali Banerjee, and Dimitrios Koutsonikolas. 2021. An Experimental Study of the Performance of IEEE 802.11ad in Smartphones. *IEEE Computer Communications* 169 (2021), 220–231.
- [8] Shivang Aggarwal, Srisai Karthik Neelamraju, Ajit Bhat, and Dimitrios Koutsonikolas. 2022. A Detailed Look at MIMO Performance in 60 GHz WLANs. *Proc. of ACM SIGMETRICS*
- [9] Shivang Aggarwal, Ujrit Satish Sardesai, Viral Sinha, and Dimitrios Koutsonikolas. 2020. LiBRA: learning-based link adaptation leveraging PHY layer information in 60 GHz WLANs. *Proc. of ACM CoNEXT*
- [10] Shivang Aggarwal, Ujrit Satish Sardesai, Viral Sinha, Deen Dayal Mohan, Moinak Ghoshal, and Dimitrios Koutsonikolas. 2020. An Experimental Study of Rate and Beam Adaptation in 60 GHz WLANs. *Proc. of ACM MSWiM*
- [11] Irfan Ahmed et al 2018. A Survey on Hybrid Beamforming Techniques in 5G: Architecture and System Model Perspectives. *IEEE Communications Survey & Tutorials* 20, 4 (2018), 3060–3097.
- [12] Ahmed Alkhateeb and Iz Beltagy. 2018. Machine Learning for Reliable mmWave Systems: Blockage Prediction and Proactive Handover. *Proc. of IEEE GlobalSIP*
- [13] Muhammad Alrabeiah, Andrew Hredzak, and Ahmed Alkhateeb. 2020. Millimeter Wave Base Stations with Cameras: Vision-Aided Beam and Blockage Prediction. *Proc. of IEEE VTC-Spring*
- [14] Je rey G. Andrews, Stefano Buzzi, Wan Choi, Stephen V. Hanly, Angel Lozano, Anthony C. K. Soong, and Jianzhong Charlie Zhang. 2014. What Will 5G Be? *IEEE Journal on Selected Areas in Communications* 32, 6 (2014), 1065–1082.
- [15] Omar El Ayach, Sridhar Rajagopal, Shadi Abu-Surra, Zhouyue Pi, and Robert W. Heath. 2014. Spatially Sparse Precoding in Millimeter Wave MIMO Systems. *IEEE Transactions on Wireless Communications* 13, 9 (2014), 1499–1513.
- [16] Tianyang Bai and Robert W. Heath. 2015. Coverage and Rate Analysis for Millimeter-Wave Cellular Networks. *Transactions on Wireless Communications* 14, 6 (2015), 1100–1114.
- [17] Omar Basit, Imran Khan, Moinak Ghoshal, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2025. 5G Metamorphosis: A Longitudinal Study of 5G Performance from the Beginning. *Proc. of ACM IMC*
- [18] Tadiio Endeshaw Bogale and Long Bao Le. 2016. Massive MIMO and MmWave for 5G Wireless HetNet: Potential Benefits and Challenges. *IEEE Vehicular Technology Magazine* 11, 1 (2016), 64–75.
- [19] Fukun Chen, Moinak Ghoshal, Enfu Nan, Phuc Dinh, Imran Khan, Z. Jonny Kong, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2025. A Large-Scale Study of the Potential of Multi-carrier Access in the 5G Era. *Proc. of PAM*
- [20] Phuc Dinh, Moinak Ghoshal, Dimitrios Koutsonikolas, and Joerg Widmer. 2022. Demystifying Resource Allocation Policies in Operational 5G mmWave Networks. *Proc. of IEEE WoWMoM*
- [21] Phuc Dinh, Tri Minh Nguyen, Chadi Assi, and Wessam Ajib. 2019. Joint Beamforming and Location Optimization for Cooperative Content-Aware UAVs. *Proceedings of IEEE WCNC*
- [22] 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.
- [23] Yufei Feng, Jin Wei, Phuc Dinh, Moinak Ghoshal, and Dimitrios Koutsonikolas. 2023. Beam Management in Operational 5G mmWave Networks. *Proc. of ACM mmNets*
- [24] Xinyu Gao, Linglong Dai, Shuangfeng Han, Chih-Lin I, and Robert W. Heath. 2016. Energy-Efficient Hybrid Analog and Digital Precoding for MmWave MIMO Systems With Large Antenna Arrays. *IEEE Journal on Selected Areas in Communications* 34, 4 (2016), 998–1009.
- [25] Zhen Gao, Linglong Dai, De Mi, Zhaocheng Wang, Muhammad Ali Imran, and Muhammad Zeeshan Shakir. 2015. MmWave Massive-MIMO-Based Wireless Backhaul for the 5G Ultra-Dense Network. *IEEE Wireless Communications* 22, 5 (2015), 13–21.
- [26] Yasaman Ghasempour, Muhammad Kumail Haider, Carlos Cordeiro, Dimitrios Koutsonikolas, and Edward Knightly. 2018. Multi-Stream Beam-Training for mmWave MIMO Networks. *Proc. of ACM MobiCom*
- [27] Yasaman Ghasempour and Edward Knightly. 2017. Decoupling Beam Steering and User Selection for Scaling Multi-User 60 GHz WLANs. *Proc. of ACM MobiHoc*
- [28] Moinak Ghoshal, Pranab Dash, Zhaoning Kong, Qiang Xu, Y. Charlie Hu, Dimitrios Koutsonikolas, and Yuanjie Li. 2022. Can 5G mmWave support Multi-User AR? *Proc. of PAM*

- [29] Moinak Ghoshal, Imran Khan, Z. Jonny Kong, Phuc Dinh, Jiayi meng, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2023. Performance of Cellular Networks on the Wheels. *Proc. of ACM IMC*
- [30] Moinak Ghoshal, Z. Jonny Kong, Qiang Xu, Zixiao Lu, Shivang Aggarwal, Imran Khan, Yuanjie Li, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2022. An In-Depth Study of Uplink Performance of 5G mmWave Networks. *Proc. of ACM SIGCOMM 5G-MeMU*
- [31] Moinak Ghoshal, Z. Jonny Kong, Qiang Xu, Zixiao Lu, Shivang Aggarwal, Imran Khan, Yuanjie Li, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2022. An In-Depth Study of Uplink Performance of 5G MmWave Networks. *Proc. of the ACM SIGCOMM 5G-MeMU Workshop*
- [32] Moinak Ghoshal, Z. Jonny Kong, Qiang Xu, Zixiao Lu, Shivang Aggarwal, Imran Khan, Jiayi Meng, Yuanjie Li, Y. Charlie Hu, and Dimitrios Koutsonikolas. 2023. Can 5G mmWave Enable Edge-Assisted Real-Time Object Detection for Augmented Reality?. *Proc. of MASCOTS*
- [33] Marco Giordani, Michele Polese, Arnab Roy, Douglas Castor, and Michele Zorzi. 2019. A Tutorial on Beam Management for 3GPP NR at mmWave Frequencies. *IEEE Communications Surveys & Tutorials* (2019), 173–196. Issue 1.
- [34] Muhammad Kumail Haider and Edward W. Knightly. 2016. Mobility resilience and overhead constrained adaptation in directional 60 GHz WLANs: protocol design and system implementation. *Proc. of ACM MobiCom* 10, 1 (2016), 61–70.
- [35] Bo Han, Yu Liu, and Feng Qian. 2020. ViVo: Visibility-Aware Mobile Volumetric Video Streaming. *Association for Computing Machinery*
- [36] 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.
- [37] Ahmad Hassan, Arvind Narayanan, Anlan Zhang, Wei Ye, Ruiyang Zhu, Shuwei Jin, Jason Carpenter, Z. Morley Mao, Feng Qian, and Zhi-Li Zhang. 2022. Vivisecting mobility management in 5G cellular networks. *Proc. of ACM SIGCOMM*
- [38] Haitham Hassanieh, Omid Abari, Michael Rodriguez, Mohammed Abdelghany, Dina Katabi, and Piotr Indyk. 2018. Fast Millimeter Wave Beam Alignment. *Proc. of ACM SIGCOMM*
- [39] Robert W. Heath, Nuria González-Prelcic, Sundeep Rangan, Wonil Roh, and Akbar M. Sayeed. 2016. An Overview of Signal Processing Techniques for Millimeter Wave MIMO Systems. *IEEE Journal of Selected Topics in Signal Processing* 10, 3 (2016), 436–453.
- [40] Sooyoung Hur, Taejoon Kim, David J. Love, James V. Krogmeier, Timothy A. Thomas, and Amitava Ghosh. 2013. Millimeter Wave Beamforming for Wireless Backhaul and Access in Small Cell Networks. *IEEE Transactions on Communications* 61, 10 (2013), 4391–4403.
- [41] Ish Kumar Jain, Raghav Subbaraman, and Dinesh Bharadia. 2021. Two beams are better than one: Towards Reliable and High Throughput mmWave Links. *Proc. of ACM SIGCOMM*
- [42] Ish Kumar Jain, Rohith Reddy Vennam, Raghav Subbaraman, and Dinesh Bharadia. 2023. mmFlexible: Flexible Directional Frequency Multiplexing for Multi-user mmWave Networks. *Proc. of IEEE INFOCOM*
- [43] Suraj Jog, Jiaming Wang, Junfeng Guan, Thomas Moon, Haitham Hassanieh, and Romit Roy Choudhury. 2019. Many-to-Many Beam Alignment in Millimeter Wave Networks. *Proc. of USENIX NSDI*
- [44] Rostand A. K. Fezeu, Claudio Fiandrino, Eman Ramadan, Jason Carpenter, Lilian Coelho de Freitas, Faaiz Bilal, Wei Ye, Joerg Widmer, Feng Qian, and Zhi-Li Zhang. 2024. Unveiling the 5G Mid-Band Landscape: From Network Deployment to Performance and Application QoE. *Proc. ACM SIGCOMM*
- [45] Shajahan Kutty and Debarati Sen. 2016. Beamforming for Millimeter Wave Communications: An Inclusive Survey. *IEEE Communications Surveys & Tutorials* 18, 2 (2016), 949–973.
- [46] Qianru Li, Zhehui Zhang, Yanbing Liu, Zhaowei Tan, Chunyi Peng, and Songwu Lu. 2023. CA++: Enhancing Carrier Aggregation Beyond 5G. *Proc. of ACM MobiCom*
- [47] Yang Li, Hao Lin, Zhenhua Li, Yunhao Liu, Feng Qian, Liangyi Gong, Xianlong Xin, and Tianyin Xu. 2021. A nationwide study on cellular reliability: measurement, analysis, and enhancement. *Proc. of ACM Sigcomm*
- [48] Cen Lin, Geoy 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.
- [49] Yanbing Liu and Chunyi Peng. 2023. A Close Look at 5G in the Wild: Unrealized Potentials and Implications. *Proc. of IEEE INFOCOM*
- [50] Sohrab Madani, Suraj Jog, Jesus O. Lacruz, Joerg Widmer, and Haitham Hassanieh. 2021. Practical null steering in millimeter wave networks. *Proc. of USENIX NSDI*
- [51] 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*
- [52] Arvind Narayanan, Xumiao Zhang, Ruiyang Zhu, Ahmad Hassan, Shuwei Jin, Xiao Zhu, Xiaoxuan Zhang, Denis Rybkin, Zhengxuan Yang, Zhuoqing Morley Mao, et al. 2021. A Variegated Look at 5G in the Wild: Performance, Power, and QoE Implications. *Proceedings of ACM SIGCOMM*

- [53] Sang-Hyun Park, Byoungnam Kim, Dong Ku Kim, Linglong Dai, Kai-Kit Wong, and Chan-Byoung Chae. 2023. Beam Squint in Ultra-Wideband mmWave Systems: RF Lens Array vs. Phase-Shifter-Based Array. *IEEE Wireless Communications* 30 (2023), 82–89. Issue 4.
- [54] Zhouyue Pi and Farooq Khan. 2011. An Introduction to Millimeter-Wave Mobile Broadband Systems. *IEEE Communications Magazine* 49, 6 (2011), 101–107.
- [55] Michele Polese, Francesco Restuccia, and Tommaso Melodia. 2021. DeepBeam: Deep Waveform Learning for Coordination-Free Beam Management in mmWave Networks. *Proc. of ACM MobiHoc*.
- [56] Sundeep Rangan, Theodore S. Rappaport, and Elza Erkip. 2014. Millimeter-Wave Cellular Wireless Networks: Potentials and Challenges. *Proc. IEEE* 102, 3 (2014), 366–385.
- [57] Muhammad Iqbal Rochman, Vanlin Sathya, Norlen Nunez, Damian Fernandez, Monisha Ghosh, Ahmed S. Ibrahim, and William Payne. 2022. A Comparison Study of Cellular Deployments in Chicago and Miami Using Apps on Smartphones. In *Proc. ACM WiNTECH*.
- [58] Muhammad Iqbal Rochman, Wei Ye, Zhi-Li Zhang, and Monisha Ghosh. 2024. A Comprehensive Real-World Evaluation of 5G Improvements over 4G in Low- and Mid-Bands. *Proc. IEEE DySPAN*.
- [59] Wonil Roh, Ji-Yun Seol, Jeongho Park, Byunghwan Lee, Jaekon Lee, Yungsoo Kim, Jaeweon Cho, Kyungwhoon Cheun, and Farshid Aryanfar. 2014. Millimeter-Wave Beamforming as an Enabling Technology for 5G Cellular Communications: Theoretical Feasibility and Prototype Results. *IEEE Communications Magazine* 52, 2 (2014), 106–113.
- [60] Swetank Kumar Saha, Shivang Aggarwal, Hany Assasa, Adrian Loch, Naveen Muralidhar Prakash, Roshan Shyamsunder Anantharamakrishna, Daniel Steinmetzer, Dimitrios Koutsonikolas, Joerg Widmer, and Matthias Hollick. 2021. Performance and Pitfalls of 60 GHz WLANs Based on Consumer-Grade Hardware. *IEEE Transactions on Mobile Computing* 20 (2021), 1543–1557. Issue 4.
- [61] Swetank Kumar Saha, Yasaman Ghasempour, Muhammad Kumail Haider, Tariq Siddiqui, Paulo De Melo, Neerad Somanchi, Luke Zakrajsek, Arjun Singh, Roshan Shyamsunder, Owen Torres, Daniel Uvaydov, Josep Miquel Jornet, Edward Knightly, Dimitrios Koutsonikolas, Dimitris Pados, Zhi Sun, and Ngwe Thawdar. 2019. X60: A Programmable Testbed for Wideband 60 GHz WLANs with Phased Arrays. *Elsevier Journal of Computer Communications (COMCOM)* 133 (2019), 77–88.
- [62] Batool Salehi, Mauro Belgiovine, Sara Garcia Sanchez, Jennifer Dy, Stratis Ioannidis, and Kaushik Chowdhury. 2020. Machine Learning on Camera Images for Fast mmWave Beamforming. *Proc. of IEEE MASS*.
- [63] 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* 59, 9 (2019), 132–137.
- [64] Maliheh Soleimani, Robert C. Elliott, Witold A. Krzymie, Jordan Melzer, and Pedram Mousavi. 2020. Hybrid Beamforming for MmWave Massive MIMO Systems Employing DFT-Assisted User Clustering. *IEEE Transactions on Vehicular Technology* 69, 10 (2020), 11646–11658.
- [65] Daniel Steinmetzer, Daniel Wegemer, Matthias Schulz, Joerg Widmer, and Matthias Hollick. 2017. Compressive Millimeter-Wave Sector Selection in O -the-Shelf. *IEEE 802.11ad Device Arch. of ACM CoNEXT*.
- [66] Shu Sun, Theodore S. Rappaport, Robert W. Heath, Andrew Nix, and Sundeep Rangan. 2014. MIMO for Millimeter-Wave Wireless Communications: Beamforming, Spatial Multiplexing, or Both? *IEEE Communications Magazine* 52, 12 (2014), 110–121.
- [67] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. 2017. WiFi-Assisted 60 GHz Wireless Networks. In *Proc. of ACM MobiCom*.
- [68] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. 2018. Towards Scalable and Ubiquitous Millimeter-Wave Wireless Networks. In *Proc. of ACM MobiCom*.
- [69] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, and Parameswaran Ramanathan. 2015. 60 GHz Indoor Networking through Flexible Beams: A Link-Level Protocol. In *Proc. of ACM SIGMETRICS*.
- [70] Sanjib Sur, Xinyu Zhang, Parameswaran Ramanathan, and Ranveer Chandra. 2016. BeamSpy: Enabling Robust 60 GHz Links Under Blockage. In *Proc. of USENIX NSDI*.
- [71] Telegraphic. 2023. A Look at the Current Status of 5G and Millimeter Waves at Home and Abroad [Part 1] Frequency Allocation and Commercial Service Launch. <https://www.telegraphic.jp/en/2023/12/19/>.
- [72] Kiran Venugopal, Ahmed Alkhateeb, Nuria González Prelcic, and Robert W. Heath. 2017. Channel Estimation for Hybrid Architecture-Based Wideband Millimeter Wave Systems. *IEEE Journal on Selected Areas in Communications* 35, 9 (2017), 1996–2009.
- [73] Dongzhu Xu, Anfu Zhou, Xinyu Zhang, Guixian Wang, Xi Liu, Congkai An, Yiming Shi, Liang Liu, and Huadong Ma. 2020. Understanding Operational 5G: A First Measurement Study on its Coverage, Performance and Energy Consumption. In *Proc. of ACM SIGCOMM*.
- [74] Xinlei Yang, Hao Lin, Zhenhua Li, Feng Qian, Xingyao Li, Zhiming He, Xudong Wu, Xianlong Wang, Yunhao Liu, Zhi Liao, Daqiang Hu, and Tianyin Xu. 2022. Mobile access bandwidth in practice: measurement, analysis, and implications. In *Proc. ACM SIGCOMM*.

- [75] Wei Ye, Xinyue Hu, Steven Sleder, Anlan Zhang, Udhaya Kumar Dayalan, Ahmad Hassan, Rostand A. K. Fezeu, Akshay Jajoo, Myungjin Lee, Eman Ramadan, Feng Qian, and Zhi-Li Zhang. 2024. Dissecting Carrier Aggregation in 5G Networks: Measurement, QoE Implications and Prediction. *Proc. ACM SIGCOMM*
- [76] Xumiao Zhang, Anlan Zhang, Jiachen Sun, Xiao Zhu, Yihua Guo, Feng Qian, and Z. Morley Mao. 2021. EMP: Edge-assisted Multi-vehicle Perception. *Proc. of ACM MobiCom*
- [77] Anfu Zhou, Xinyu Zhang, and Huadong Ma. 2017. Beam-forecast: Facilitating mobile 60 GHz networks via model-driven beam steering. In *Proc. IEEE INFOCOM*
- [78] 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.
- [79] Pei Zhou, Xuming Fang, Xianbin Wang, Yan Long, Rong He, and Xiao Han. 2019. Deep Learning-Based Beam Management and Interference Coordination in Dense MmWave Networks. *IEEE Transactions on Vehicular Technology*, 68, 1 (2019), 592–603.

(a) Boston. (b) Charlotte.

Fig. 19. Constructed angular coordinates for two gNBs from Verizon in Boston and Charlotte.

Appendix

A Ethics

This study was carried out by PhD students and faculty. We purchased multiple unlimited cellular data plans from two US carriers and our experiments comply with their customer agreements. This work does not raise any ethical concerns.

B Inferring Beamwidths

Throughout the paper, we analyze various beam management parameters among gNBs with different beam resolutions under the assumption that a higher beam resolution (larger number of available SSB beams) results in narrower beamwidths. Since information about the supported beamwidth of phased arrays from different vendors is not publicly available, in this section, we describe our systematic experiments and analysis to verify this assumption.

The key to our analysis is to build a mapping between angular positions relative to a gNB and the active SSB indices used in each position. If we observe the same index used across consecutive angular positions, we can obtain an estimate of how wide a given beam is.

Experiment procedures. We first created our own coordinate system for different types of gNBs as shown in Fig. 19. For each surveyed gNB, we divided the surrounding area into sectors. We denoted the middle of a gNB panel as θ_0 and recorded this direction with a high-accuracy RTK GPS. Using Google Earth Pro, we generated the coordinates using this reference. We then generated DL traffic for a total of 40 s with the user standing at a given sector and changing their orientation with respect to the gNB by 45° every 5 s. We repeated this process at each 10 sector, with 8 user orientations at each sector, while using XCAL to log the SSB indices used by the gNB to send data at that sector, enabling an analysis of the distribution of observed SSB indices across the defined sectors.

Beamwidth inference. Fig. 20 shows the different SSB indices recorded at various angles for two gNBs with different resolutions (144 beams and 36 beams) as an example. Recall that our assumption throughout this paper a gNB with a beam resolution of 144 beams has narrower beams than a gNB with a beam resolution of 24, 32, or 36 beams. Since many beams can be observed at the same angle for different user orientations, due to environmental factors such as blockage causing the signal to reach the UE over nLoS paths (reflections), we focus on the most commonly observed beams, annotated for each specific angle. We observe that, for a gNB with 36 beams (Fig. 20b), the same dominant SSB indices are recorded at 2-3 consecutive sectors indicating beamwidths

(a) 144-beam gNB. (b) 36-beam gNB.

Fig. 20. Distributions of SSB indices (denoted by different colors) recorded at different angles for two gNB types.

of 20° - 30° . For instance, SSB index 4 is observed from 160° - 180° . On the other hand, for a gNB with 144 beams (Fig. 20a), the same dominant beams are observed at a single sector or at most two consecutive sectors, indicating a beamwidth of 120° . Overall, the results indicate that a gNB with fewer available beams has to use the same SSB index (beam) to cover a wider range of angles. We also observe that a gNB with a resolution of 144 beams uses a much larger number of beams to cover a sector compared to a gNB with a resolution of 36 beams. The results for resolutions of 24 and 32 beams are similar to those with a resolution of 36 beams.

C Additional Experimental Details and Results

C.1 Trajectories of Controlled Walking Experiments

Fig. 21 shows the trajectories of the controlled walking experiments. In the case of walking towards the gNB, we started at P4 and stopped at P5. In the case of walking away from the gNB, we started at P5 and stopped at P4. For lateral motion, we started at P1, walked until we reached in front of the gNB at P3, and continued walking until P2.

C.2 Examples of Commercial 5G mmWave gNBs and Antenna Panels

Fig. 22 shows three examples of commercial 5G mmWave gNBs and antenna panels. A typical 5G mmWave gNB structure consists of 3 antenna panels, each covering a 120° sector. Figs. 22a, 22b) show two such examples of a Verizon and AT&T gNB, respectively, in Boston. The AT&T gNBs in Las Vegas have a different structure consisting of 4 panels, each covering a 90° sector, as shown in Fig. 22c.

Fig. 21. Trajectories for controlled walking experiments.

(a) Verizon, Boston. (b) AT&T, Boston. (c) AT&T, Las Vegas.

Fig. 22. gNBs and antenna panel examples.

(a) DL traffic. (b) UL traffic.

Fig. 23. Avg. # of unique Tx-Rx beam pairs included in beam measurement reports per run for different gNB beam resolutions and mobility patterns.

C.3 Additional Results

C.3.1 Number of Reported Beam Pairs. We explore the number of Tx-Rx beam pairs included in the UE beam reports during typical mobility patterns. N_g and N_u denote the total number of available beams on the gNB and UE, respectively, then there are a total of $N_g \times N_u$ beam pairs. With a gNB beam resolution of 24/32/36/144 beams and a UE beam resolution of 36 beams, this number is very high (864/1152/1296/5184 beam pairs). However, Fig. 23, shows that the average number of beam pairs included in UE reports during a run is very small, 35-117 on average in the DL direction and 51-164 on average in the UL direction, which account for 0.7-14% (1-19%) of the total beam pairs in the DL (UL) direction. The numbers of reported beam pairs are higher for the UL direction than the DL direction 80.5 vs. 53, 70 vs. 60, 77 vs. 71, and 150 vs. 103 for 24, 32, 36, and 144 gNB beams, respectively, which also aligns with our previous observations about UL vs. DL traffic. We also observe that the number of reported beam pairs for a given beam resolution can vary significantly for different mobility patterns, but the mobility pattern that yields the largest/smallest number of reported beam pairs is different for different gNBs resolutions, due to different implementations of 3D beamforming, as noted in §4.1.

C.3.2 Impact of Motion Type. Fig. 24 complements Fig. 12 comparing the number of beam changes per second under walking and driving mobility, when walking mobility does not include the "walking away" pattern. The results show lower beam change rates for walking mobility in this case compared to Fig. 12, although still higher median values compared to driving.

(a) gNB side. (b) UE side.

Fig. 24. Comparison of beam changes per second between walking and driving (w/o away mobility).

Fig. 25 further plots the beam change rate separately for each city. We observe that in Atlanta, where we only conducted random walking experiments, the beam change rate under walking is lower than under driving. In contrast, in Boston and Miami, the "walking away" mobility increases the overall beam change rate under walking.

C.4 Driving Trajectories

Fig. 26 shows our driving trajectories in 5 cities. Note that mmWave 5G coverage is intermittent due to either sparse deployments or building blockage. When mmWave coverage is unavailable, the connection typically falls back to lower-band 5G technologies or LTE.

Received December 2024; revised April 2025; accepted April 2025

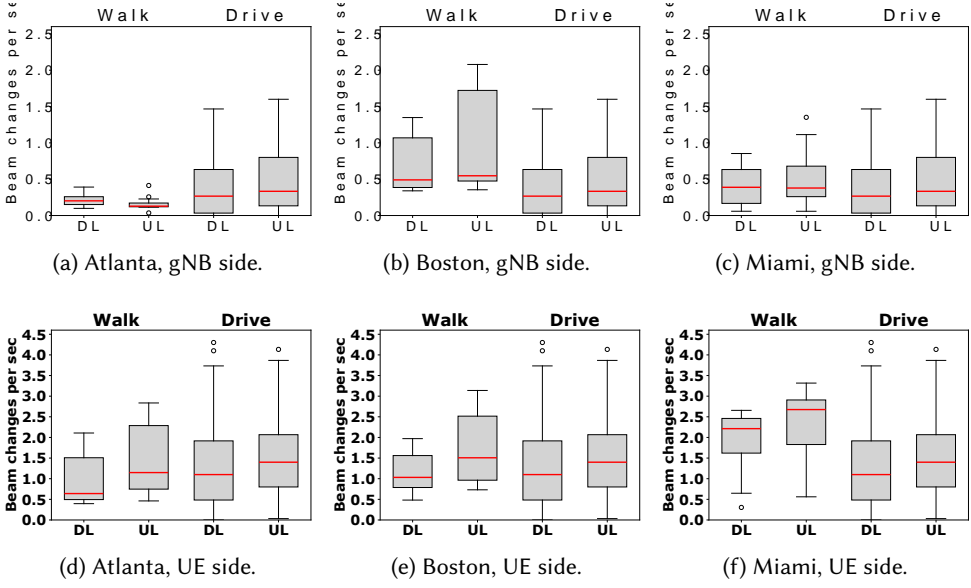


Fig. 25. Comparison of beam changes per second between walking and driving in 3 different cities.

(a) Boston driving trajectory. (b) Miami driving trajectory. (c) Las Vegas driving trajectory.

(d) Atlanta driving trajectory.

(e) Chicago driving trajectory.

Fig. 26. Driving trajectories in different cities. Red: 5G mmWave, blue: 5G low/mid.