Exploring Upper-6GHz and mmWave in Real-World 5G Networks: A Direct on-Field Comparison

Marcello Morini, Eugenio Moro, Ilario Filippini, Antonio Capone, Danilo De Donno

Abstract—The spectrum crunch challenge poses a vital threat to the progress of cellular networks and recently prompted the inclusion of millimeter wave (mmWave) and Upper 6GHz (U6G) in the 3GPP standards. These two bands promise to unlock a large portion of untapped spectrum, but the harsh propagation due to the increased carrier frequency might negatively impact the performance of urban Radio Access Network (RAN) deployments. Within the span of a year, two co-located 5G networks operating in these frequency bands were deployed at Politecnico di Milano, Milan, Italy, entirely dedicated to the dense urban performance assessment of the two systems. This paper presents an in-depth analysis of the measurement campaigns conducted on them, with the U6G campaign representing the first of its kind. A benchmark is provided by ray-tracing simulations. The results suggest that networks operating in these frequency bands provide good indoor and outdoor coverage and throughput in urban scenarios, even when deployed in the macro base station setup common to lower frequencies. In addition, a comparative performance analysis of these two key technologies is provided, offering insights on their relative strengths, weaknesses and improvement margins and informing on which bands is better suited for urban macro coverage.

Index Terms—5G, measurements, millimeter-wave, upper-6GHz, commercial deployment, Milan.

I. INTRODUCTION

The ongoing rise in the number of mobile users and their requirements for bandwidth makes the exploration of new spectrum necessary [1]. The overcrowding of lower frequency bands, where the propagation is good and the technology is familiar, led to what is known as *spectrum crunch*, which jeopardizes the future performances of mobile radio networks [2]. Two bands in particular attracted attention as a prompt way out to mitigate such shortage in the next years, namely the 6-GHz and the mmWave bands.

The frequency range from 5.925 to 7.125 GHz, known as the 6 GHz band, possesses good coverage features that characterize the mid-band spectrum. It provides sufficient bandwidths to ensure high-speed data transfer — letting users reach gigabits per second throughput — without the need to resort to higher frequencies. These qualities led to the introduction of the Upper 6GHz (U6G) among New Radio (NR) bands in 3rd Generation Partnership Project (3GPP) Release 17 and, in parallel, to the birth of Wi-Fi 6E within the Wi-Fi Alliance. However, the different and incompatible types of spectrum access that these two standards require (i.e., licensed and unlicensed, respectively), brought the 6-GHz licensing process

M. Morini, E. Moro, I. Filippini and A. Capone were with DEIB, Politecnico di Milano, Milan, Italy. {*name.surname*}@*polimi.it*

D. De Donno was with Milan Research Center, Huawei Technologies Italia S.r.l., Milan, Italy. *danilo.dedonno@huawei.com*

on a winding path. In April 2020, the United States Federal Communications Commission (FCC) enabled the use of the entire 1.200 MHz spectrum for unlicensed use [3], particularly for low-power indoor applications under an Automatic Frequency Coordination (AFC) framework. In China, the Ministry of Industry and Information Technology started supporting licensing policies at the end of June 2023, when it officially endorsed the U6G (or portions thereof) for licensed systems[1]. For Europe, Africa, and part of Asia, the discussion had a turning point in December 2023, at the World Radiocommunication Conference 2023. After years of technical analysis and discussions, the International Telecommunication Union decided to split the 6-GHz band into lower (5.945–6.425 GHz) and upper part (6.425–7.125 GHz) and to allocate only the latter for licensed use [4]. Europe is expected to adopt this decision according to the Radio Spectrum Policy Group (i.e., the high advisory body of the European Commission) opinion on World Radiocommunication Conference (WRC) [5].

MmWave represents another asset for future mobile networks. Firstly embedded in mobile access networks in 2012 to support high throughputs in Wi-Fi WiGig networks, mmWave were then brought into 5G starting in 2017 with Release 15. Frequencies between 24.25 GHz and 71 GHz are currently supported by 3GPP 5G standards as well as IEEE 802.11ad, aj, and ay. According to 5G standards, the total bandwidth in Frequency Range 2 (FR2) is around 29 GHz, which is six times more than that available in Frequency Range 1 (FR1) [6], [7]. This free spectrum real estate is meant to be one of the key assets to reach the goals of IMT-2020 [8]. Its role is vital for Enhanced Mobile Broadband (eMBB) application, in particular in hotspots and dense urban scenarios. However, its adoption has not reached the expected scale so far. To December 2023, only a minority of the 146000 5G base stations[2] deployed around the world use FR2 [9].

High-frequency bands in general, suffer from harsher propagation than low and mid bands. However, when put into the context of mobile radio networks, some encouraging signals should be considered. Free-space path loss increases with frequency, but smaller wavelengths facilitate integrating very directive Multiple Input, Multiple Output (MIMO) antennas, eventually generating higher gain and narrow beams that offset the path loss. Atmospheric gaseous absorptions oscillate from a minimum of 0.1 dB/km at 35HGz to a maximum of 10 dB/km at 60GHz due to a peak in oxygen absorption [10].

Still, these values should not be too alarming, considering that today's cell radius in urban environments is around 200m [1]. The same reasoning can be extended for rain attenuation: while it is true that at 26 GHz, heavy rain (25mm/h) attenuates around 4dB/km [10], this value does not heavily affect the connection on cell-sized distances. Certainly, the power originated by diffraction is smaller at mmWave than at lower frequencies. This per-se limits Non-Line of Sight (NLoS) propagation. However, as shown in the following, the effect might not be so severe in urban environments. Outdoor-to-indoor propagation is perhaps the most severe scenario, and high throughput connections can only be established in the presence of clear, thin glass [1], [11], as it will be shown in this work.

We believe that, as in any engineering discussion, there is a need for experiments and results to understand the benefits and challenges of the alternatives under investigation. Pursuing this objective, we deployed two standard-compliant 5G networks working in the U6G and mmWave bands, both configured as macro base stations. The co-location of the two deployments at the exact same point in the city of Milan, Italy, also allows a direct and meaningful comparison. Both networks present distinctive features. At the time of writing (Dec. 23), no other contributions on U6G 5G deployments are present, making this the first measurement campaign available in the literature at such frequencies [12]. Furthermore, our mmWave network stands out for not relying on anchoring base stations at lower frequencies (i.e., fallback to sub-6GHz connection), thus providing a clear and unbiased perspective on this band's performances. Finally, the comparative analysis enabled by the co-location of the two systems reveals further insights on the ideal application scenarios and the performance improvement margins.

The remainder of this paper is organized as follows. *Section II* introduces the measurement scenario and characterizes it with ray-tracing simulations. *Section III* and *IV* present the results of the U6G and mmWave measurement campaigns, respectively, which are then compared in *Section V*. The related works are reported in *Section VI*, just before concluding the paper in *Section VII*.

II. Measurement scenarios and equipment

The Upper 6GHz and millimeter-wave networks were mounted on the rooftop of one building in the main campus of Politecnico di Milano, in Milan, Italy. In this section, the common aspects between the deployments are stated, followed by the ray-tracing simulation results. Deployments' specific features are then outlined and the measurement methodology is eventually reported.

Both networks are made of three main components: an Active Antenna Unit (AAU), a Base Band Unit (BBU) and a 5G Core Network (CN) deployment. The AAU was installed at a height of 22 meters, down-tilted by 2°, and covering a 120° sector with a center azimuth of 135°. The spot and the orientation were chosen exactly equal to enable a meaningful comparison. The selected site location, along with the main hardware pieces composing the testbed, are depicted in *Figure*

1. Equipment details are reported in *Tab. I* and commented in the remaining of this section.

In both deployments, the BBUs and AAUs had been connected by a 25Gbps fiber fronthaul and mounted on the rooftop. The BBUs was then connected to the CN using a 10 Gbps backhaul link.

On the user side, the two campaigns leveraged different equipment, which are described below. Each User Equipment (UE) was mounted on a holder, connected to a server for traffic generation and data extraction, and powered by a battery. All these pieces of equipment were placed inside a cart and carried around the campus to perform cell coverage measurements.

The usage of a private core network allowed us to ease practical issues usually encountered by similar contributions. We used a set of ad-hoc SIM cards that could access the network, removing the need to purchase multiple network subscriptions. Moreover, while most of the related works rely on a speed test website, we installed a speed test tool directly inside the core network. Using remote servers to perform speed tests is a matter of concern because other bottlenecks along the network might alter the results. With our deployment, this obstacle is cleared and the only element impacting our measures is the wireless link.

This setup is considered *optimal* since there are no interferents in the area and only one UE can connect to the network. In such conditions, the signal strength is directly related to the achievable Modulation and Coding Scheme (MCS) and throughput. Otherwise, the management of more than one user would have originated non-linear relationships between these indicators [13]. Moreover, the choice to deploy a macro base station instead of other architectures makes the outcomes of this campaign very effective for mobile radio network operators since it is the most common network roll-out option and allows co-location with existing facilities. Milan, as the location for the testbed, is especially fitting to examine urban environments. Parks, trees, and relatively tall buildings characterize the neighborhood and allowed us to study several propagation scenarios, such as the conventional *Line of Sight* and *Non Line of Sight*, urban canyons, and Outdoor-to-Indoor propagation.

A. Ray-tracing simulations

Before starting the experimental campaign, we acquired the 3D digital maps of the area to perform ray-tracing simulations with the double objective of theoretically verifying the expected coverage as well as driving the selection of interesting positions for the subsequent measurement phase. We used the S_5GChannel tool by Siradel [14] powered by Volcano Urban ray-based model that is able to simulate multiple propagation paths from reflections, diffractions, transmissions and scattering with the objects described by raster and 3D vector layers. In particular, in the software, we placed and oriented the AAU antenna as in the real testbed and used the realistic radiation pattern of traffic beams to perform ray launching. We carried out point-to-area predictions by placing outdoor measurement points (i.e., user equipment locations) on a bi-dimensional regular grid with a resolution of 5 meters

Fig. 1: Aerial view of the test area with equipment details

(a) U6G simulation　　　(b) mmWave simulation

Fig. 2: Ray tracing prediction based on 3D vector maps of the Politecnico di Milano premises

along both x and y-axis. In such measurement points, we assumed UEs with omnidirectional antenna at 1.5m height.

The results of ray tracing predictions in terms of the Reference Signal Received Power (RSRP) heatmaps are reported in *Figure 2a* and *2b* for U6G and mmWave, respectively. *Figure 2a* shows RSRP peaks of -55dBm in the running track placed in perfect Line of Sight (LoS) with the base station, which gently degrades entering in the urbanized blocks, up to around $600m$ from the base station. At this distance, the RSRP decreases to -115dBm, considered the minimum power that can be received. Only limited portions of the area are in an outage, mostly placed at the cell edge where buildings shadow the measurement point.

Figure 2b follows a behavior similar to the U6G simulation but exhibits lower RSRP values and a more steeped decrease in NLoS conditions. The maximum RSRP achieved attests to -60dBm but sharply drops to less than -100dBm when the signal encounters buildings. As a result, the cell radius is smaller. However, slightly better coverage is expected in the real implementation of mmWave network thanks to the gain of the UE antenna we used.[3]

[3]We did not consider this effect in our ray-tracing simulations due to the challenges and uncertainties in properly modeling handset antennas.

TABLE I: Hardware specifications and 5G NR parameters

Parameter	U6G	mmWave
AAU coordinates (lat,lon,h)	45.478671, 9.232550, 22 m	
AAU azimuth, down tilt	135°, 2°	
Center frequency	6.8 GHz (n104)	27.2 GHz (n257)
Channel bandwidth	80 MHz	200 MHz
Subcarrier spacing	30 kHz	120 kHz
Frame structure	TDD 4:1 (DDDSU)	
Max QAM order (D/U)	256/64	
AAU TX power	37 dBm	37.5 dBm
AAU gain	33 dBi	32.5 dBi
AAU EIRP	70 dBm	70 dBm
AAU MIMO	128T-128R	8T-8R
UE gain	0 dBi (isotropic)	20 dBi
UE EIRP	22 dBm	45 dBm
UE MIMO	2T-4R	2T-2R

B. Upper 6GHz equipment

The U6G network equipment stands out for being a prototype in some of its parts. Specifically, the AAU and the Test UE (TUE), both standard-compliant, were made specifically for a U6G demonstration. On the other hand, the BBU is a commercially-available product. The base station works at a center frequency of 6.8 GHz, on a band of 80 MHz. The AAU has a gain of 33 dBm and it is equipped with 128 elements for both the transmission and the reception chains. The TUE was made of an omnidirectional antenna housed in a commercial smartphone chassis. The antenna was equipped with 4 elements in reception and 2 in transmission, inherently favoring downlink. The radio-frequency front-end was connected to a baseband unit processor, a server for traffic generation, and powered by a generator. All these pieces of equipment were placed inside a cart to reproduce a mobile user. More details are reported in *Table I*.

C. Millimeter-wave equipment

The mmWave network is fully standard-compliant and commercially available. The RAN consists of one Huawei HAAU5323, interfaced through a 25Gbps eCPRI fiber fronthaul to the baseband unit BBU5900 from the same vendor. The BBU is then connected to the virtualized 5G CN. A commercial Customer Premises Equipment (CPE) was used as a mobile terminal. More details regarding the mmWave hardware are reported in *Table I*. The CPE was mounted on a holder in the cart, powered by a power bank with AC output, and connected to a laptop. The laptop was extracting the measurements seen by the CPE through a drive test log software (Keysight's NEMO) that can access the information available in the CPE's chipset. Differently from the U6G TUE, CPE are directional, giving angular resolution to this campaign.

D. Measurement methodology

The measurement methodology is shared between the two campaigns. To carry out the tests, the User Equipment was brought to a measurement point, and two speed tests were launched (both in DL and UL, subsequently). Measurements of the aforementioned Key Performance Indicator were captured in 15-second windows. The communication data, ranging from the application to the physical layer, were captured and recorded by the log software. The antenna radiation pattern dictates one major difference in the data collection. With the U6G UE antenna, there was no need to test the TUE over different orientations, given its omnidirectional pattern. Oppositely, the mmWave CPE antenna is directive, so a full capture of one point requires orienting the antenna in more directions. In particular, we chose to point them toward the four cardinal points. In the results figure, these directions are indicated by an arrow. The directional pattern increased the angular resolution but required repeating the speed test procedure for each direction.

III. 6GHz

In the following, we will present an assessment of the upper 6 GHz 5G NR deployment's performance in both outdoor and indoor environments. Our evaluation encompasses the following metrics: RSRP, uplink/downlink throughput (measured at the application layer), and statistics related to the active Modulation and Coding Scheme (MCS).

A. Outdoor-to-outdoor performance

We start our analysis by showcasing the downlink RSRP measurements obtained by the TUE within the outdoor pedestrian zone of the testing area. *Figure 3a* visualizes these measurements superimposed on an aerial view of the testing region. Notably, we can observe that test points benefiting from an unobstructed LoS connection with the Base Station (BS) experience the highest RSRP values, aligning with our expectations. Interestingly, NLoS conditions do not necessarily translate to diminished RSRP values. Relatively robust RSRP measurements were recorded even at test locations where a single building partially obscured the LoS. However, with multiple buildings blocking the LoS, performance decreases rapidly starting from 300 m from the BS, as for *points C* and *D*. On the map, below *point C* we can observe a *street-canyon* effect that increases the received signal strength to relatively high values even with multiple blockages and up to the cell edge.

The RSRP values translate in the downlink peak throughput shown in *Fig. 3b*. Here it is shown how, even where the signal strength is at its lowest, the peak supported throughput remains well above 200 Mbps. In *Fig. 4a* a quantitative analysis is given through the cumulative distribution functions of measured RSRP and downlink throughput. The results illustrate that even with an RSRP as low as -110 dBm, a common value found at the cell edge, a peak downlink throughput of approximately 330 Mbps is achievable. It's important to note that these values were only recorded at fewer than 5% of the test points, while more than 50% of the test points attain peak throughputs exceeding 800 Mbps.

To discuss uplink performance, consider *Fig. 4b*, where the empirical Cumulative Distribution Function (CDF) for the measured uplink throughput is presented. In this context, it is evident that uplink throughput values are generally lower in comparison to the downlink. This discrepancy was anticipated
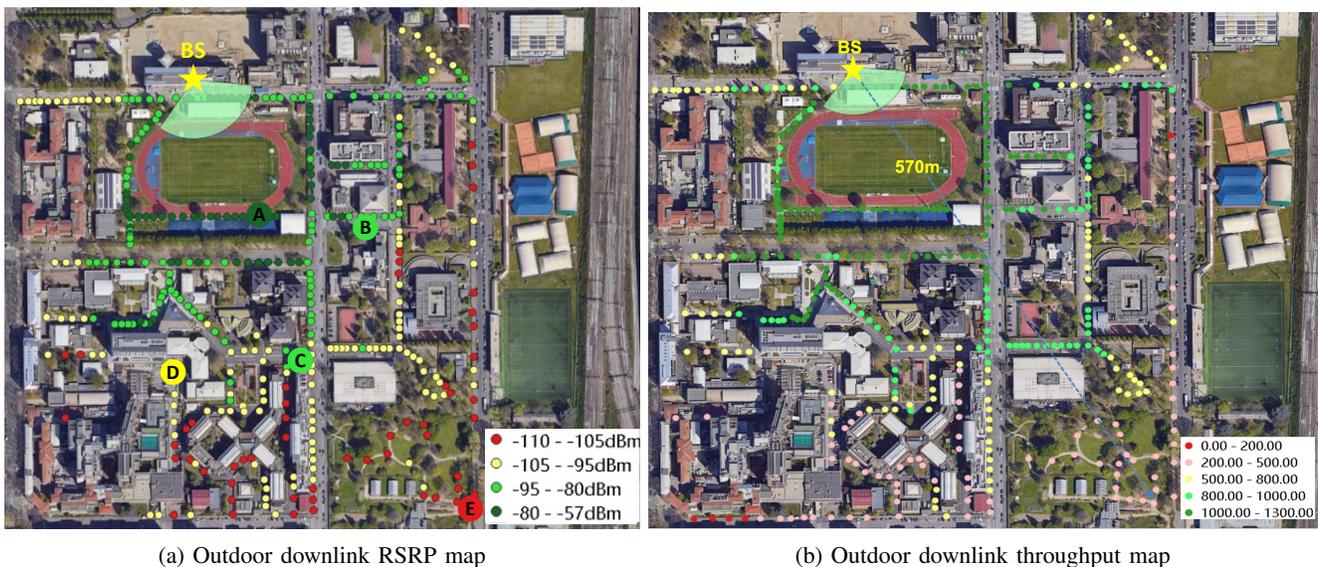
(a) Outdoor downlink RSRP map

(b) Outdoor downlink throughput map

Fig. 3: Outdoor downlink RSRP and throughput measurements.

(a) RSRP and downlink throughput

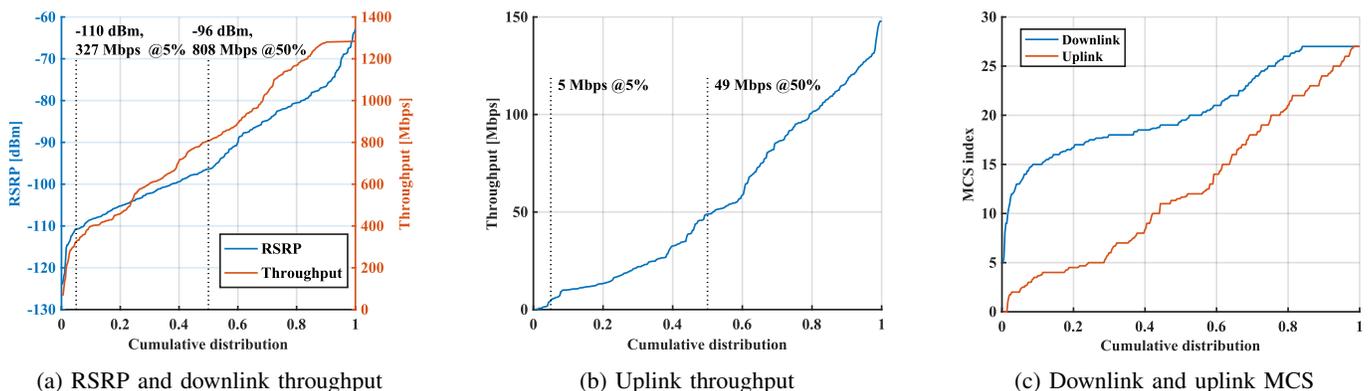(b) Uplink throughput

(c) Downlink and uplink MCS

Fig. 4: Empirical cumulative distribution function of network parameters

and can be attributed to three main factors. One contribution is the uplink-unfavorable 4:1 Time-Division Duplex (TDD) radio frame configuration. Then, the MIMO capabilities of the Test User Equipment (TUE) are more limited in the uplink, as outlined in *Table I*. Moreover, the reduced transmit power of the TUE in the uplink further impacts performance. This trend is clearly illustrated in *Fig. 4c*, where the empirical CDFs of the active MCS index in both the downlink and uplink are depicted. As the uplink transmission power diminishes, the RSRP is likewise reduced, resulting in consistently lower supported MCS indices compared to the downlink direction.

In summary, the system exhibits a degradation in uplink performance, particularly in challenging propagation conditions. This experiment suggests that enhancing uplink performance may require higher transmission power and additional MIMO layers for user equipment. However, this may pose challenges, especially in mobile terminals with power consumption constraints. Alternatively, the use of frame structures that prioritize the uplink can help compensate for the reduction in spectral efficiency, potentially through the implementation of dynamic Time-Division Duplex (TDD) techniques to prevent
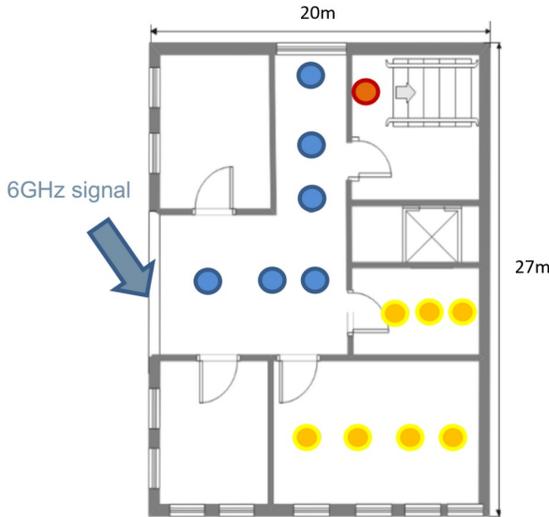
TP	Dist.	LoS	DL TP	UL TP	Rank (D/U)
A	175 m	yes	1282 Mbps	132 Mbps	4/2
B	236 m	no	992 Mbps	95 Mbps	4/2
C	310 m	no	770 Mbps	55 Mbps	4/2
D	344 m	no	550 Mbps	12 Mbps	3/1
E	570 m	no	332 Mbps	6 Mbps	2/1

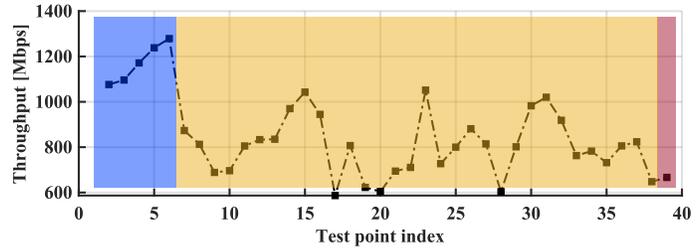TABLE II: Specific test point details and measurements.

unnecessary degradation in downlink performance (e.g., TDD 2:3) [15].

We conclude this section of the analysis by offering a more detailed examination of the system's performance in a range of distinct scenarios. These scenarios correspond to specific test points, highlighted in *Fig. 3a*, which exhibit noteworthy characteristics from a propagation environment perspective. The test point details and measurements are summarized in *Table II*.
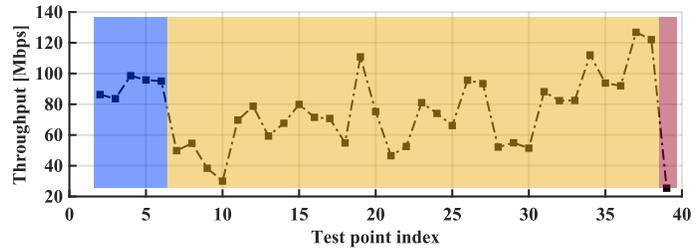
Test point A is situated 175 meters from the base station and enjoys a clear Line-of-Sight (LoS) connection, making it one of the test points with the most favorable propagation conditions. Here, the downlink throughput reaches 1282 Mbps

(a) Test points inside the building.

(b) Downlink throughput.

(c) Uplink throughput.

Fig. 5: O2I measurements. Blue, yellow, and red denote test points in the lobby, rooms with windows, and rooms without windows, respectively.

with a channel rank of 4, while the uplink throughput is 132 Mbps with a channel rank of 2. These throughputs set the performance benchmark for this study. Test point B is positioned slightly farther from the base station, at 236 meters. However, it is no longer within the LoS field, as it is obstructed by one building. In this scenario, the downlink and uplink throughput values decrease by about 22% to 28% compared to the benchmark, reaching 992 Mbps and 95 Mbps, respectively. The channel rank achieves the maximum in both directions. This drop in performance represents the typical degradation experienced when a single building obstructs the LoS, which is consistent with other test points exhibiting similar propagation characteristics.

Test points C and D are situated at slightly more than 300 meters from the base station and both have the LoS blocked by two buildings. However, they experience significantly different performances due to the "urban canyon" effect. Test point C is located on a road flanked by buildings and achieves performances of 770 Mbps for downlink and 55 Mbps for uplink, with no decrease in channel rank. In contrast, test point D is located in a "cul-de-sac" and does not benefit from the same "urban canyon" effect. Here, the throughput measurements are 550 Mbps in downlink and 12 Mbps in uplink, with both channel ranks reduced by 1. This scenario gives clues on the impact of building layout.

Lastly, we detail the performance observed at test point E. This spot is placed at the cell's edge, 570 meters away from the base station, with LoS obstructed by multiple buildings. Furthermore, it is situated in a park, devoid of a strong "urban canyon" effect in this case. Ultimately, this represents one of the less favorable propagation scenarios. The measured performance here is 332 Mbps for downlink throughput and

6 Mbps for uplink throughput, with channel ranks reduced to 2 and 1, respectively. While the peak downlink data rate might still accommodate most applications, the uplink data rate appears to be more adversely affected, as previously noted in the general analysis.

B. Outdoor-to-indoor performance

To evaluate the Indoor-to-Outdoor (O2I) performance, we conducted a similar analysis inside the building highlighted in *Fig. 1*. This building enjoys a LoS connection with the BS and is approximately 200 meters far from it. The LoS condition is perfect except for a minor coverage of the facade due to vegetation. Just outside of the building, the downlink and uplink throughputs measure around 1200 Mbps and 130 Mbps.

Fig. 5a illustrates a selection of crucial test points overlaid on the floor plan of the building's first floor.

Test points marked in blue are distributed in the lobby area, where a large glass window offers a direct LoS view with the BS. These points benefit from the best propagation characteristics, as electromagnetic waves only need to penetrate the glass windows. Consequently, both downlink and uplink performances are strong, as depicted in *Figures 5b* and *5c*, with average values closely resembling those measured outside the building. This confirms the signal's glass penetration capabilities.

Test points marked in yellow are situated in rooms that still have windows, but these do not overlook the BS. They enable radio signals to reach the TUE with relatively high strength through reflection and other propagation effects.

The red test point is located in a room without any windows but with a door opened over the lobby. Here, results demonstrate a significant decrease in performance, with downlink
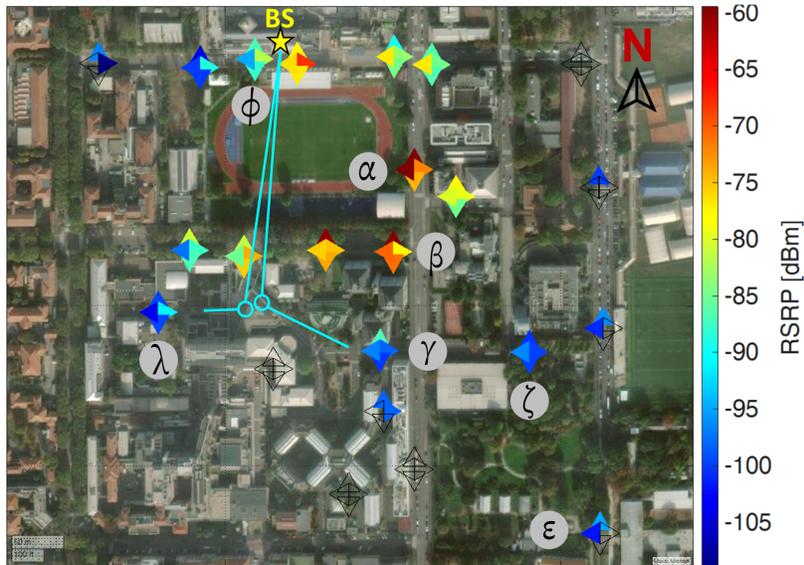
Fig. 6: Reference Signal Received Power heatmap

throughput halved and uplink throughput reduced by approximately 80%.

As for the previous analysis, it's evident that uplink performance is more sensitive to particularly challenging propagation conditions. Nonetheless, the system demonstrates commendable overall outdoor-to-indoor performance. An essential role in the propagation is played by the presence of windows, even when they are placed in NLoS with the base station.

IV. MMWAVE

This section reports the results of the measurement campaign at mmWave. Heatmaps and CDF will help provide both qualitative and quantitative samples of millimeter-wave network behavior around the deployment area. Through the beam identifiers, we were also able to reconstruct the approximate path of the beams toward the reception points. This information is also discussed in the most relevant cases. A glimpse of outdoor to indoor propagation is also provided at the end of the section.

A. Outdoor-to-outdoor performance

Consider the RSRP superimposed to the test area in *Fig. 6*. This parameter is an indicator of the strength of the signal received by the UE.

The highest received power occurs under the direct sector illumination. The spots in this region (e.g, α, β) reach RSRP values in the order of -60 dBm, which sets as the maximum received signal strength and corresponds to peak throughputs of slightly less than 1.3Gbps in Downlink (DL) and 250 Mbps in Uplink (UL). Once again, this unbalance in favor of DL throughputs originates from the Time Division Duplexing (TDD) 4:1 slot configuration, from the lower transmission power of the CPE with respect to the BS and the maximum QAM order. The pointing direction is also relevant. Point α in the north and west direction stably connects to a beam

enjoying -59 dBm and decreases to around -73 dBm when pointing east and south. Point β overlooks a street skirted by trees that were mostly bare at the time of the campaign. Here a direct LoS connection is established with the BS on the rooftop. In the most favorable direction, the RSRP is equal to the above-mentioned one, with a lower peak towards the west, probably due to the absence of close buildings that can provide major reflections.

The points immediately under the base station, shadowed by the building hosting the site, do not enjoy a direct LoS link. Nevertheless, a strong urban canyon effect greatly improves the coverage. For example, spot ϕ, confined between two buildings, reaches performances that coincide with the perfect LoS of point α, thanks to the aforementioned effect. However, moving a few meters west, where buildings do not surround the road, is sufficient to make this effect fade, and the signal rapidly worsens. The beam choice (not shown in the figure) in the point exactly under the BS is distinctive: the best beam quickly switches every time the CPE's orientation is changed. This is due to the separation between beams being not pronounced as a consequence of the reflection, thus creating a crowded scene where beams have similar propagation conditions. On the top-right side of the image, the signal rapidly decreases, denoting the sector's edge.

On the center-left part of the image, above λ, we enter a *soft* NLoS condition, where signals as high as -74 dBm are received from reflections, as the south and west-pointing arrows suggest.

Shifting toward the center of the picture, a dense block of buildings is encountered. Interestingly, the signal can infiltrate inside the block and reach locations where the LoS is completely blocked. For example, spot λ reaches -91 dBm while pointing east, thanks to a strong reflection on the building on its right, also highlighted in *Figure 6*.

On the opposite side of the block, the spot γ gets its maximum power when pointing north. Given the inability of the beam to
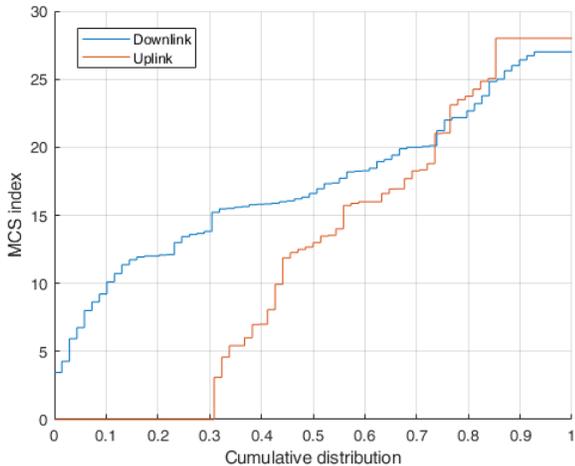
Fig. 7: Downlink and uplink Modulation and Coding Scheme CDF

Fig. 8: Downlink and RSRP CDF

TP	Dist.	LoS	DL TP	UL TP	RSRP
α	190 m	yes	1239 Mbps	235 Mbps	-59 dBm
λ	285 m	no	708 Mbps	0 Mbps	-91 dBm
γ	315 m	no	355 Mbps	17 Mbps	-87 dBm
ζ	395 m	no	407 Mbps	0 Mbps	-96 dBm
ϵ	568 m	no	435 Mbps	10 Mbps	-95 dBm

TABLE III: Specific test point details and measurements.

pass through the building ahead, the best beam must bounce on the building on the left and is eventually reflected, as highlighted in the figure.

Point ζ is also placed in soft NLoS: the direct path is covered by a 12m-tall building and foliage. Still, the connection is kept in every direction, reached by a decent signal strength which translates into downlink throughputs of 407 Mbps.

The presence of the spots on the right side of the picture is surprising. The existence itself of a connection in these points is notable since the ray tracer simulation foresees RSRP values as low as -115 dBm or even no coverage. Point ϵ is placed at slightly less than 600m, at the edge of the 6GHz cell, and its direct path is covered by several buildings and trees. The signal here is very low, indeed it cannot be received in every direction. Still, one orientation reaches an RSRP of -100dBm, which can supply a peak throughput of 435Mbps in DL. Part of the merit belongs to the antenna of the CPE, which gain helped in extracting the mmWave signal out. The uplink transmission in this area is afflicted by the harsh propagation, demonstrating once again that the UL is more delicate. More details regarding some of these peculiar test points are reported in *Tab. III*.

Finally, some spots where the signal cannot reach the UE are displayed with a black empty arrow.

Fig. 7 reports the MCS cumulative distribution function for both downlink and uplink. Note that these MCS indexes originate from two different MCS tables for uplink and downlink, due to the hardware configuration. The amount of locations where only downlink transmission is withstanded impacts the distribution. Around 30% of the points where the CPE can attach to the base station can transmit in the downlink but do not support the uplink transmission (represented with MCS equal to 0). Interestingly, around 15% of the UL measurements could reach the highest MCS, which is around twice the amount that can achieve the same in DL. This behavior is dictated by the specific implementation of the hardware used. The number of missed connections makes the UL MCS curve steeper. Interpreting this result, it could be argued that the UL
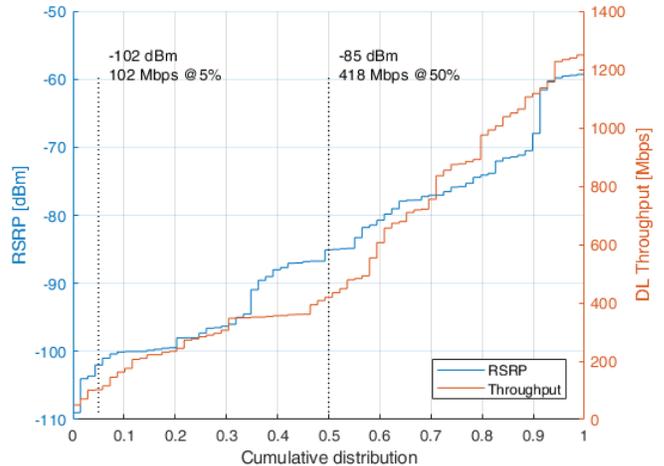
connection is more fragile since the margin between maximum performance and a missed connection is small. In *Fig. 8* the comparison between downlink throughput and RSRP is reported. The trend is as expected: high RSRP corresponds to high throughput values. This remarks that with our setup, the speed of the connection can be in general inferred by the reference signal strength.

From the results, some general trends can be derived. The angular resolution of the setup gives clues on the relevance of reflections to get the signal in NLoS conditions. Points reported in *Fig. 6* show that outside the LoS-illuminated sector, the highest-quality signal often arrives from reflections. One main difference between the LoS and NLoS conditions is the variance of the throughput values during the 15s-long capture. While the signal power is rather stable in LoS, the throughput and RSRP values fluctuate significantly in NLoS, as well as the anchor beam.

B. Outdoor-to-indoor performance

The campaign is completed with indoor measurements, shown in *Fig. 9*. Those measurements were captured in the same building discussed in the previous section (highlighted in Fig. 1). The entrance is made of a two-layer glass window. *Point 1* was chosen close to the door, to observe its effect on the signal. The peak throughput does not vary from what was obtained outside of the same building, indeed it reaches the upper limit peak of 1.3Gbps in DL and 251Mbps in UL towards the optimal direction (west). However, this speed fluctuates more than the corresponding measurement taken just outside of the glass, and the RSRP tops at -73dBm, which means a 14 dBm loss. This gives clues regarding the
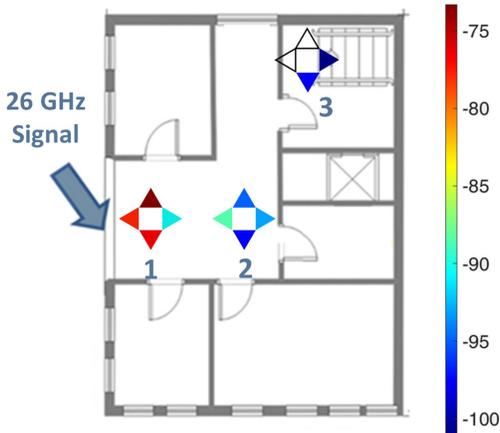
Fig. 9: Outdoor to indoor RSRP results [dBm]

	U6G		mmWave	
	UL	DL	UL	DL
Theoretical [Mbps]	170.33	1379.94	253.74	1354.34
Measured [Mbps]	132	1282	235	1239
Loss [%]	22.5	7.1	7.4	8.5

TABLE IV: Maximum achievable throughput in Mbps computed with 3GPP formula [16] and maximum measured values.

absorption of the window. Comparing *point 1* north and east directions, one can also conclude that the signal reflected inside the lobby is 17 dBm lower than the direct path.

Point 2 has less favorable conditions thus, as expected, the throughputs are slightly lower, reaching up to 560 Mbps in DL and 70 in UL.

Point 3 is instead placed in a room with no windows but with a door open on the lobby. Here the CPE only manages to connect in two directions (those closest to the door), where it maintains a good connection in DL with 425 Mbps and keeps the link with 20 Mbps in UL. These measurements once again highlight how the presence of windows positively impacts indoor penetration, while the concrete completely blocks the signal. It is worth noting that in almost every connectable direction, the downlink throughput exceeds 400Mbps throughput.

V. COMPARISON

The measurement campaigns summarized in the previous sections prove how both U6G and mmWave have the potential of bringing Gigabit-level performance to the RAN. However, the two deployments still retain some fundamental differences. Exploiting the co-location of the two cells, in this section we give a comparative analysis of both systems' performance to provide a comprehensive understanding of their respective strengths and limitations.

A. Expected vs measured throughput and overhead

The U6G and the mmWave cells present the maximum achievable throughput values reported in Tab. IV. The value for U6G was computed using the 3GPP formula in [16], while, for mmWave, the 3GPP formula in [17] was used. This choice was obliged since [16] was evidently not precise for high frequencies, and resulted in an unjustifiedly high value. These theoretical values are computed considering the cell configuration (i.e. bandwidth, MIMO layers, etc.) and represent the achievable MAC throughput for both deployments. Note that the UE have different capabilities. In the UL direction, both cells have the same number of layers and the expected

throughput difference is mainly given by the larger bandwidth available at mmWave. On the other hand, the U6G cell has 4 MIMO layers in the DL direction, two more than the mmWave cell, which balances the larger bandwidth gains in mmWave. Tab. IV also reports the maximum measured throughput in our campaigns, showing how these values are slightly lower than the theoretical maximum. This result was expected, and it is known to be caused by several factors. First, it is not easy to estimate the overhead caused by the control plane. 3GPP suggests empirical values that account for it, but some margin remains. In realistic settings, the control plane overhead is tightly related to both the cell configuration (e.g. PRACH, channel estimation and positioning signals density) and the instantaneous network conditions (e.g. handovers, user attachments). Thus, a mismatch is to be expected. Second, this formula computes the MAC throughput but does not consider additional bottlenecks generated by higher layers. Among these, the transport layer is most notable, as it is known that TCP under-performs over less-than-stable wireless links, especially mmWave [18]. Overall, the gap between theoretical and measured results can be attributed to reasonable deteriorations related to the system implementation and the real-world evironment.

We observed this behavior for both U6G and mmWave, but the effects are slightly more pronounced for the latter. This is, once again, expected due to the harsher propagation, the increased control plane overhead, and the lower TCP-over-mmWave performance.

B. Performance and coverage comparison

We continue the analysis by comparing the performance and coverage of both systems at different test points, which we report in Tab. V. The two technologies show fundamental differences in the cell configurations and device capabilities, such as bandwidth, numerology and MIMO layers. While these differences are representative of typical cell deployments, they strongly impact the final performance. As such, alongside a direct throughput comparison, we also report the spectral efficiency computed in terms of capacity over Hertz per MIMO channel. Furthermore, the values reported for the mmWave case are taken selecting the best pointing direction for the mmWave UE.

Points A and α represent the position with the highest measured performance in both campaigns, where both technologies can establish a full-rank LoS connection. Here the lower downlink spectral efficiency and reduced MIMO layers of mmWave are compensated by the larger available bandwidth, making the downlink throughput comparable. On the other hand, the uplink throughput is almost doubled in mmWave since 2 MIMO layers are active in both cells.

Points	Distance [m]	LoS	Bandwidth [MHz]		Rank (D/U)		Throughput (D/U) [Mbps]		Spectral efficiency (D/U) [bps/Hz/ch]	
			U6G	mmWave	U6G	mmWave	U6G	mmWave	U6G	mmWave
A, α	175	y			4/2	2/2	1282/132	1239/235	5.4/3.3	4.2/2.3
B, β	236	n			4/2	2/2	992/95	891/112	4.2/2.5	3.0/1.2
C, γ	310	n	80	200	4/2	2/2	770/55	355/17	3.2/1.3	1.2/0.4
D, n/a	344	n			3/1	n/a	550/12	0/0	3.1/0.4	n/a
E, ϵ	570	n			2/1	2/2	332/6	435/10	2.8/0.4	1.5/0.4

TABLE V: Comparison between U6G and mmWave capabilities

Points B, β are in a NLoS condition caused by a single building. Despite the obstruction, we still observe a relatively high throughput, with a more pronounced impact on the mmWave system, as expected. Nonetheless, the mmWave uplink is still higher than the U6G one. Points C and γ are obstructed by multiple buildings, but they benefit from a urban canyon effect, as previously mentioned. Here we observe a sharp mmWave performance degradation.

More in detail, the mmWave connection is still at full rank, but the equivalent spectral efficiency is more than halved with respect to the previous case. On the contrary, the U6G system experiences a less pronounced performance degradation. Point D represents a particularly harsh test, as multiple tall buildings obstruct the LoS with no urban canyon effect. Here the mmWave UE fails to attach to the base station, while the U6G system can still provide almost half of the full capacity.

Points E and ϵ are the furthest from the base station. However, the LoS obstruction here is less severe with respect to the previous case. In this case, with a maximum achievable channel rank of 2 for both system, the higher mmWave bandwidth compensates for the reduced spectral efficiency and allows higher performance in both directions.

According to the detailed comparison given above, some general trends can be observed. Despite the different cell configuration and device capabilities, the two systems show comparable best case performance. As expected, mmWave shows higher sensitivity to penetration losses and achieves lower performance than U6G under severe NLoS. On the other hand, the U6G system heavily relies on multiple active MIMO streams to offer a performance level on par with mmWave. Consequently, MIMO-adverse propagation environments can be better exploited by mmWave. Indeed, channel separation performs generally better at higher frequencies [19], allowing the mmWave system to potentially enjoy a higher channel rank. At the same time, a single mmWave spatial stream has higher potential throughput with respect to U6G, making the loss of spatial diversity less impactful on the mmWave system. Such behaviour is expected and it is confirmed by our results in points E and ϵ.

For what concerns the Outdoor-to-indoor (O2I) scenario, both systems show a good penetration of the building's glass window. As the test point is moved deeper inside the building, the performance degradation follows the same trend for both. Most notably, however, the direction of the mmWave UE has an impact on the indoor performance, while this is not true for the ominidirectional U6G UE.

C. Performance improvement margins

From the comparative analysis given in this section, it appears that U6G dominates mmWave almost entirely in the context of macro urban coverage. Such result is not surprising, mostly due to the well-known harsher propagation at mmWave. However, the cell configuration and the technology maturity level of the involved devices have to be taken into account when forecasting for realistic future performance.

Indeed, in those propagation environments not heavily dominated by penetration losses, U6G performs better because up to 4 MIMO layers can be activated, as opposed to mmWave which has only 2. However, statistical data from urban macro coverage at lower frequencies, thus representative of an high technological maturity and realistic UE capabilities, shows that up to 2 MIMO layers are active for most of the connections, even when 4 layers are available [20]. At the same time, we can expect 4 MIMO layers being available also for mmWave UEs with higher technological maturity. Additionally, up to 400 MHz of cumulative bandwidth are available for mmWave, potentially doubling the overall performance. On the other hand, the 80 MHz bandwidth configuration of our U6G deployment is already reasonably close to the maximum bandwidth availability in U6G. All together, the suggestion goes towards a larger performance improvement potential for mmWave deployments, especially if higher MIMO capabilities will be made available by improved technological maturity.

Finally, uplink traffic is gaining increasing attention [21] but both systems show overall poor performance. This can be mitigated by selecting a more favourable TDD frame structure in those cells where enhanced uplink is required. However, this option is only viable for mmWave networks, as U6G coverage is expected to operate at the macro level, making heterogeneous frame structures impossible due to inter-cell synchronization constraints.

VI. RELATED WORKS

Given the recent introduction of the U6G in the standard and the WRC-23 licensing decision, there are no other articles that document a demonstration of a cellular network working at these frequencies. Therefore, this is the first work that reports measurement data on a U6G 5G deployment.

The same does not hold for mmWave, for which few contributions are instead present in the literature. Still, the relatively small number of mmWave commercial 5G deployments delayed the measurement effort. At the time of writing, the literature containing such measurements is composed of only a few contributions, briefly commented in the following. This shortage hinders the study and the optimization of such networks, which is one of the impairments for the adoption of

this technology.

Authors in [22] carry out several measurements over commercial, Non Stand Alone (NSA), 5G network deployed in Chicago, Illinois, working at 28 and 39 GHz. While a user is moving, the signal's physical data and throughputs are extracted, and emphasis is put on the derived beam management techniques. Their dataset is publicly available[4]. The work in [23] studies the first mmWave networks deployed in the U.S., namely in Minneapolis, Chicago and Atlanta in 2019. It offers insights on both stationary and moving UEs, including an analysis of handoffs. Collected physical-layer parameters are limited to the RSRP, while some more information is extracted from upper layers. The dataset is publicly available[5]. Authors in [24] perform measurements in Boston, Chicago, and Indianapolis at 28 and 39 GHz, both static and in mobility, but limited to uplink and with a main focus on upper layers[6]. Authors in [25] conducted a large test drive to map the handoffs behavior, passing also through mmWave stations. Finally, [26] extends the already rich set of campaigns in Chicago with one in Miami, Florida.

This literature produced a significant dataset for few major cities in the U.S. The only campaign that took place in Europe is discussed in [27]. Coverage measurements are done in Oslo, Norway, and the results include interesting insights on foliage, human body attenuation, and propagation close to water.

The lack of measurement campaigns in Europe is an important issue: while it is, in general, difficult to extend propagation characteristics through different environments, this is especially true for European cities, which significantly differ from those of the available campaigns. Therefore, more data and tests are still needed.

Finally, a few more contributions focus on more specific aspects of the network operations. The work in [11] reports signal strength, throughput and latency of an outdoor-to-indoor measurement performed in Chicago, also comparing mmWave performances (28 GHz) with LTE ones. Authors in [28] measure high-level characteristics and UE power consumption in 28/39 GHz 5G networks in the U.S. and compare them to those of 4G in the same location. Their dataset is publicly available[7]. Authors in [29] collected data[8] and developed a machine learning model to predict throughputs based on them. Finally, [30] focuses on the measurement of latency and other end-to-end performance indicators.

VII. Conclusion

This paper presents a thorough analysis of the results of two measurement campaigns on 5G networks working in upper 6GHz and millimeter-wave bands.

The U6G network, operating at a central frequency of 6.8 GHz, covered an area with a 600m radius and achieved throughputs of up to 1.3 Gbps in downlink and 150 Mbps in uplink. The network demonstrated stable performance,

maintaining throughputs higher than 200 Mbps in downlink for the majority of points, with exceptions due to signal strength issues constituting less than 5%. The ray-tracer results overall comply with the empirical results. Indoor penetration is larger than expected, particularly in the presence of windows.

The network working at 26GHz exhibits a coverage area compliant with the one predicted by the ray-tracing simulations. Maximum performances achieved are about 1.25 Gbps in downlink and 230 Mbps in uplink, demonstrating the high potential of the mmWave to offer unprecedented uplink speeds in LoS. This technology can also boast a high margin for improvement given the large bandwidth available at those frequencies. Despite the harsher propagation at mmWave, the base station deployment covered an area comparable to the existing commercial macro cells in the neighborhood, working at sub-6GHz frequencies. Outdoor to indoor propagation is limited to spaces where windows are present.

5G and beyond networks heavily rely on the availability of additional spectrum, and U6G and mmWave bands are relevant assets. This paper showcases the capabilities of radio access networks operating at such frequencies and compares them within a realistic scenario for which only limited information is currently available.

A. Acknowledgments

References

[1] T. S. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. N. Wong, J. K. Schulz, M. Samimi, and F. Gutierrez, "Millimeter wave mobile communications for 5g cellular: It will work!," *IEEE access*, vol. 1, pp. 335–349, 2013.

[2] N. I. of Standards and Technologies, "The spectrum crunch," 2022.

[3] Federal Communication Commission (FCC), "Unlicensed Use of the 6 GHz Band; Expanding Flexible Use in Mid-Band Spectrum Between 3.7 and 24 GHz," Report and order FCC-20-51, 04 2020. ET Docket No. 18-295; GN Docket No. 17-183.

[4] W.-F. Alliance, "Wi-fi alliance celebrates the 2023 wrc-23 decision on the 6.425-7.125 ghz frequency band," 2023.

[5] R. S. P. G. (RSPG), "Opinion on the ITU-R World Radiocommunication Conference 2023," report, European Commission, 12 2022.

[6] 3GPP, "5g; nr; user equipment (ue) radio transmission and reception; part 2: Range 2 standalone (3gpp ts 38.101-2 version 17.11.0 release 17)," Tech. Rep. TS 138.101-2 V17.11.0 (2023-10), 3GPP, 2023.

[7] 3GPP, "5g; nr; user equipment (ue) radio transmission and reception; part 1: Range 1 standalone (3gpp ts 38.101-1 version 17.11.0 release 17)," Tech. Rep. TS 138 101-1 V17.11.0 (2023-10), 3GPP, 2023.

[8] ITU-R, "Minimum requirements related to technical performance for imt-2020 radio interface(s)," Tech. Rep. Report ITU-R M.2410-0 (11/2017), ITU-R, 2017.

[9] Y. Chang, "Mmwave development in 2023: Where's it going & what are the challenges," 2023.

[10] FCC, "Millimeter wave propagation: Spectrum management implications," Tech. Rep. Bulletin number 70 (06/1997), 1997.

---

[4]https://5gbeams.umn.edu

[5]https://fivegophers.umn.edu/www20/

[6]https://github.com/NUWiNS/sigcomm-5gmemu-5g-mmWave-uplink-data

[7]https://github.com/SIGCOMM21-5G/artifact

[8]https://lumos5g.umn.edu/

[11] M. I. Rochman, V. Sathya, and M. Ghosh, "Outdoor-to-indoor performance analysis of a commercial deployment of 5g mmwave," in *2022 IEEE Future Networks World Forum (FNWF)*, pp. 519–525, 2022.

[12] M. Morini, E. Moro, I. Filippini, and A. Capone, "Will the upper 6 ghz bands work for 5g nr? a urban field trial," in *Proceedings of the 17th ACM Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization*, WiNTECH '23, (New York, NY, USA), p. 49–55, Association for Computing Machinery, 2023.

[13] M. Amini, R. Stanica, and C. Rosenberg, "Where are the (cellular) data?," *ACM Comput. Surv.*, vol. 56, sep 2023.

[14] SIRADEL, "S_5GChannel," 2023. Available: https://www.siradel.com/s5gchannel-for-advanced-5g-mmw-propagation-modeling/ [Accessed: 2023-18-12].

[15] H. Kim, J. Kim, and D. Hong, "Dynamic tdd systems for 5g and beyond: A survey of cross-link interference mitigation," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 4, pp. 2315–2348, 2020.

[16] 3GPP, "5g; nr; user equipment (ue) radio access capabilities (3gpp ts 38.306 version 17.6.0 release 17)," Tech. Rep. TS 138 306 V17.6.0 (2023-10), 3GPP, 2023.

[17] 3GPP, "5g; nr; physical layer procedures for data (3gpp ts 38.214 version 17.8.0 release 17)," Tech. Rep. TS 138 214 V17.8.0 (2024-02), 3GPP, 2024.

[18] Y. Ren, W. Yang, X. Zhou, H. Chen, and B. Liu, "A survey on tcp over mmwave," *Computer Communications*, vol. 171, pp. 80–88, 2021.

[19] M. Mizmizi, D. Tagliaferri, D. Badini, C. Mazzucco, and U. Spagnolini, "Channel estimation for 6g v2x hybrid systems using multi-vehicular learning," *IEEE Access*, vol. 9, pp. 95775–95790, 2021.

[20] 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," *arXiv preprint arXiv:2312.00957*, 2023.

[21] NGMN, "5G TDD Uplink v1.0," report, NGMN, 02 2022.

[22] 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 *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications*, pp. 800–809, 2022.

[23] 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 *Proceedings of The Web Conference 2020*, WWW '20, (New York, NY, USA), p. 894–905, Association for Computing Machinery, 2020.

[24] 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 *Proceedings of the ACM SIGCOMM Workshop on 5G and Beyond Network Measurements, Modeling, and Use Cases*, 5G-MeMU '22, (New York, NY, USA), p. 29–35, Association for Computing Machinery, 2022.

[25] A. Hassan, A. Narayanan, A. Zhang, W. Ye, R. Zhu, S. Jin, J. Carpenter, Z. M. Mao, F. Qian, and Z.-L. Zhang, "Vivisecting mobility management in 5g cellular networks," in *Proceedings of the ACM SIGCOMM 2022 Conference*, SIGCOMM '22, (New York, NY, USA), p. 86–100, Association for Computing Machinery, 2022.

[26] M. I. Rochman, V. Sathya, N. Nunez, D. Fernandez, M. Ghosh, A. S. Ibrahim, and W. Payne, "A comparison study of cellular deployments in chicago and miami using apps on smartphones," in *Proceedings of the 15th ACM Workshop on Wireless Network Testbeds, Experimental Evaluation & CHaracterization*, WiNTECH '21, (New York, NY, USA), p. 61–68, Association for Computing Machinery, 2021.

[27] S. Mohebi, F. Michelinakis, A. Elmokashfi, O. Grøndalen, K. Mahmood, and A. Zanella, "Sectors, beams and environmental impact on the performance of commercial 5g mmwave cells: An empirical study," *IEEE Access*, vol. 10, pp. 133309–133323, 2022.

[28] 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 *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, SIGCOMM '21, (New York, NY, USA), p. 610–625, Association for Computing Machinery, 2021.

[29] A. Narayanan, E. Ramadan, R. Mehta, X. Hu, Q. Liu, R. A. K. Fezeu, U. K. Dayalan, S. Verma, P. Ji, T. Li, F. Qian, and Z.-L. Zhang, "Lumos5g: Mapping and predicting commercial mmwave 5g throughput," in *Proceedings of the ACM Internet Measurement Conference*, IMC '20, (New York, NY, USA), p. 176–193, Association for Computing Machinery, 2020.

[30] R. A. K. Fezeu, E. Ramadan, W. Ye, B. Minneci, J. Xie, A. Narayanan, A. Hassan, F. Qian, Z.-L. Zhang, J. Chandrashekar, and M. Lee, "An in-depth measurement analysis of 5g mmwave phy latency and its impact on end-to-end delay," in *Passive and Active Measurement* (A. Brunstrom, M. Flores, and M. Fiore, eds.), (Cham), pp. 284–312, Springer Nature Switzerland, 2023.

Marcello Morini (Member, IEEE) is currently a PhD student at Politecnico di Milano, Department of Electronics, Information and Bioengineering. His research area is Telecommunication, with a focus on high-frequency radio access networks and smart propagation environments. He received his BSc degree in Electronic Engineering in 2020 at Università di Modena e Reggio Emilia and he completed his MSc in Telecommunication Engineering in 2022 at Politecnico di Milano.

Eugenio Moro (Member, IEEE) is currently an Assistant Professor at Politecnico di Milano, Milan, Italy. He received the M.Sc. and Ph.D. cum laude in 2019 and 2023, respectively. He was a visiting researcher in Nokia Bell Labs, Stuttgart, Germany and Northeastern University, Boston MA, USA. His research interests are wireless networks, with a focus on high-frequency radio access, wireless network programmability and optimization and smart propagation environments.

Ilario Filippini (Senior Member, IEEE) is an Associate Professor at the Dipartimento di Elettronica, Informazione e Bioingegneria of Politecnico di Milano. He received an M.S. in Telecommunication Engineering and a Ph.D. in Information Engineering from Politecnico di Milano in 2005 and 2009, respectively. From February 2008 to August 2008, he was visiting the Department of Electrical and Computer Engineering of The Ohio State University in Columbus (OH) working on routing in Cognitive Radio Networks. He works on networking topics, his main research activities include radio resource management and optimization in wireless networks, programmable networks, and smart radio environments. He has co-authored more than 70 peer-reviewed conference and journal papers. He serves in the Technical Program Committee of major conferences in Networking and as an Editor of Computer Networks (Elsevier).

Antonio Capone (Fellow, IEEE) is currently the Dean of the School of Industrial and Information Engineering, Politecnico di Milano (Technical University of Milan). His main research activities include radio resource management in wireless networks, traffic management in software defined networks, network planning, and optimization. He is author of over 250 publications on these topics. He contributes to major international conferences on networking as a technical program committee member. He is an Editor of the IEEE TRANSACTIONS ON MOBILE COMPUTING, Computer Networks, and Computer Communications and served as an Editor for the ACM/IEEE TRANSACTIONS ON NETWORKING from 2010 to 2014.

Danilo De Donno received the B.Sc. and M.Sc. degrees in Telecommunication Engineering from Politecnico di Milano, Italy, in 2005 and 2008, respectively, and the Ph.D. degree in Information Engineering from the University of Salento, Lecce, Italy, in 2012. He was a Post-Doctoral Fellow with the ElectroMagnetic Lab Lecce (EML2) of the University of Salento from 2012 to 2015 and a Post-Doc Researcher with the IMDEA Networks Institute, Madrid, Spain, from 2015 to 2017. In July 2017, he joined the Huawei Research Center in Milan, Italy, as a Wireless System Engineer. His main research focus is on mmWave technologies, systems and networks, beyond-5G RAN topologies, optimization of Smart Electromagnetic Environment.