The Chronicles of Deploying a 5G Campus Network: Challenges, Insights, and Experiences

Philip Diederich*, Valentin Thomas Haider*, and Wolfgang Kellerer
Chair of Communication Networks, Technical University of Munich, Germany
{firstname.lastname}@tum.de

***Abstract*—Private 5G Campus Networks (CNs) are rapidly gaining momentum as emerging industrial applications and open-source ecosystems push beyond the capabilities of public mobile networks. However, despite this growing interest, real-world insights into how private 5G networks behave at scale remain limited. Providing such insights is difficult because building and operating a 5G CN requires substantial hardware investments and operational expertise, making large-scale deployments rare. In this work, we introduce the TUM 5G CN, which, to the best of our knowledge, is the first documented deployment spanning multiple university campuses and integrating a heterogeneous set of Base Stations (BSs). We describe its physical and network architecture, assess throughput, latency, and coverage, and outline planned extensions. The presented findings support researchers and practitioners in deploying their own private 5G networks.**

***Index Terms*—Campus Network, 5G, 6G, Cellular Network**

I. Introduction

With the introduction of the 5[th] generation (5G) of mobile networks and the availability of locally licensed spectrum for industrial and research environments, private cellular networks have gained significant attention in recent years [1], [2]. These private 5G networks, often referred to as 5G Campus Networks (CNs), are envisioned to enable a wide range of new use cases, including smart factories [3], [4], smart agriculture and construction [1], remote driving and flying for maintenance and rescue [5], as well as tele-medicine applications such as patient monitoring or robotic surgery [6]. However, to successfully deploy 5G CNs for such use cases, a detailed understanding of the underlying network architectures, deployment and maintenance complexity, and the performance and cost characteristics of different CN types is required.

In general, similar to 5G networks operated by Mobile Network Operators (MNOs), a 5G CN consists of a core network, a Radio Access Network (RAN) comprising one or several Base Stations (BSs) operating in licensed spectrum, and an underlying transport network that interconnects the core and all RAN components. Two approaches exist for deploying a 5G CN: (i) using open-source SW running on commodity servers in combination with Software Defined Radios (SDRs), or (ii) deploying commercially available HW and SW from established vendors or start-ups. These approaches naturally

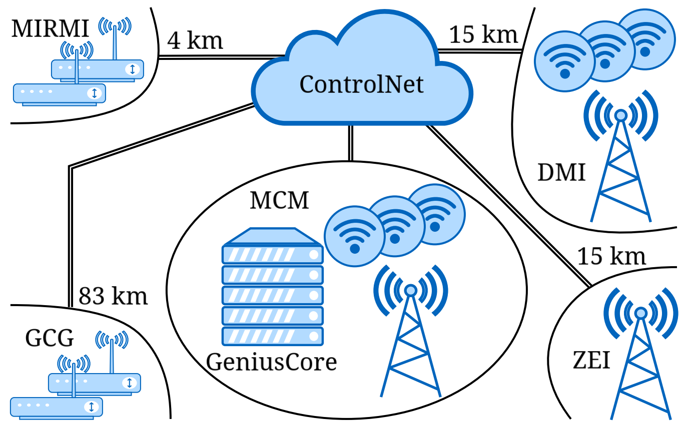

Fig. 1: The physical setup of the TUM 5G CN consisting of the core network (GeniusCore), indoor cells (MCM, DMI), outdoor cells (MCM, DMI, ZEI), and all-in-one BSs (MIRMI, GCG) connected through a control network (ControlNet).

differ in terms of performance, cost, flexibility, and scalability. While open-source based deployments are cheaper and offer greater flexibility, they may exhibit limitations in scalability, robustness, or performance compared to commercial solutions.

Both deployment options have already been explored in the literature. Several works [4], [5], [7]–[11] present campus networks that are fully or partially based on open-source software components like OpenAirInterface (OAI) [12], srsRAN [13], free5GC [14], or Open5GS [15]. These deployments are evaluated with respect to their system performance, including throughput, latency, jitter, coverage, and CPU or memory utilization. In contrast, there also exist CNs that are completely based on commercially available products. In [16], the authors evaluate the influence of logical link design on the transmission latency of packets. However, the vendors and detailed information about the 5G CN are not disclosed. The authors of [17], [18] present a CN that is completely based on commercial products and evaluate the system in terms of electromagnetic field (EMF) strength measurements, throughput, latency, and jitter measurements, and radio-level metrics such as Signal to Interference plus Noise Ratio (SINR) and Reference Signal Received Power (RSRP). While this system resembles a similar 5G CN to the one presented in this work, it differs substantially in its architecture, as our CN spans multiple campuses and incorporates a diverse set of BSs.

In this paper, we provide a comprehensive characterization of a large-scale, multi-site 5G CN deployment and present

This work was supported by the Federal Ministry of Research, Technology and Space of Germany (BMFTR) under the projects “6G-ANNA” and “6G-life” with project identification numbers 16KISK107 and 16KISK002 as well as by the Bavarian Ministry of Economic Affairs, Regional Development and Energy as part of the project 6G Future Lab Bavaria.

*The first two authors contributed equally to this work.

the TUM (Technical University of Munich) 5G CN. Its system diagram is shown in Fig. 1. We describe the physical and network architecture of our CN (Sec. II), which poses unique challenges due to the substantial inter-site distances between the BSs and the core. Moreover, our setup comprises a variety of indoor and outdoor BSs, including all-in-one BSs, conventional Baseband Unit (BBU)/Remote Radio Head (RRH)-based base stations, and massive MIMO (mMIMO) BSs. Alongside these descriptions, we share our experiences and the lessons learned during the deployment phase. To provide insights into the performance of our CN, we present results from throughput, latency, and coverage measurements (Sec. III). Finally, we outline future expansion plans (Sec. IV) before concluding the paper (Sec. V).

II. System Setup

In this section, the physical system deployment of the TUM 5G CN is first described. Subsequently, the underlying IP network for connecting all system parts is explained in detail.

A. Physical System Deployment

The heart of the TUM 5G CN is the core network, which is located at TUM Main Campus *in Munich* (MCM). In particular, the deployed core network is the GeniusCore solution, which is developed and maintained by CampusGenius [19]. The GeniusCore is built of a set of Network Functions (NFs) according to the 3GPP standardization [20]. Specifically, the NFs Network Repository Function (NRF), User Plane Function (UPF), Access and Mobility Management Function (AMF), Session Management Function (SMF), Unified Data Management (UDM), Unified Data Repository (UDR), Authentication Server Function (AUSF), Policy Control Function (PCF), and 5G-Equipment Identity Register (5G-EIR) are implemented and deployed in a dockerized [21] environment.[1]

The other main part of the TUM 5G CN is its RAN. It consists of a total of seven BSs (gNBs) located at various locations across the different TUM campuses. An overview of the physical network architecture is depicted in Fig. 1. In general, the 5G CN operates in the n78 frequency band at 3.7 GHz with a channel bandwidth of 100 MHz and a Subcarrier Spacing (SCS) of 30 kHz [22]. Depending on the type of BS, each BS is managing and controlling one or more radio cells. At MCM, the BS consists of a BBU, two Smart Hubs connecting to a total of 10 indoor pico Remote Radio Heads (pRRHs), and two outdoor RRHs, which are all supplied by Nokia. The indoor pRRHs serve a total of four radio cells, while the two outdoor RRHs serve one cell each. A similar setup is located at TUM Departments of Mathematics and Computer Science *in Garching* (DMI), where the BS is managing two indoor radio cells served by four pRRHs each, as well as three outdoor radio cells which are each served by an RRH. The TUM Zentrum für Energie und Information (Center for Energy and Information) *in Garching* (ZEI) is supplied with a BS that manages two outdoor radio cells, each served by a 64T64R (64 transmit and 64 receive) mMIMO Radio Unit (RU). All cells at MCM, DMI, and ZEI operate with a 4D1U Time Division Duplex (TDD) slot pattern (four downlink and one uplink slot). Furthermore, the Multiple Input Multiple Output (MIMO) mode in the Downlink (DL) is 4×4, while it is closed-loop 2×2 in the Uplink (UL). The MIMO mode specifies the number of supported spatial layers per User Equipment (UE). Finally, TUM Munich Institute of Robotics and Machine Intelligence *in Munich* (MIRMI) and TUM Geriatronics Campus *in Garmisch-Patenkirchen* (GCG) are each connected to the CN via two all-in-one BSs, where each BS serves one indoor radio cell. These BSs are supplied by LiteOn and operate in a 4D1U TDD slot pattern and use a 2×2 MIMO mode in the DL. Although all sites are served by the same 5G core network, they are quite geographically dispersed across the campuses. The resulting air-distances between each BS location and the core network are summarized in Table I.[2]

TABLE I: Air-distances between BS locations and the core network located at MCM

BS Location	MCM	DMI	ZEI	MIRMI	GCG
Distance [km]	0.00	14.69	15.13	3.92	82.65

The indoor radio cells at MCM and at DMI are each served by several pRRHs and provide coverage for three hallways (one long hallway at MCM and two shorter hallways at DMI). Importantly, the assignment of pRRHs to individual radio cells is not fixed. As operators of the 5G CN, we are able to dynamically modify these associations. This flexibility enables cell configurations that would typically not be deployed in public mobile networks. For example, we may assign each pRRH placed along a single hallway to a different radio cell. Although such a configuration is clearly suboptimal from a conventional network performance perspective, it substantially increases the number of handovers experienced by a UE when moving along the hallway. This behavior allows us to collect detailed mobility and handover performance data for research purposes. Further important parameters that can be adjusted depending on the research purpose include, e.g., the transmission power of the radio heads, the used TDD pattern, the channel bandwidth, the Channel State Information (CSI) reporting periodicity, the Demodulation Reference Signal (DMRS) configuration, or the handover triggering condition and its parameters.

B. IP Network Configuration

All base stations maintain continuous connectivity to the core network at MCM for stable operation. However, connecting the remote locations with the core is challenging. The base stations and the core require either their own isolated IP domain or Layer 2 connectivity. Due to the geographic distances, remote base stations are connected in a heterogeneous way to the university's network provider Leibniz-Rechenzentrum (LRZ). As a consequence, the backbone network fragments

[1]For a more detailed description and explanation of the functionalities of these NFs, the reader is referred to [20].

[2]The exact device model names of all components are omitted for brevity. Interested readers may contact us by email for additional information.

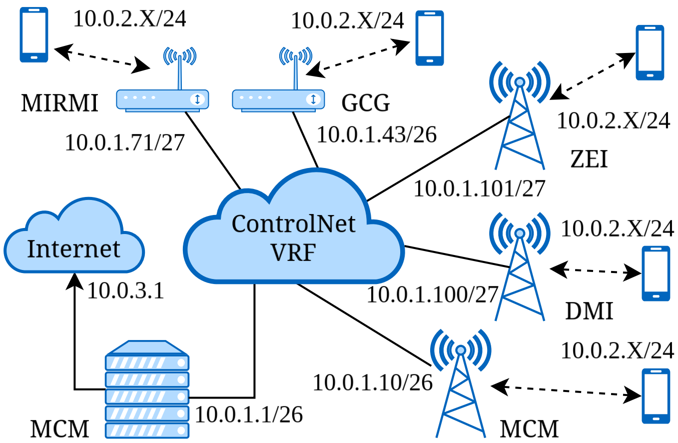

Fig. 2: The IP-level system architecture of the TUM 5G CN. The VRF function connects all locations with the core server at MCM. Control and user traffic is forwarded via the VRF to the respective network functions in the core.

locations into individual domains. A Virtual Routing and Forwarding (VRF) function solves the connection issue and creates an isolated IP domain. The VRF function provides IP reachability to all base stations and the core network, enabling the exchange of control and user traffic. Fig. 2 visualizes the configuration of the campus network setup. The VRF setup provides a control network that is divided into one /26 and two /27 IP networks. Subnets are assigned to remote locations based on their physical connection to the LRZ backbone. For example, MIRMI and ZEI connect to the LRZ backbone via two different core routers. Therefore, MIRMI and ZEI have different IP ranges and addresses in the control net[3]. Consequently, each base station has an IP address in its IP subnet. Furthermore, the core at MCM connects to the control network and the internet for connectivity. Finally, as shown in Fig. 2, all UEs are assigned IP addresses within the same IP subnet regardless of their physical location.

C. *Challenges and Lessons Learned*

Deploying the TUM CN provided several practical insights. First, reliable GPS reception is essential, particularly for all-in-one BSs. Insufficient or unstable GPS synchronization can lead to unpredictable system behavior, thereby complicating debugging and troubleshooting.

Second, the placement of the pRRHs in the hallway was predetermined well before the CN installation. In retrospect, the chosen positions were unnecessarily close to one another. Greater spacing between individual pRRHs would have been feasible without incurring significant performance degradation and would have resulted in either greater coverage or a reduction in capital expenditure.

Commercial Nokia equipment allows configuration of nearly all radio parameters, but their interdependencies make the setup process prone to errors. In contrast, all-in-one BSs provide fewer configuration options but offer a more user-friendly interface with sufficient controls for typical users (e.g., TDD patterns and transmit power).

Finally, the control network and its associated IP address space management proved to be difficult to extend. Newly added or relocated sites cannot be seamlessly integrated into the existing IP addressing scheme, which further complicates the management and troubleshooting process.

[3]The IP address ranges are anonymized.

III. Measurements

In the following, we evaluate the performance of our CN based on throughput, latency, and coverage measurements.

A. *HTTP Speedtest*

Online speed tests are widely used as an initial diagnostic tool for network performance. Furthermore, they provide a user-agnostic view on the CN. To minimize outside influences, we host our own LibreSpeed [23] instance that provides a download and upload speed test.

LibreSpeed characterizes user-perceived performance since speed calculations are based on HTTP traffic and throughput. HTTP itself runs on top of the TCP/IP protocol. For the download test, the user requests a large file. Then, LibreSpeed provides the file by generating random data. Random data is necessary to make the file incompressible. Next, the user downloads the file for 20 seconds. Following the transfer, the speed test computes the effective data rate based on the volume of successfully received HTTP payload data. Because the objective is to estimate the steady-state throughput, the calculation excludes TCP/IP start-up effects. To remove these start-up effects, all data received during the first four seconds of the transfer is omitted from the throughput evaluation. The upload test is similar to the download test. The difference is that the user generates data and uploads it to the speed test server. Similarly, evaluation results are based on a 20 s test, where the first four seconds are omitted.

Methodology: The measurement campaign quantifies how throughput degrades with increasing distance from a serving RRH. The UE (iPhone 14) is placed on a table and positioned at fixed horizontal distances from the RRH, ranging from 0 m to 20 m in increments of 2 m. The vertical distance between the RRH and the UE on the table is 2.29 m. Furthermore, the maximum transmission power of the RRH is 18 dBm. At each position, the device performs a speed test run, waits for 5 s, and repeats this procedure until 45 runs are completed.

Results and Evaluation: In Fig. 3, the download and upload application layer rates are presented in the form of violin plots. Additionally, the mean of 40 Secondary Synchronization Signal (SSS) RSRP measurement values is shown for the downlink. These data were measured with a R&S TSMA6, a device that can passively measure 5G radio signals.

From Fig. 3a, it is clearly observable that the experienced download rates and the RSRP degrade with an increasing distance between the UE and the RRH. However, there are measurement points that do not correspond to this trend that require further explanation. Although the average RSRP

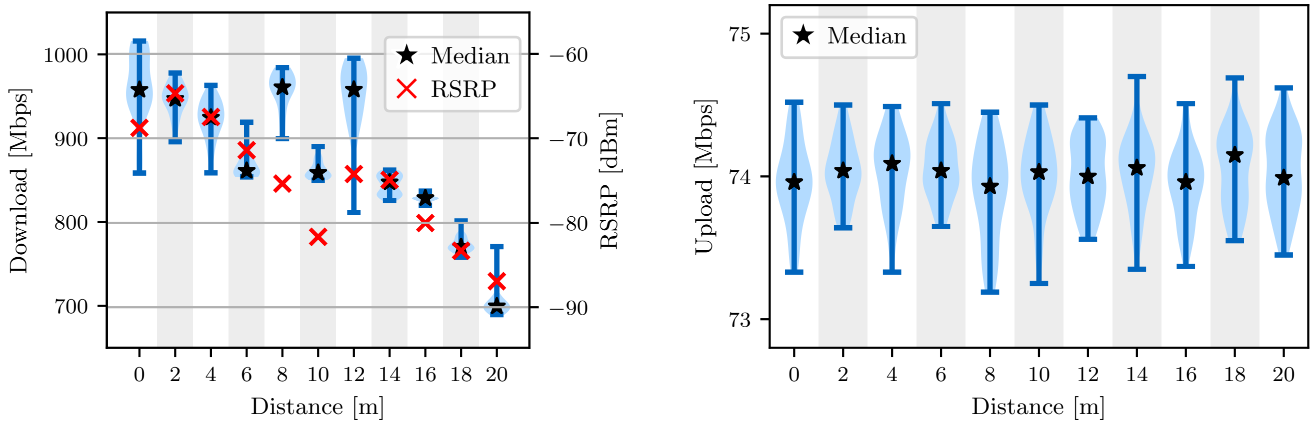

(a) Download speeds together with RSRP average values.

(b) Upload speeds.

Fig. 3: Violin plot of 45 speed test measurements for both the uplink and downlink conducted for horizontal distances between the UE and the RRH ranging from 0 m to 20 m in 2 m steps. For the downlink, the average values of 40 RSRP measurements per UE to RRH distance are also shown.

strictly decreases between the horizontal distances of 6 m and 10 m, the download rates at 8 m are higher than at 6 m and stay roughly the same at 10 m compared to 6 m. This behavior can be explained by a change in the selected Modulation and Coding Scheme (MCS) between measurements. At 6 m, the UE may operate close to the lower threshold of a MCS level, leading to an increased Block Error Rate (BLER) and consequently more retransmissions. At 8 m and 10 m, the RSRP may lie closer to the upper bound of a lower MCS level, resulting in a more stable link with fewer retransmissions. As a result, the throughput at 8 m and 10 m can exceed the throughput at 6 m, despite the lower RSRP values.

Furthermore, the average RSRP does not decrease monotonically with increasing distance. Since these measurements are conducted in a hallway with metallic walls on one side, strong reflections from the walls, floor, and ceiling create multipath components that can interfere either constructively or destructively at different positions along the hallway. As a result, the received signal strength fluctuates rather than decaying smoothly, and peaks or dips in RSRP occur depending on how the reflected paths combine at a given position.

Finally, it can be seen that peak downlink data rates exceeding 1000 Mbps were measured (maximum: 1015.91 Mbps). These values are in line with the theoretical upper bound for the physical layer information bit rate for the configured system settings. Using a 100 MHz bandwidth and a SCS of 30 kHz results in 273 usable Physical Resource Blocks (PRBs). Furthermore, the system is configured with one additional DMRS, leading to 24 Resource Elements (REs) per PRB being reserved for the DMRS. Assuming a code rate and modulation order of 0.926 and 8, respectively (corresponding to MCS 28 in Table 5.1.3.1-2 [24]), and taking into account the MIMO configuration and the 4:1 TDD pattern, the theoretical peak data rate is

$$R = (273 \cdot (168 - 24) \cdot 0.926 \cdot 8 \cdot 4 \cdot (4/5) \cdot 2000)\ \text{bps} = 1863.83\ \text{Mbps}$$

when neglecting all other reference signals and control data beyond the DMRS and assuming perfect channel conditions and error-free packet reception. Once TCP window behavior, congestion control dynamics, and protocol overheads are taken into account, limiting the achievable application layer rate, the observed throughput values are in line with expectations.

The corresponding upload application layer rates over the distance between UE and RRH are shown in Fig. 3b. The results show that the rates are very stable and independent of the distance, with an average median value of 74.02 Mbps. Using a modulation order of 6 (corresponding to MCS 27 in Table 6.1.4.1-1 [24]) and assuming two layers are used by the UE, the theoretical upper bound for the physical layer information bit rate is calculated as

$$R = (273 \cdot (168 - 24) \cdot 0.926 \cdot 6 \cdot 2 \cdot (1/5) \cdot 2000)\ \text{bps} = 174.73\ \text{Mbps},$$

again neglecting all other reference signals and control data beyond the DMRS and assuming error-free communication. Therefore, the measurement results are consistent with expectations and show that the channel offers sufficiently high quality to be fully utilized even at a UE–RRH distance of 20 m.

Header Overhead: Fig. 3a and Fig. 3b show the user-centric application-layer throughput over HTTP, but physical and link-layer rates differ across interfaces due to varying protocol headers. As packets traverse the 5G system, they pass through multiple segments, each adding its own overhead. Fig. 4 summarizes the captured headers and their sizes for a representative speed-test packet.

From the UE to the BS, each packet carries a TCP (32 B) and an IPv4 (20 B) header, adding $\frac{52\,\text{B}}{1398\,\text{B}} = 3.7\%$. Furthermore, additional overhead is introduced by the 5G NR protocol stack (PDCP, RLC, MAC) and by the segmentation and aggregation of packets into transport blocks on this segment. A detailed breakdown of these mechanisms is outside the scope of this discussion. Next, packets travel from the BS to the core server. The BS to core connection contains a special tunnel, the

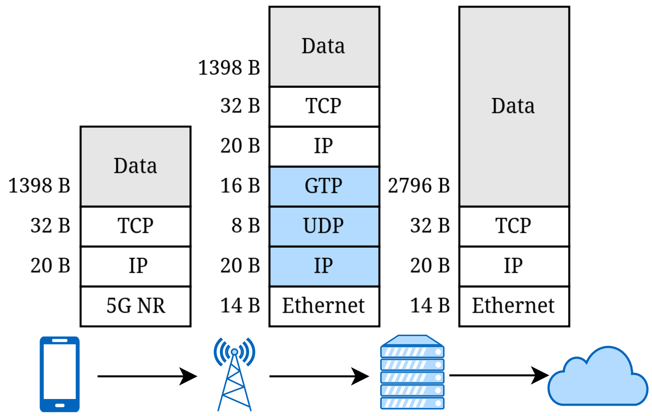

Fig. 4: Overview of a packet's path and the corresponding headers through the physical section from a UE to the internet.

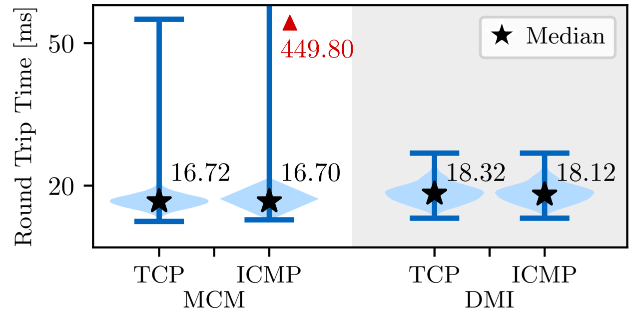

Fig. 5: Violin plot of round-trip delay measurements from UEs located at MCM and DMI to the speed test server and back.

GPRS Tunneling Protocol (GTP) tunnel, which encapsulates the UE packet. All headers involved are shown in Fig. 4 and total 110 B, increasing the data by $\frac{110\,\text{B}}{1398\,\text{B}} = 7.9\%$. In the backhaul uplink, Ethernet (14 B), IP (20 B), and TCP (32 B) headers apply. Due to a larger maximum packet size, multiple UE packets can be aggregated; in the capture, two UE packets were combined, yielding an extra overhead of $\frac{66\,\text{B}}{2796\,\text{B}} = 2.4\%$.

B. Ping Measurements

Latency is a significant factor in evaluating the usability of a network. Since all traffic routes through the core network at MCM, we expect to observe increased latency for traffic from the remote locations. To quantify the latency, we measure the round-trip time for two types of packets. First, the test measures the average delay between a TCP-SYN request and the server's answer for ten requests. Next, the test measures the average latency for ten ICMP echo requests and replies (commonly known as ping). The sequence then restarts with the TCP latency test and runs continuously until stopped.

We show latency results for MCM and DMI in Fig. 5 gathered over a time window of 20 hours. In both locations, we place a UE with the test running. It is observable for both locations that the TCP and the ICMP delays are very similar. Furthermore, we observe that the median delay from DMI with 18.32 ms and 18.12 ms is slightly higher than the delay from MCM with 16.72 ms and 16.70 ms. The additional delay arises from the physical distance and the additional networking hardware between DMI and MCM. The high outlier at MCM comes from a connection error that significantly increased the delay of five ICMP packets.

C. Coverage Tests

To exemplarily evaluate the radio coverage of the outdoor RRHs, we conducted a measurement campaign of the SSS RSRP in the inner courtyard at MCM using the R&S TSMA6[4]. Two RRHs are mounted under the roof of a building at the top of the third floor, with an antenna downtilt of 45°, and a transmission power of 27 dBm. Since GPS data were highly distorted due to reflections from several surrounding (metallic) buildings, we conducted 26 measurements along predefined straight vertical lines spaced 2 m apart. Measurements were taken at walking speed with a sampling rate of approximately 10 Hz, and the receiver antenna was approximately 1 m above the ground. For each measurement, the recorded data were evenly spaced along the corresponding predefined path.

The resulting coverage map is shown in Fig. 6, where grid interpolation is used between the measurement points. The locations of the RRHs are indicated by red dots. We plotted three coverage maps: (i) the RSRP from the left RRH, (ii) the RSRP from the right RRH, and (iii) the best RSRP across both RRHs (combined). The RSRP values were separated by the cell identity, as each RRH serves its own cell.

It is clearly visible that the left RRH provides significantly weaker coverage than the right RRH. Weaker coverage is caused by several large trees that obstruct the line of sight from the left RRH, which considerably attenuate the signal. Due to structural and legal constraints, alternative mounting positions were not feasible. As expected, increasing distance from the RRHs generally results in lower RSRP values. Overall, the measured RSRP ranges between –80 dBm and –110 dBm, with the combined coverage of both RRHs only rarely dropping below –100 dBm, indicating robust signal strength and very good overall coverage.

IV. Future Expansion Plans

As a research-oriented 5G CN, the TUM 5G CN is under continuous development to integrate relevant State of the Art (SotA) technologies and capabilities, thereby enhancing research possibilities toward future cellular network generations. In the near future, two major extensions are planned: (i) the deployment of local breakout UPFs and (ii) the implementation of roaming through an interconnection with the 5G CN of the Technical University of Dresden (TUD).

Since the distance between the core network of our 5G CN and the connected BSs can be relatively large (see Table I), we plan to deploy local breakout UPFs that are co-located with the BSs. To this end, dedicated servers capable of hosting UPFs will be installed at MIRMI, DMI, and GCG. The use of local breakout UPFs enables, for example, investigations of the trade-off between reduced latency and the additional CPU and energy overhead introduced by deploying additional UPF

[4]For reference, the exact geographic coordinates of the inner courtyard are 48°08'56.6"N 11°34'02.2"E.

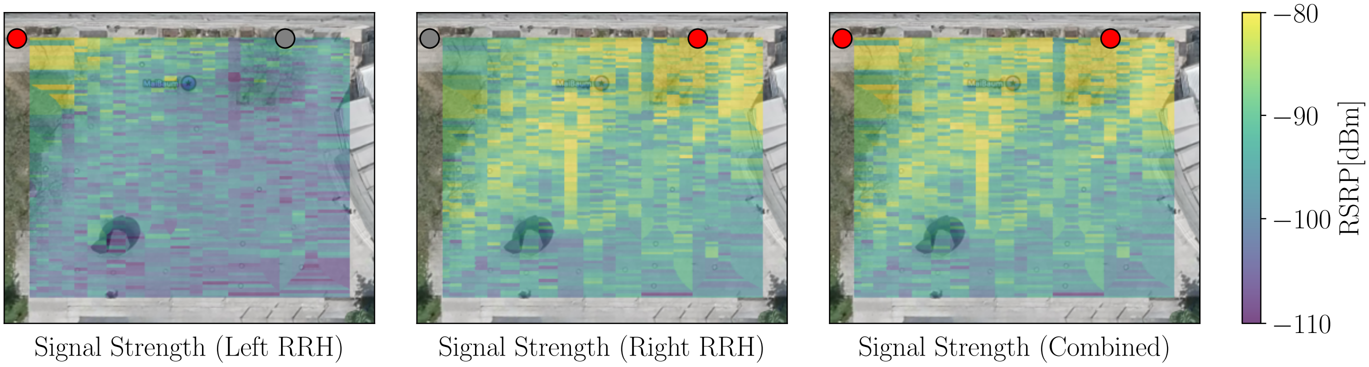

Fig. 6: The coverage map shows the RSRP values of the SSS at MCM's inner courtyard. RRH locations are marked with red dots. The left and middle plots present the signal strength from the left and right RRH, respectively. The right plot shows the maximum received power across both RRHs, i.e., it chooses the maximum value from the left and middle plots.

instances. Furthermore, the inter-working of multiple UPF instances can be explored in detail.

To facilitate research on current and future roaming architectures and functions, we will interconnect parts of our RAN with the core network of a CN located at TUD, and vice versa. This interconnection will be realized via an (already existing) dedicated 100 Gbps fiber link that is connecting TUM with the Deutsche Telekom Chair of Communication Networks at TUD. Specifically, one or more pRRHs of our CN will be connected to the TUD CN, allowing for an inter-Public Land Mobile Network (PLMN) handover, i.e., a transition from the visited network to the home network, while traversing the hallway at MCM. Possible research projects leveraging this setup include measurements of roaming transition latencies and investigations of the overhead introduced by different security protocols in a real-world testbed.

V. Conclusion

This paper presents the TUM 5G Campus Network, a large-scale, multi-site deployment built entirely with commercial products. By illustrating its physical and network architecture, we demonstrated how heterogeneous BS types, diverse control connection paths, and substantial inter-site distances can be handled within a unified 5G system. The performance evaluation indicates that the network achieves high throughput, stable operation, predictable latency, and robust coverage across both indoor and outdoor scenarios. Beyond these quantitative insights, we reported practical lessons learned that highlight real-world challenges often absent from controlled lab environments and elaborated on future expansion plans. Finally, we invite interested researchers to contact us regarding opportunities to conduct real-world 5G experiments using our 5G CN, which is based on commercially available products.

References

[1] M. Emmelmann, M. Corici, F. Eichhorn, M. Hauswirth, and T. Magedanz, "Open 5G campus networks: key drivers for 6G innovations," *Elektrotech. Inftech.*, vol. 139, 2022.

[2] S. Homayouni, M. Paier, G. Stangelmayer, C. Kaipl, C. Sulz, T. Schweeger, and J. Rehak, "Design and development of private 5G standalone network for vertical industries," in *Proc. IWCMC*, 2023.

[3] A. Khalil, B. Becker, L. Wernet, R. Kundel, B. Richerzhagen, T. Meuser, and R. Steinmetz, "Dependability: Enablers in 5G campus networks for industry 4.0," in *Proc. DRCN*, 2023.

[4] S. Senk, S. A. W. Itting, J. Gabriel, C. Lehmann, T. Hoeschele, F. H. P. Fitzek, and M. Reisslein, "5G NSA and SA campus network testbeds for evaluating industrial automation," in *Proc. European Wireless*, 2021.

[5] A. Yaqot, A. Soukieh, R. Schlonsak, A. Hanemann, and H. Hellbrueck, "Installation of multicellular 5G campus network ‘ for remote driving and flying," in *Proc. ICCCNT*, 2024.

[6] D. Jakob, S. Schmidt, and F. Wahl, "Planning and installation of 5G campus networks in hospitals in rural areas and possible use cases: A practical example," in *Advances in Information and Communication*, K. Arai, Ed. Springer Nature Switzerland, 2025.

[7] M. Chepkoech, E.-R. Modroiu, J. Mwangama, M. Corici, and T. Magedanz, "Evaluation of OSS-enabled OpenRAN compliant 5G StandAlone campus networks," in *Proc. ICECET*, 2023.

[8] D. Hamidovic, A. Hadziaganovic, R. Muzaffar, and H.-P. Bernhard, "5G campus network factory floor measurements with varying channel and QoS flow priorities," in *Proc. IECON*, 2023.

[9] F. John, J. Schuljak, L. B. Vosteen, B. Sievers, A. Hanemann, and H. Hellbrück, "A reference deployment of a minimal open-source private industry and campus 5G standalone (SA) system," in *Proc. ICICN*, 2022.

[10] A. Ferguson, U. Pawar, T. Wang, and M. Marina, "Demo: A campus scale private 5G Open RAN testbed," in *Proc. ACM MOBICOM '25*. Association for Computing Machinery, 2025.

[11] K. Mubasier, F. Y. Li, J. A. S. Øgaard, and M.-C. Vochin, "Campus-based full-scale and portable open-source 5G SA networks: Prototyping and experiments," in *Proc. WPMC*, 2023.

[12] OpenAirInterface Software Alliance. (2025) OPENAIRINTERFACE. [Online]. Available: openairinterface.org

[13] Sofware Radio Systems. (2025) srsRAN Project: Open Source RAN. [Online]. Available: srslte.com

[14] LF Projects LLC. (2025) free5GC. [Online]. Available: free5gc.org

[15] S. Lee. (2025) Open5GS. [Online]. Available: open5gs.org

[16] G. Cainelli, L. Underberg, and L. Rauchhaupt, "Influences of logical link design in 5G campus systems," in *Proc. FNWF*, 2022.

[17] A. Fellan, C. Hobelsberger, C. Schellenberger, D. Lindenschmitt, and H. D. Schotten, "Electromagnetic field strength measurements in a private 5g campus network," in *Proc. Q2SWinet*, 2022.

[18] S. B. Mallikarjun, C. Schellenberger, C. Hobelsberger, and H. D. Schotten, "Performance analysis of a private 5g sa campus network," in *Mobile Communication - Technologies and Applications; 26th ITG-Symposium*, 2022.

[19] CampusGenius GmbH. (2025) GeniusCore: The heart of your 5G network. [Online]. Available: campusgenius.com/geniuscore/

[20] ETSI, "5G; system architecture for the 5G system (5GS) (3GPP TS 23.501 version 18.11.0 release 18)," etsi.org, 2025, Tech. Specification.

[21] Docker Inc. (2025) docker. [Online]. Available: docker.com

[22] ETSI, "5G; NR; user equipment (UE) radio transmission and reception; part 1: Range 1 standalone (3GPP TS 38.101-1 version 19.3.1 release 19)," etsi.org, 2025, Tech. Specification.

[23] Federico Dossena. (2025) Librespeed. [Online]. Available: github.com/librespeed/speedtest

[24] ETSI, "5G; NR; physical layer procedures for data (3GPP TS 38.214 version 19.1.0 Release 19)," etsi.org, 2025, Tech. Specification.