

How Mature is 5G Deployment? A Cross-Sectional, Year-Long Study of 5G Uplink Performance

Original

How Mature is 5G Deployment? A Cross-Sectional, Year-Long Study of 5G Uplink Performance / Khan, I.; Ghoshal, M.; Angjo, J.; Dimce, S.; Hussain, M.; Parastar, P.; Yu, Y.; Fiandrino, C.; Orfanidis, C.; Aggarwal, S.; Aguiar, A.; Alay, O.; Chiasserini, C. F.; Dressler, F.; Hu, Y. C.; Ko, S. Y.; Koutsonikolas, D.; Widmer, J.. - ELETTRONICO. - (2024).
(Intervento presentato al convegno IFIP Networking 2024 tenutosi a Thessaloniki (Greece) nel June 2024).

Availability:

This version is available at: 11583/2987800 since: 2024-04-15T06:57:54Z

Publisher:

IFIP

Published

DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2024 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

How Mature is 5G Deployment? A Cross-Sectional, Year-Long Study of 5G Uplink Performance

Abstract—After a rapid deployment worldwide over the past few years, 5G is expected to have reached a mature deployment stage to provide measurable improvement of network performance and user experience over its predecessors. In this study, we aim to assess 5G deployment maturity via three metrics: (1) Does 5G performance remain stable over a long time span? (2) Does 5G provide better performance than its predecessor LTE? (3) Does the technology offer similar performance across diverse geographic areas and cellular operators? We answer this important question by conducting a cross-sectional, year-long measurement study of 5G uplink performance. Leveraging a custom Android App, we collected 5G uplink performance measurements (of critical importance to latency-critical apps) spanning 8 major cities in 7 countries and two different continents. Our measurements show that 5G deployment in major cities appears to have matured, with no major performance improvements observed over a one-year period, but 5G does not provide consistent, superior measurable performance over LTE, especially in terms of latency, and further there exist clear uneven 5G performance across the 8 cities. Our study suggests while 5G deployment appears to have stagnated, it is short of delivering its promised performance and user experience gain over its predecessor.

I. INTRODUCTION

The most recent generation of cellular networks, 5G, promises ultra-high bandwidth and ultra-low latency, far surpassing the performance of 4G LTE networks, via a combination of PHY layer innovations such as higher modulation schemes, beamforming, (massive) MIMO, and wider channels. Such high data rates, combined with low latency, hold the promise to finally support *latency-critical* applications such as Augment Reality (AR), Mixed Reality (XR), Connected Autonomous Vehicles (CAVs), cooperative sensing, and the Metaverse, which demand ultra high network bandwidth and low network latency to support offloading of compute-intensive tasks to the edge cloud.

5G rollout started in 2019 and the wide-scale deployment has been rapid and aggressively marketed by all mobile network operators. As such, after a rapid deployment worldwide over the past few years, it is highly anticipated that 5G has reached a deployment stage mature enough to significantly improve the performance of mobile networks and, more importantly, the user experience, in particular, when running the class of latency-critical apps that could not be supported by LTE.

To answer this question, there have been a number of measurement studies of 5G networks in recent years [1]–[7], [10], [12]–[20]. However, most of these studies have focused on measuring the 5G *downlink* performance while the *uplink* performance of 5G networks remain largely unknown. Understanding the 5G uplink performance is important, since

TABLE I: Overview of the collected data.

City (Country)	Operator	Tests	Duration	Cell IDs	Radius of Gyration (km)
Berlin (Germany)	Telekom	341	11/22-09/23	194	6.156
Turin (Italy)	TIM, WINDTRE	90	11/22-09/23	41	6.819
Oslo (Norway)	Telenor, Telia	1429	09/22-09/23	276	2.179
Porto (Portugal)	MEO	241	01/23-08/23	57	1.191
Madrid (Spain)	Vodafone	7096	10/22-09/23	525	8.734
Vancouver (Canada)	Bell, Shaw Comm.	561	11/22-09/23	206	14.516
Boston (USA)	ATT, Verizon, T-Mobile	328	07/22-04/23	93	8.71
Bay Area (USA)	T-Mobile	80	07/22-07/23	30	6.34
Total	-	10166	07/22-09/23	1422	-

most latency-critical "5G killer" apps distinguish themselves from legacy apps for their heavy, bursty *uplink* data transfers, and 5G, similar to all its predecessors, has provisioned much higher downlink bandwidth than uplink bandwidth.

In this work, we aim to fill this gap by answering two questions: (1) *How mature is today's 5G deployment?* and (2) *Is today's 5G uplink performance sufficient to enable latency-critical uplink-oriented apps such as AR or CAVs?* We consider that a technology deployment is "mature" when the following three conditions are satisfied: (i) *Its performance remains stable over a long time span.* Previous works performed measurements within a short time span, ranging from a few days up to a couple of months. However, any findings from such studies might be short-lived and lead to wrong conclusions about the potential of 5G in the long term. (ii) *The technology offers higher coverage and better performance than its predecessor.* In its mature stage, 5G should offer extended coverage replacing LTE and significantly higher throughput and lower latency than LTE, as promised. (iii) *The technology offers similar coverage and performance across diverse geographic areas and cellular operators* (in the same frequency band). Several previous works performed studies limited to one or a couple of cities or with a single operator. Such studies only provide a partial view of 5G performance, as hardware, configurations, and policies can differ not only across operators but also across cities for the same operator [2], [3]. Consequently, these two questions cannot be answered without a detailed, *longitudinal and cross-sectional* study of 5G uplink performance.

To answer these questions, in this work we conduct a cross-sectional, year-long measurement study of 5G uplink performance. Leveraging an IRB-approved custom Android app, we collected a large dataset of 5G performance (uplink TCP throughput and RTT) along with various metadata (cell IDs, handovers, GPS, coordinates, signal strength, mobility status,

etc.). Our dataset, summarized in Table I, spans 8 major cities in 7 different countries and two different continents – Berlin (Germany), Turin (Italy), Oslo (Norway), Porto (Portugal), and Madrid (Spain) in Europe, Boston (USA), Bay Area (USA), and Vancouver (Canada) in North America – and 12 operators. In each of these cities, volunteers used our app to perform weekly measurements at their convenience. As such, our dataset reflects the average performance experienced by a user at home, office, or during their regular commute over a whole year. *We plan to release the dataset upon acceptance of the paper.*

Leveraging this unique dataset, we first looked at the evolution of 5G performance in each city over the past year in terms of uplink throughput and latency. Somewhat surprisingly, we did not observe any increasing/decreasing trend for either metric, which suggests that condition (i) for maturity is satisfied; the technology deployment appears to have reached a mature stage and there are no major updates over the past one year.

We then looked at 5G performance (throughput and latency) in each city and compared it with the corresponding LTE performance. Surprisingly, our analysis reveals that 5G does not always yield better performance than LTE, suggesting that condition (ii) for maturity is not met. In particular, 5G throughput is lower than LTE throughput in 1 city and the 5G-LTE throughput gap across the remaining 7 cities varies significantly from 2.36 Mbps to 52.23 Mbps in the median case. More importantly, the 5G latency is lower than the LTE latency only in 3/8 cities and higher in 3/8 cities. Further, our dataset reveals very diverse 5G performance across the 8 cities, suggesting that condition (iii) for maturity is not met either. Overall, our study suggests while 5G deployment appears to have stagnated, it is short of delivering its promised performance gain over its predecessor and is not ready to support the next generation of latency-critical apps.

II. RELATED WORK

Since the initial 5G rollout in 2019, a large number of studies have measured various aspects of 5G performance [1]–[7], [10]–[20]. Most of them have a limited geographic coverage, conducting measurements in one [1], [5], [10], [11], [13]–[15], [19] or a few cities [2], [3], [7], [12], [17], [18]. While these studies provide valuable insights into 5G performance within localized areas, they may not necessarily reflect worldwide performance trends. Additionally, most of these studies conduct measurements over a limited time span, from a few days to a couple of weeks [2]–[4], [6], [7], [12]–[15], [17]–[19] and they do not investigate the evolution of 5G performance over extended periods. Finally, most of them (with the exception of [6], [7]) focus primarily on downlink performance.

A small number of studies conduct measurements over a larger number of geographic locations. The work in [6] performs a measurement study of 5G and LTE performance during a cross-country drive from Los Angeles, CA to Boston, MA. The work in [4] measures roaming performance in the EU while driving across four European countries. Finally, the work in [3] compares midband performance in the US and Europe via measurements in a total of 5 cities over multiple cellular

TABLE II: List of metrics collected with our Android app.

Metric	Description
GPS	User's City, Country
Network Type	5G (mmWave) /5G (sub-6 GHz)/ LTE
RSRP	Reference Signal Received Power
RSRQ	Reference Signal Received Quality
RSSI	Received Signal Strength Indicator
Cell-ID, EARFCN/ARFCN	Connected cell id and frequency
Operator	User's cellular operator

operators. However, these studies limit their measurement campaigns within a short time span of at most a few weeks and do not analyze performance evolution over time.

On the other hand, the works in [1], [5], [10], [16], [20] conduct studies over a longer time span, from several weeks up to two years. The work in [5] studies 5G performance on public transit systems over a 3-month period in Madrid and the work in [10] studies 5G NSA performance over a 7-week period in Rome. The work in [1] also studies 5G NSA performance in Rome in two different periods, December 2020 to January 2021 and March-May 2023 and shows a slight decrease in performance in 2023. The works in [16], [20] conduct large-scale measurement campaigns over a 2-year period. The work in [16] studies 5G across a whole country (UK) from a mobile operator's perspective, focusing on network usage and temporal evolution. The authors in [20] study mobile access bandwidth in China by collaborating with a commercial bandwidth testing app, and similar to the work in [1], show a decrease in both LTE and 5G throughput in 2023 compared to 2022. In our work, we do not observe any notable difference in performance during 2023. All these works focus on a single city or country.

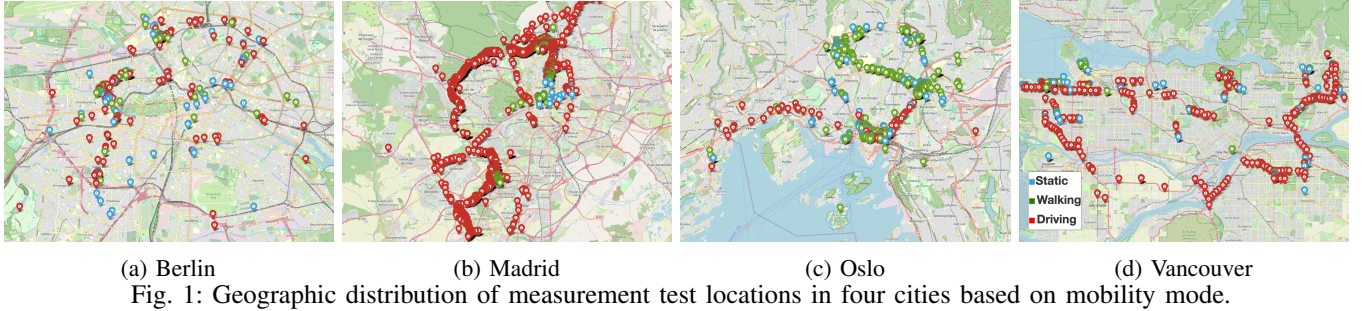
Our work, to our best knowledge, is the first to perform a longitudinal and cross-sectional measurement campaign of 5G performance, spanning 8 cities in 7 countries and 2 continents over a one-year period.

III. METHODOLOGY

Measurement servers. To enable throughput and latency measurements, we deployed three AWS Cloud servers, two in the US (Northern Virginia and Oregon) and one in Europe (Frankfurt, Germany). Additionally, for measurements in Boston with Verizon, we deployed an AWS Wavelength server in Boston. Wavelength servers are located *inside* Verizon's network in selected cities and specially designed for edge computing.

Measurement app. Our Android measurement app has two main functionalities. It measures uplink TCP throughput and RTT while collecting various cellular network metrics. We leverage Android-provided APIs to retrieve the network metrics that require the users to grant permission to access certain data on the phone (TELEPHONY, GPS, etc.). A detailed list of the collected metrics is shown in Table II.

The app initially collects the user's location in the background and selects the nearest server based on this information. Subsequently, the user is prompted to choose between three test types: static, walking, or driving. The app checks whether the UE's WiFi is turned off, exclusively focusing on cellular network performance. Once these checks are completed, the application measures uplink TCP throughput using nttcp-8.1.4



(a) Berlin (b) Madrid (c) Oslo (d) Vancouver
Fig. 1: Geographic distribution of measurement test locations in four cities based on mobility mode.

over a 10-second period. Following this, the app initiates an RTT test using the ping utility, sending 11 ICMP packets spaced 200 ms apart. The app features a lightweight design, with an image size of 6.5 MB and utilizing less than 250 MB of memory while running, ensuring efficient performance and minimal resource consumption on user devices.

Measurements. We reached out to our research community to recruit volunteers to participate in the measurement study for a one-year period. We received responses and data from 16 countries. However, we only present results from cities from which we have a sufficiently large number of measurement tests spanning several months. Our final dataset, summarized in Table I, consists of data from 8 cities in 7 different countries across Europe and North America, 1422 unique cell IDs, and 12 different operators.

In each city, one or two volunteers used our app to perform measurements with different mobility modes (static, walking, driving). Our dataset captures the average performance experienced by a user during their daily routine at home, office, or during their regular commute. The volunteers were asked to use all three mobility modes and perform at least a few measurements every week, however, they performed the tests at their convenience. As such, the total number of tests, their geographic spread (expressed as the radius of gyration [8]), and the number of tests for each mobility mode vary significantly across cities (see Table I). Fig. 1 shows the geographic distribution of measurement tests in the four cities with the largest number of measurements. In some cities, e.g., Berlin (Fig. 3a), the number of tests is roughly balanced across the three mobility modes; in others, e.g., we observe a dominant mobility mode, e.g., driving in Madrid (Fig. 3b) and Vancouver (Fig. 3d) or walking in Oslo (Fig. 3c).

IV. LONGITUDINAL STUDY

In this section, we explore the first condition for calling a technology mature, as defined in §I: does the 5G performance remain stable over a long time, without an increasing trend? To answer this question, we perform linear regression on the LTE and 5G throughput and latency values (averaged over each week) over time and show the results (slope and p-value) in Tables III and IV, respectively. We observe that the slopes for both technologies and both metrics are very close to 0 in all cities, indicating no increasing/decreasing trend of throughput and latency over the one-year period we consider in our study. Similarly, p values are typically (much) higher

TABLE III: Linear regression with time (Throughput).

City	slope		p-value	
	LTE	5G	LTE	5G
Berlin	5.3e-07	-5.6e-07	0.08	0.05
Turin	1.4e-06	1.3e-06	3e-4	4e-4
Oslo	6.6e-07	6.0e-07	0.008	0.04
Porto	-7.6e-07	9.9e-07	0.006	0.1
Madrid	1.77e-07	-3.2e-08	1.56e-13	0.44
Vancouver	-1.6e-07	6.6e-08	0.22	0.62
Boston	-6.4e-07	-1.4e-06	0.12	0.03
Bay Area	-4.1e-07	-7.7e-08	0.68	0.74

TABLE IV: Linear regression with time (Latency).

City	slope		p-value	
	LTE	5G	LTE	5G
Berlin	-1.5e-06	-9.8e-07	0.39	0.13
Turin	-1.6e-06	4.5e-07	0.09	0.22
Oslo	-4.3e-07	2.5e-07	0.7	3.9e-19
Porto	-8.6e-07	8.6e-07	0.01	0.18
Madrid	-1.9e-06	-3.9e-06	0.29	0.01
Vancouver	1e-05	3.9e-06	0.01	0.3
Boston	6.1e-06	4.4e-06	3.1e-12	2.6e-06
Bay Area	-1.7e-07	-1.2e-07	0.8	0.8

than 0.05 meaning that the throughput and latency do not show a statistically significant relationship with time. While this is expected for LTE (a mature technology), it is rather surprising for 5G four years after its initial rollout.

We further show examples of the evolution of 5G throughput over time for three cities – Madrid (the city with the largest number of tests), Oslo (the city with the second largest number of tests and p-value lower than 0.05), and Berlin (p-value 0.05) – in Fig. 2. For each city, we plot the average throughput per week over all the tests and over the dominant mobility mode – driving in Madrid and Berlin, and walking in Oslo. The plots confirm our conclusions from the linear regression study. While throughput can vary significantly from one week to the next, we observe no increasing trend.

Overall, our results show that the first condition for maturity is satisfied: *5G deployment appears to have reached a mature stage in major cities in Europe and North America with no major performance improvements over the past one year.*

V. CROSS-SECTIONAL STUDY

We now turn our attention to the remaining two conditions for maturity: does 5G offer higher coverage and better performance than LTE? Is the 5G coverage and performance similar across diverse geographic locations and operators? We study coverage in §V-A and performance in §V-B–§V-F.




Fig. 2: Evolution of 5G in terms of throughput.




Fig. 3: Geographic distribution of measurement test locations in four cities based on cellular technology.

TABLE V: Technology coverage, expressed as the fraction of the number of throughput/RTT samples over a particular technology out of the total number of samples.

City (Country)	Throughput				RTT				Mobility Mode (LTE / 5G)		
	LTE	5G			LTE	5G			Static	Walking	Driving
Berlin (Germany)	0.53	0.47			0.53	0.47			0.17/0.14	0.12/0.10	0.24/0.23
Turin (Italy)	0.68	0.32			0.66	0.34			0.36/0.14	0.29/0.16	0.02/0.03
Oslo (Norway)	0.36	0.64			0.31	0.69			0.1/0.16	0.24/0.45	0.02/0.03
Porto (Portugal)	0.18	0.82			0.16	0.84			0.27/0.05	0.06/0.30	0.07/0.25
Madrid (Spain)	0.53	0.47			0.55	0.45			0.01/0.01	0.03/0.12	0.51/0.32
Vancouver (Canada)	0.41	high -	mid 0.57	low 0.02	0.38	high -	mid 0.60	low 0.02	0.13/0.13	0.06/0.13	0.21/0.34
Boston (USA)	0.60	high 0.11	mid 0.26	low 0.03	0.62	high 0.24	mid 0.13	low 0.01	0.33/0.12	0.17/0.25	0.11/0.02
Bay Area (USA)	0.08	high -	mid 0.88	low 0.04	0.10	high -	mid 0.86	low 0.04	0.30/0.03	0.03/0.62	0.01/0.01
Total	0.52	0.48			0.52	0.48			0.08/0.06	0.08/0.17	0.37/0.24

A. 5G coverage

We calculate coverage for a particular technology as the fraction of throughput or RTT samples over that technology out of the total number of samples. Table V shows the results for each city as well as the overall results. We observe that the results are very similar with both metrics; hence, we focus on the throughput results in the remainder of this section.

Table V shows that *the overall 5G coverage is moderate*; in total, 52% of the throughput samples were collected while the UE was connected to a 5G cell. However, *coverage varies significantly across cities and operators*. The largest 5G coverage is observed in the Bay Area with T-Mobile (92%) and Porto with MEO (82%), and the lowest in Turin with TIM and WINDTRE combined (only 32%). Interestingly, the two US cities exhibit very different 5G coverage – 92% in the Bay Area with T-Mobile vs. 40% in Boston with all three major US operators combined.

We also break down the 5G coverage based on the frequency band in 5G-low, 5G-mid, and 5G-high (mmWave) using the Absolute Radio Frequency Channel Number (ARFCN),

recorded by our app. Unfortunately, the Android API that returns the ARFCN failed in all the tests conducted in Europe; hence, this information is only available for tests in North America. Nonetheless, 5G in Europe is primarily deployed in the midband (band n78) [9]. When we compare the North American locations, we observe almost exclusively 5G-midband in the Bay Area with T-Mobile and Vancouver with Bell and Shaw Comm., as these operators do not use mmWave. On the other hand, in Boston, we observe 11% of the 5G throughput samples and 24% of the RTT samples over 5G-high, mainly with Verizon and AT&T. This result is in sharp contrast with a recent study [6] that reported a significant 5G-low coverage, mainly with T-Mobile and AT&T during a cross-country drive, suggesting that 5G-low is mainly used in highways thanks to its longer coverage, while the mid and high bands are preferred in cities to provide high throughput.

We next look at the geographic coverage of the two technologies, focusing on the four cities from which we collected the largest number of measurements in Fig. 3. In Berlin, which has a balanced coverage for the two technologies (53% LTE,

47% 5G), interestingly, we observe a large aggregation of tests over 5G southwest of the city center, while most of the tests around the city center were done over LTE (Fig. 3a). In Madrid (Fig. 3b), with similar 5G coverage as Berlin, we observe two major areas of high 5G coverage and one area with mostly LTE coverage, but also areas with both technologies present. In contrast, in Oslo (Fig. 3c) and Vancouver (Fig. 3d), where 5G coverage is significantly higher compared to Berlin and Madrid (64% and 59%, respectively), we observe no area where LTE is the prevalent technology. In areas with both technologies present, we observe tests over different technologies at locations geographically very close to each other.

We also explore the relationship between 5G coverage and the geographic spread of the measurements in each city. Tables I and V show that the two cities with the shortest radius of gyration (Oslo and Porto) have the 2nd and 3rd highest 5G coverage among the 8 cities (82% and 64%, respectively). However, we also observe cities with similar radius of gyration (Berlin, Turin, Bay Area), where the 5G coverage varies significantly (from 32% to 88%). We also note that the city with the largest radius of gyration (Vancouver) has much higher 5G coverage (59%) than other cities with much smaller radius. Overall, we do not observe any clear relationship between 5G coverage and the the geographical spread of the measurements.

We finally explore the impact of the user’s mobility mode on coverage. Table V shows that the coverage for a given mobility mode typically follows the same trend as the overall coverage. The only exception is Madrid, where 5G coverage is higher than LTE coverage during walking but lower during driving. While the same is also true for Boston, 5G coverage is also much lower than LTE coverage in Boston for static scenarios, suggesting that the user speed is not a critical factor.

In summary, our results in this section show that conditions (ii) and (iii) are not satisfied with respect to coverage across the 8 cities in our study. *Users are still connected to LTE about 50% of the time on average and coverage is very different across different locations and operators*, ranging from an impressive 92% to a disappointing 32%.

B. Throughput

Fig. 4 plots the CDFs of uplink 5G and LTE throughput in each of the 8 cities. For the 3 cities in North America, we further break down the 5G throughput into 5G-low, 5G-mid, and 5G-high. We observe that *5G offers higher throughput than LTE in 7/8 cities*. However, *the median gain varies significantly across cities*, from 2.36 Mbps in the Bay Area to 52.23 Mbps in Oslo, showing that *four years after its initial rollout, 5G does not always deliver the high throughput gains it promised*. Interestingly, in these two cities, the maximum 5G throughput is similar to the LTE throughput. In all the other cities (with the exception of Turin), the maximum 5G throughput is higher than the maximum LTE throughput, typically by several tens of Mbps up to 100 Mbps.

Two exceptions are worth noting – Bay Area and Turin.¹ In

¹Note that these are the two locations with the smallest number of runs, and hence, the results may not be fully representative.

Bay Area, the location with the highest 5G coverage (92%), 5G throughput is largely similar to LTE throughput, although it exhibits a much longer tail, indicating that better coverage does not necessarily translate to better user experience. Even more surprisingly, in Turin, 5G offers lower throughput than LTE. After contacting the volunteer in Turin, we found out that initially they used WindTre with a 5G subscription of a maximum rate of 10 Mbps throughput, and later they switched to using TIM as an operator, with an unlimited subscription. While the rate limiting imposed by WindTre explains the lower 30% of the samples in Fig. 4b, the remaining samples also exhibit very low throughput values of at most 85 Mbps.

When we compare the three different 5G bands in North America, we observe that, as expected, 5G mmWave offers the highest throughput among the three bands, followed by 5G-mid and then by 5G-low. Interestingly, our small number of 5G-low samples exhibit lower median and maximum throughput than LTE in all three cities. Further, even though 5G midband is viewed as the band that offers the best tradeoff between range and performance, our results show that the gains over LTE in the uplink direction are quite low – 2.69 Mbps in the Bay Area, 3.91 Mbps in Vancouver, and 13.65 Mbps in Boston in the median case. Interestingly, in Boston, we observed a maximum 5G-mid throughput of 80 Mbps while the maximum LTE throughput exceeded 150 Mbps.

C. Latency

Fig. 5 plots the CDFs of uplink 5G and LTE latency in each of the 8 cities. For the 3 cities in North America, we further break down the 5G throughput into 5G-low, 5G-mid, and 5G-high. *Although 5G promises a significantly lower latency than LTE, our results in Fig. 5 surprisingly show that this is typically not the case*. 5G offers lower latency than LTE only in 3 out of 8 cities and the improvements are marginal. The median values for 5G vs. LTE latency in these three cities are – 46.1 ms vs. 49.9 ms in Oslo, 63.57 ms vs. 66.8 ms in Porto, and 34.33 ms vs. 40.85 ms in Vancouver. In the remaining 5 cities, the 5G latency is similar to or higher than the LTE latency. In Boston, latency is similar for the two technologies, although 5G offers lower best-case latency (25.38 ms vs. 34.46 ms at the 20-th percentile). In Madrid, 5G offers lower latency than LTE in the median case (55.22 ms vs. 59.6 ms) but significantly higher at the 80-th percentile (102.47 ms vs. 69.26 ms). In the Bay Area, latency is similar for the two technologies, but 5G has a much higher worst-case latency (e.g., 141.67 ms vs. 89.77 ms at the 90-th percentile). Finally, in Berlin and Turin, 5G latency is higher than LTE latency – 42.76 ms vs 30.94 ms and 57.49 ms vs. 47.41 ms in the median case, respectively. In fact, in Berlin, the upper quartile of the LTE latency is equal to lower quartile of the 5G latency.

We ran a few traceroute tests to the AWS Frankfurt server in Berlin (the city with the largest gap between 5G and LTE latency) over 5G and LTE and found that the path is the same over both technologies. This suggests that the root cause for the higher 5G latency lies in the RAN. We plan to further investigate this as part of our future work.

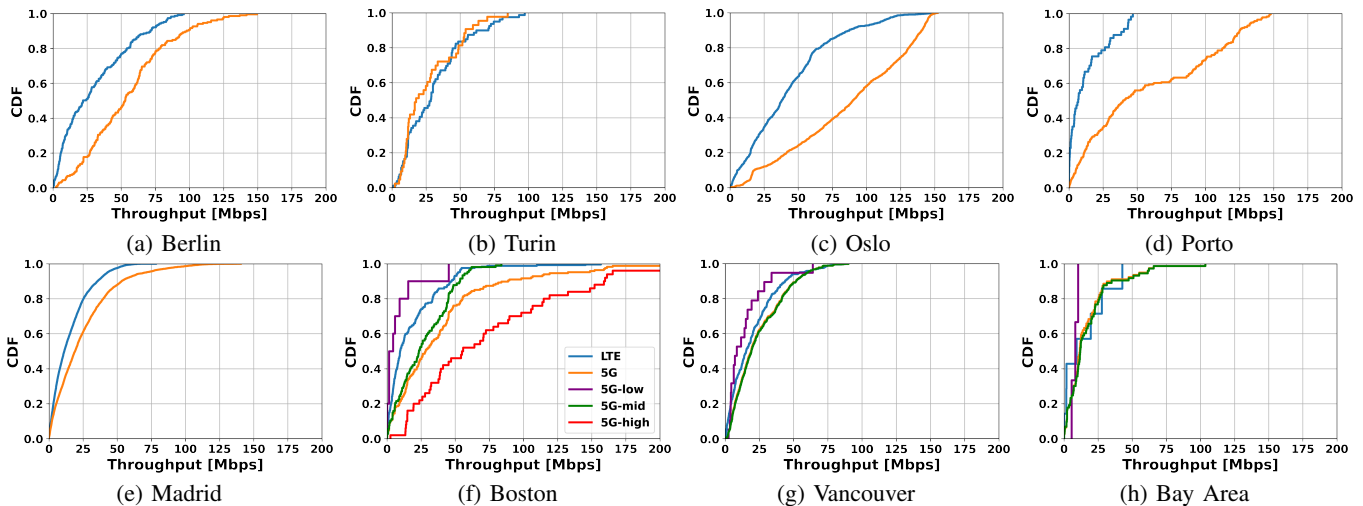


Fig. 4: Throughput comparison across different cities.

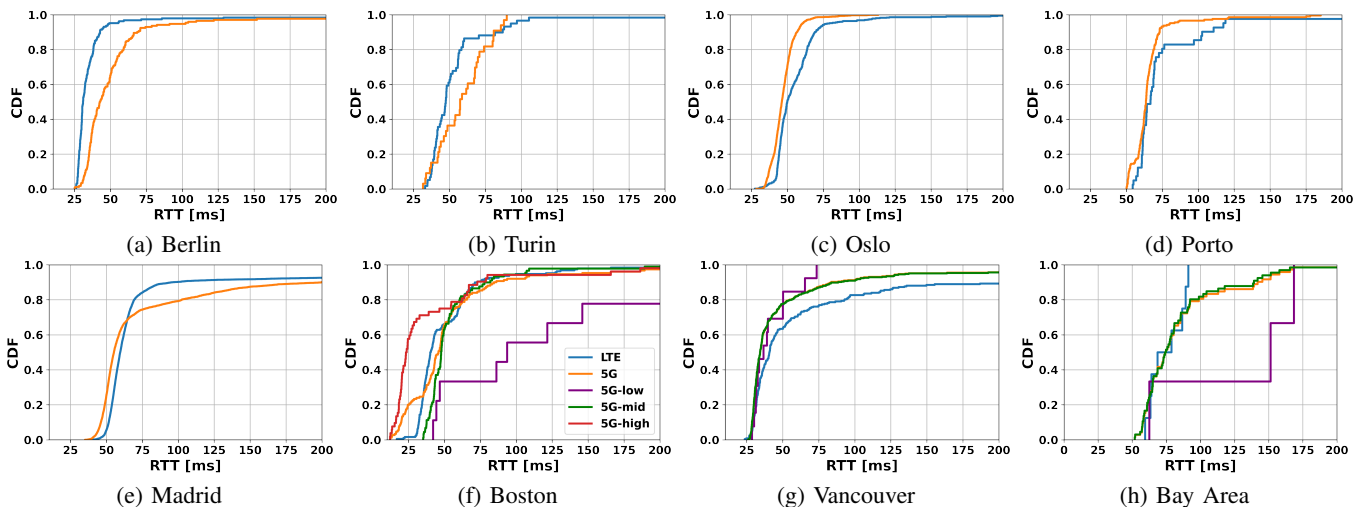


Fig. 5: Latency comparison across different cities.

When we compare the three different bands in North America, we observe that 5G-high in Boston over Verizon, combined with an edge AWS Wavelength server, offers significantly lower latency than all the other technologies and is responsible for the lowest 10-th percentile of the overall 5G latency in Boston in Fig. 5f. On the other hand, the 5G-low and 5G-mid latency is higher than the LTE latency in the two US locations but lower in Vancouver. In particular, the 5G-low latency is very high in Boston and Bay Area, but given the very small number of samples, it does not contribute significantly to the overall latency, which is mainly affected by the 5G-mid samples.

D. Impact of signal strength

In this section, we compare the signal strength of the two technologies and their correlation with performance. Fig. 6 plots the CDFs of the Reference Signal Received Power (RSRP) for 5G and LTE in each of the 8 cities. We observe that the RSRP is lower over 5G than over LTE in 7/8 cities; the gap varies from -4 dB (Porto) to -10 dB (Turin) in the median case. Boston and Vancouver are the only two exceptions. However, the impact of RSRP is different in throughput and latency across different cities.

Throughput. In Boston and Vancouver, the two cities where RSRP is higher over 5G than over LTE, the 5G throughput is also higher than the LTE throughput (Figs. 4f, 4g). Among the remaining 6 cities (Figs. 4a, 4b, 4c, 4d, 4e, 4h), 5G yields higher throughput than LTE in four of them (Berlin, Oslo, Porto, Madrid) but lower or similar in the other two (Turin, Bay Area). The availability of wider channel bandwidths in 5G NR than in LTE is the main reason for the overall higher throughput observed with 5G than with LTE in spite of the lower signal strength. 5G NR channel bandwidths of the operators under analysis are at least four times bigger than the maximum LTE channel bandwidth (i.e., 20 MHz). For example, previous measurement studies in Spain, France, Germany, and Italy show channel bandwidths in the range 80-100 MHz [3]. As operators try to allocate the maximum number of frequency resources per user with bulk transfers like our throughput experiments [3], the use of robust modulation schemes is sufficient to explain the reason behind the reported higher throughput with 5G despite a lower signal strength.

Latency. The higher 5G RSRP results in lower 5G latency in Vancouver (Fig. 5g), but only improves the worst-case 5G latency compared to the LTE latency in Boston (Fig. 5f).

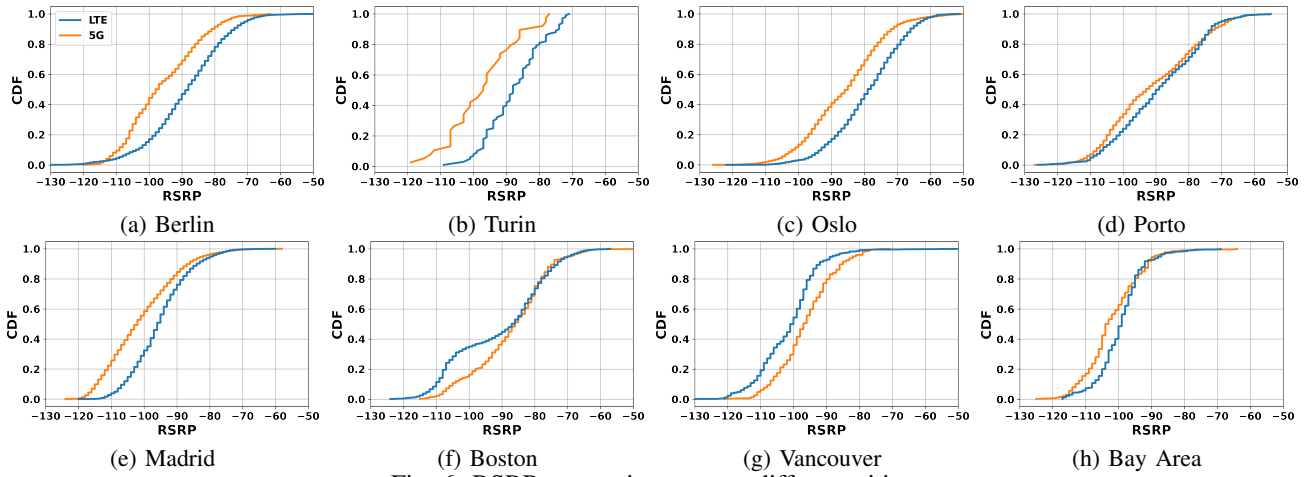


Fig. 6: RSRP comparison across different cities.

Note that in Boston (Fig. 6f) the 5G RSRP is higher than the LTE RSRP only for the lower half of the CDFs. Among the remaining 6 cities (Figs. 5a, 5b, 5c, 5d, 5e, 5h), the latency is lower over 5G than over LTE in two of them (Oslo, Porto), but similar or worse in the remaining four (Berlin, Turin, Madrid, Bay Area).

Overall, we observe that *RSRP has a weak correlation with performance but it appears to affect the latency more than the throughput*.

E. In-depth analysis of select cities

In this section, we analyze in depth the performance in three cities and explore the impact of mobility mode. We select Madrid and Vancouver, the two cities in Europe and North America, respectively, with the largest number of measurement tests, and Berlin as an example of a city with a good balance of tests with each mobility mode. Figs. 7 & 8 plot the technology-wise CDFs of throughput, latency, and RSRP, respectively, for each mobility mode in these three cities.

Berlin. Figs. 7a, 7b and 7c, 7d show that in Berlin both LTE and 5G exhibit the best performance (highest throughput and lowest latency) under walking. In contrast, the performance under static conditions is poor with both technologies and similar to that under driving, especially over 5G. Although in the previous section we concluded that RSRP alone cannot explain the performance difference between the two technologies, Figs. 8a, 8d show that RSRP can explain the performance for a given technology. These figures show that in Berlin, RSRP was high during walking tests and low during static and driving tests. Our volunteers in Berlin did the majority of the static tests indoors, which explains the low RSRP values and the low performance in static conditions.

Madrid. Figs. 7e, 7f and 7g, 7h show that in Madrid, static tests exhibit the worst performance over LTE but the best performance over 5G. Interestingly, driving exhibits the best performance over LTE but the worst over 5G. Walking also exhibits poor performance – worse than driving over LTE and similar to driving over 5G. However, the RSRP in Madrid is similar across all three mobility modes for each technology (Figs. 8b, 8e), and hence, it cannot explain the performance,

unlike in Berlin. Several 5G walking tests were run outdoors around the volunteer’s apartment building where there is a 5G a tower installation from a different operator (Orange) than the one used for the measurements (Vodafone). Since the two operators have a RAN sharing agreement, we conjecture that interference from the other operator is responsible for the low 5G performance in that area.

Vancouver. Figs. 7i, 7j and 7k, 7l show that in Vancouver, driving exhibits the worst throughput over both LTE and 5G and the worst latency over 5G, but surprisingly not over LTE. Static and driving tests, conducted outdoors in Vancouver, exhibit similar throughput, better than driving tests over both technologies. However, that latency is the best over 5G but the worst over LTE (in the median case). Fig. 8c shows that the LTE RSRP is similar for static and walking tests and much higher than for driving tests, which explains the throughput results but not the latency results. Fig. 8f shows that 5G RSRP was the lowest under static conditions (much lower than under walking), yet static tests exhibit the best latency and similar throughput to walking tests over 5G.

Overall, we observe that *cellular performance is the result of the complex interplay among a large number of factors and cannot be explained by looking individually at a single factor*. Previous works also arrived at a similar conclusion, showing a poor correlation of cellular throughput with RSRP [6], [11] and UE speed [6].

F. Overall performance across all cities

In the previous section, we focused on the comparison between 5G and LTE performance and showed that the second condition for maturity is not satisfied. In this section, we turn our attention to the third condition and compare the performance of a given technology across cities in Fig. 9.

Throughput. Fig. 9a shows that Oslo has the highest overall 5G throughput across the 8 cities, with a median/75-th percentile of 88/125 Mbps. Berlin comes second in terms of median throughput (52 Mbps vs. Porto’s 38 Mbps), but Porto has a much higher 75-th percentile (101 Mbps vs. 74 Mbps). On the other hand, Bay Area has the lowest 5G throughput among the 8 cities, both at the median and 75-th percentile (11.64 Mbps

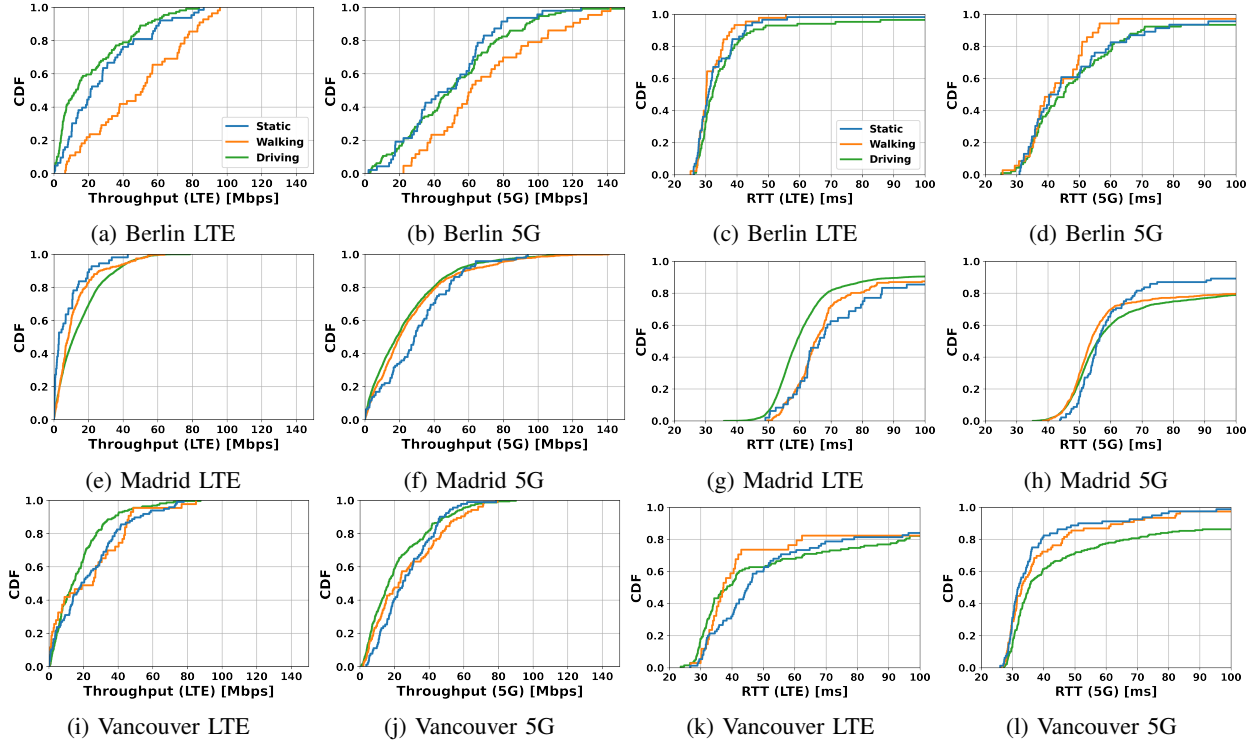


Fig. 7: Throughput & Latency Comparison across different mobility modes.

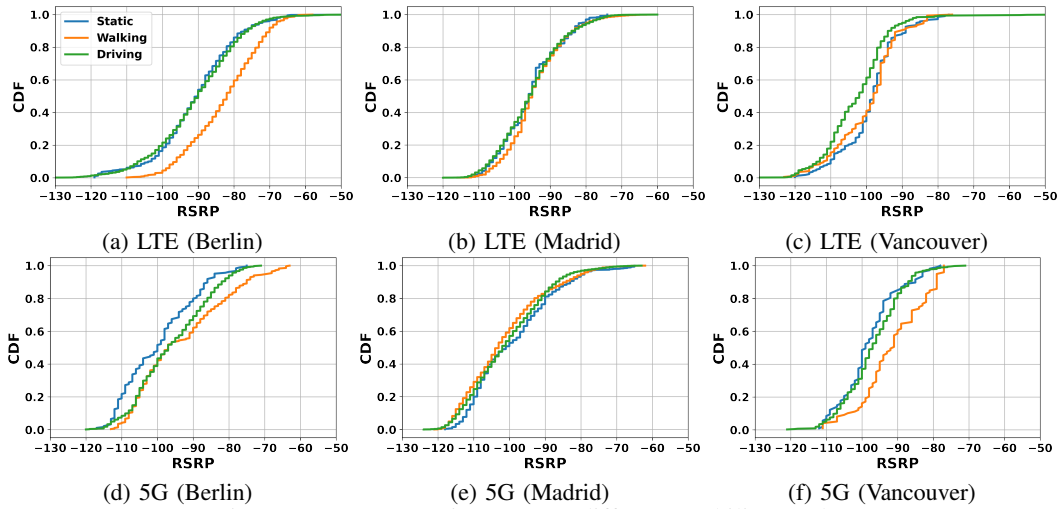


Fig. 8: RSRP comparison across different mobility modes.

and 22.64 Mbps, respectively). Note that Oslo’s lower quartile of 5G throughput is higher than the upper quartile of all cities except Berlin and Porto. Fig. 9a also shows that Oslo exhibits the highest LTE throughput with a median/75-th percentile of 40/59.18 Mbps, followed by Berlin and Turin. Interestingly, the median LTE throughput in Oslo matches the median 5G throughput in Porto and is higher than the 75-th percentile of the 5G throughput in Turin, Madrid, Vancouver, Boston, and Bay Area.

Overall, we observe a large disparity among the 5G throughput values across the 8 cities, suggesting that the third condition for maturity is not satisfied. We also observe a much larger spread of throughput values for 5G compared to LTE. Oslo and Porto, the two cities with the highest 75-

th percentiles also exhibit the largest IQR (73.84 Mbps and 87.06 Mbps, respectively). Note that these two cities have the lowest geographic sample spread, indicating that 5G throughput exhibits strong variations even in quite limited areas, and further reinforcing our conclusion that the third condition for maturity is not met yet.

Latency. A direct latency comparison among different cities is challenging, as the server location has a much higher impact on RTT than on throughput. For example, it is not surprising that Berlin exhibits the lowest median and lower quartile values for both 5G and LTE latencies among the 5 European cities in Fig. 9b, given that its distance to the Frankfurt AWS server we used for the measurements in Europe is the shortest. Yet, a few interesting observations are worth noting. First, Berlin’s




Fig. 9: Technology-wise comparison across different cities.

upper quartile for the 5G latency is higher than Oslo's, even though Oslo's distance from the Frankfurt AWS server is much longer. Note that Berlin is the city with the largest disparity between 5G and LTE latency. Second, in Boston 24% of the 5G RTT measurements were done over Verizon to an AWS Wavelength server located in the same city resulting in very low latency (notice the low whisker of the 5G boxplot in Boston in Fig. 9b), but for the remaining tests to an AWS server in North Virginia, the 5G latency is higher than in Vancouver, where the measurements were performed to a server located in Oregon. Third, we observe again a large disparity in the IQRs among different cities. Oslo and Porto, two cities with a large distance to the Frankfurt AWS server exhibit low IQRs for both 5G and LTE, suggesting the the latency is dominated by the wired network. On the other hand, for Madrid, which is also located far from the Frankfurt server, we observe a low IQR for LTE but not for 5G.

VI. CONCLUSION

In this paper, we conducted a cross-sectional, year-long measurement study of 5G aiming to assess its deployment maturity via three metrics: stability of its performance over a long time span, performance comparison with its predecessor LTE, and performance diversity in geographic locations and operators. Our measurements show that 5G deployment in major cities appears matured, with no major performance improvements observed over a one-year period, however, 5G uplink throughput often exhibits erratic and suboptimal behavior, and in some cases, is inferior to LTE. Further, 5G has not demonstrated significant improvements over LTE in terms of latency. Surprisingly, in certain cities worldwide, latency over LTE networks is comparable to or even lower than that of 5G. These findings suggest that while 5G holds promise for transformative enhancements in mobile networks, its full potential has yet to be realized.

REFERENCES

- [1] G. Caso, M. Rajiullah, K. Kousias, U. Ali, L. De Nardis, A. Brunstrom, O. Alay, M. Neri, and M.-G. Di Benedetto. An Initial Look into the Performance Evolution of 5G Non-Standalone Networks. In *Proc. of IFIP/IEEE TMA*, 2023.
- [2] P. Dinh, M. Ghoshal, D. Koutsonikolas, and J. Widmer. Demystifying Resource Allocation Policies in Operational 5G mmWave Networks. In *Proc. of IEEE WoWMoM*, 2022.
- [3] R. A. K. Fezeu, J. Carpenter, C. Fiandrino, E. Ramadan, W. Ye, J. Widmer, F. Qian, and Z.-L. Zhang. Mid-Band 5G: A Measurement Study in Europe and US. Technical Report 2310.11000 [cs.NI], arXiv, 2023.
- [4] R. A. K. Fezeu, C. Fiandrino, E. Ramadan, J. Carpenter, D. Chen, Y. Tan, F. Qian, J. Widmer, and Z.-L. Zhang. Roaming across the European Union in the 5G era: Performance, challenges, and opportunities. In *In Proc. of IEEE INFOCOM*, 2023.
- [5] C. Fiandrino, D. Juárez Martínez-Villanueva, and J. Widmer. Uncovers 5G Performance on Public Transit Systems with an App-based Measurement Study. In *Proc. of ACM MSWiM*, 2022.
- [6] M. Ghoshal, I. Khan, Z. J. Kong, P. Dinh, J. Meng, Y. C. Hu, and D. Koutsonikolas. Performance of Cellular Networks on the Wheels. In *Proc. of ACM IMC*, 2023.
- [7] M. Ghoshal, Z. J. Kong, Q. Xu, Z. Lu, S. Aggarwal, I. Khan, Y. Li, Y. C. Hu, and D. Koutsonikolas. An In-Depth Study of Uplink Performance of 5G MmWave Networks. In *Proc. of the ACM SIGCOMM 5G-MeMU Workshop*, 2022.
- [8] M. C. González, C. A. Hidalgo, and A.-L. Barabási. Understanding individual human mobility patterns. *Nature*, 453(7196), 2008.
- [9] R. S. P. GROUP. 5G developments and possible implications for 6G spectrum needs and guidance on the rollout of future wireless broadband networks. https://radio-spectrum-policy-group.ec.europa.eu/system/files/2023-10/RSPG23-040final-RSPG_Opinion_on_5G_developments_and_6G_spectrum_needs.pdf.
- [10] K. Kousias, M. Rajiullah, G. Caso, O. Alay, A. Brunstrom, L. D. Nardis, M. Neri, U. Ali, and M.-G. D. Benedetto. Coverage and Performance Analysis of 5G Non-Standalone Deployments. In *Proc. of ACM WiNTECH*, 2022.
- [11] J. Moreno, M. Contini, and A. Aguiar. 5G NSA Performance: A Measurement Study. In *Proc. of IFIP WONS*, 2024.
- [12] A. Narayanan, E. Ramadan, J. Carpenter, Q. Liu, Y. Liu, F. Qian, and Z.-L. Zhang. A first look at commercial 5G performance on smartphones. In *Proc. of ACM WWW*, 2020.
- [13] A. Narayanan, E. Ramadan, R. Mehta, X. Hu, Q. Liu, R. A. Fezeu, U. K. Dayalan, S. Verma, P. Ji, T. Li, et al. Lumos5G: Mapping and predicting commercial mmWave 5G throughput. In *Proc. of ACM IMC*, 2020.
- [14] A. Narayanan, M. I. Rochman, A. Hassan, B. S. Firmansyah, V. Sathya, M. Ghosh, F. Qian, and Z.-L. Zhang. A Comparative Measurement Study of Commercial 5G mmWave Deployments. In *In Proc. of IEEE INFOCOM*, 2022.
- [15] A. Narayanan, X. Zhang, R. Zhu, A. Hassan, S. Jin, X. Zhu, X. Zhang, D. Rybkin, Z. Yang, Z. M. Mao, F. Qian, and Z.-L. Zhang. A Variegated Look at 5G in the Wild: Performance, Power, and QoE Implications. In *Proc. of ACM SIGCOMM*, 2021.
- [16] P. Parastar, A. L. O. O. Alay, G. Caso, and D. Perino. Spotlight on 5G: Performance, Device Evolution and Challenges from a Mobile Operator Perspective. In *Proc. of IEEE INFOCOM*, 2023.
- [17] M. I. Rochman, V. Sathya, D. Fernandez, N. Nunez, A. S. Ibrahim, W. Payne, and M. Ghosh. A comprehensive analysis of the coverage and performance of 4G and 5G deployments. *Computer Networks*, 237, 2023.
- [18] M. I. Rochman, W. Ye, Z.-L. Zhang, and M. Ghosh. A Comprehensive Real-World Evaluation of 5G Improvements over 4G in Low- and Mid-Bands. Technical Report 2312.00957[cs.NI], 2023.
- [19] D. Xu, A. Zhou, X. Zhang, G. Wang, X. Liu, C. An, Y. Shi, L. Liu, and H. Ma. Understanding operational 5G: A first measurement study on its coverage, performance and energy consumption. In *Proc. of ACM SIGCOMM*, 2020.
- [20] X. Yang, H. Lin, Z. Li, F. Qian, X. Li, Z. He, X. Wu, X. Wang, Y. Liu, Z. Liao, et al. Mobile access bandwidth in practice: Measurement, analysis, and implications. In *Proc. of ACM SIGCOMM*, 2022.