Accepted for presentation in: 2024 CIRED Chicago Workshop

RELIABLE 5G NETWORK SLICING FOR DISTRIBUTED ENERGY SYSTEMS: VALIDATION OF LIVE PERFORMANCE DURING HIGH-LOAD EVENT

Dennis OVERBECK
TU Dortmund University – Germany
dennis.overbeck@tu-dortmund.de

Hendrik SCHIPPERS
TU Dortmund University – Germany
hendrik.schippers@tu-dortmund.de

Fabian KURTZ
TU Dortmund University – Germany
fabian.kurtz@tu-dortmund.de

Christian WIETFELD
TU Dortmund University – Germany
christian.wietfeld@tu-dortmund.de

ABSTRACT

The shift towards renewable energy sources is accompanied by increased use of decentralized monitoring and control systems for handling, e.g., distributed energy resources (DERs). In turn, this requires pervasive and reliable communication throughout all levels of the energy grid. Here, traditional approaches need dedicated infrastructures to enable robust and scalable Information and Communication Technologies (ICT). However, the fifth generation of mobile radio networks (5G) enables telecommunication providers to provision virtually dedicated and thus isolated network slices to user equipment, enabling mixed-critical communications of energy devices via existing public infrastructure. Within this work, we present real-world performance measurements of a live network slice during a high-load scenario, thereby illustrating the current state of commercial network slicing solutions as well as providing an outlook on ongoing research.

INTRODUCTION

Transforming the electrical grid into a smart grid relies on sophisticated Information and Communication Technology (ICT) to cope with new grid architectures and resulting data streams. However, to meet service guarantees as required by mission critical use cases such as Distributed Energy Resources (DER) traditionally demands dedicated communication networks. These are in turn associated with high costs and slow rollout cycles, making them unfeasible to support the predicted scale and speed for the rollout of thousands of intelligent energy devices (IEDs) at low voltage level. The fifth generation of mobile radio networks (5G) defines three distinct use cases mapping specific requirements on the communication network. For data-intensive applications such as video surveillance of substations, the Enhanced Mobile Broadband (eMBB) use case is defined, which further focuses on enhanced throughput. However, current communication networks are also faced with large numbers of connected devices, for example, for smart metering or the Internet of Things (IoT), which is covered by the Massive Machine Type Communications (mMTC) profile. Finally, for mission-critical communications dependent on low latency and robustness, the Ultra-Reliable and Low Latency Communications (URLLC) is defined, supporting, e.g., Supervisory Control and Data Acquisition (SCADA) systems, aiming at a 5 ms latency threshold for smart grid communications [1]. Unlike previous generations, the goal of 5G is to provide resources tailored to each use case. To enable different Quality of Service (QoS) requirements co-existing in the same physical communication network, 5G introduces the novel concept of end-to-end network slicing. As shown in Fig. 1, it enables the provision of virtually isolated resources for tenants of a network slice, such as DERs, thus increasing the robustness and reliability of the connections. While hard QoS guarantees can be provided by statically allocating resources there is also the concept of having dynamic network slices, which is currently under research, by utilizing machine learning methods to allocate resources tailored to the needs of the end device according to the current requirements. Similar to dedicated communication networks, slicing provides reliable connectivity for field devices (e.g. IED). Unlike traditional cellular networks with no network slicing, the slice itself can be managed by, e.g., the local utility, yielding direct control while reducing deployment time and costs. As a result, dedicated communication networks become obsolete. However, if there is an abundance of network resources available for the connected devices, i.e., the cell is empty, the network slicing concept does not differ in terms of performance. Hence, the communication network needs to be loaded to prove the capability to enable robust over-the-air communication by utilizing the novel concept. For this purpose, a high-load scenario is chosen to measure performance in a real-world environment.

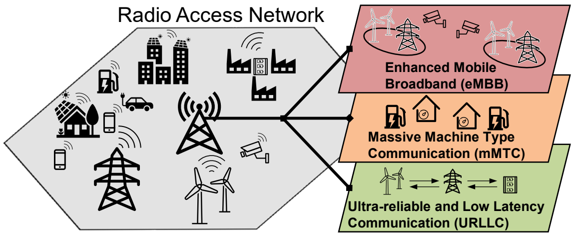

Figure 1: 5G enables virtually isolated resources for distinct service requirements within public ICT infrastructures.

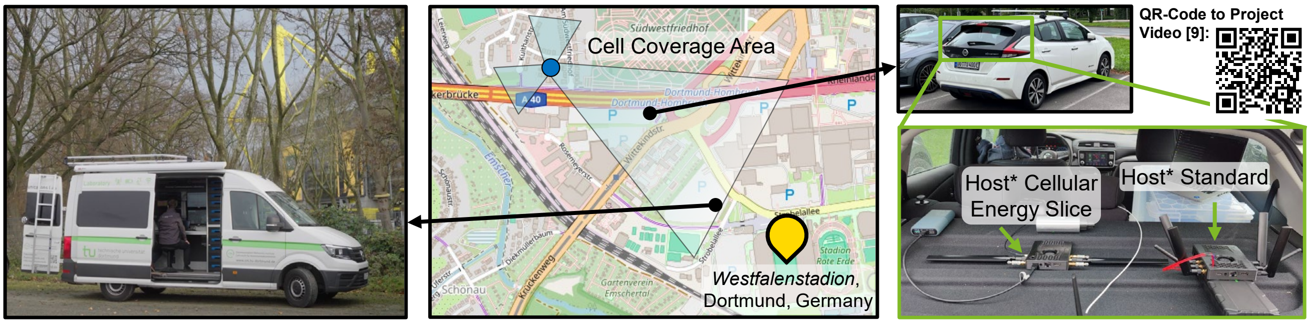

Source: Projektträger Jülich (PTJ)

© Map: OpenStreetMap contributors (openstreetmap.org/copyright)

*Hostsystem based on LattePanda 3 Delta with Quectel RM520N Modem

Figure 2: Depiction of measurement setup for testing network slice performance within high-load scenario. QR Code Link: http://tiny.cc/ReliableNetworkSlicing

In the following, the state of the art for network slicing in the context of energy systems is highlighted. Next, the evaluation scenario and methodology are described in detail. The results are discussed in the subsequent section and an outlook on future research is given.

5G NETWORK SLICING IN PRACTICE

The concept of slicing a communication network into virtually isolated resources for its tenants according to their specific service requirements has been researched for a while [2]. However, the initial focus has been on the wired core network, where the Software-Defined Networking (SDN) and Network Function Virtualization (NFV) technologies provide the foundation. SDN enables more efficient and easily scalable management of communication networks by separating the control and data plane for centralized operation of wide-area communication networks. On the other hand, NFV describes the abstraction and virtualization of specific network functions such as routing or firewall on top of regular Commercial-of-the-Shelf (COTS) hardware, reducing the cost of using proprietary hardware instead. With Release 15, the 3rd Generation Partnership Project (3GPP) introduced the network slicing technology to 5G, building on the aforementioned technologies. The core network now supports the Network Slice Selection Function (NSSF), which provides end-to-end network slicing capabilities. Current research focuses on resource provisioning in the wireless domain, where channel conditions and external influences must be considered for precise resource allocation. This development allows for grid operators to connect energy systems directly via public infrastructure using dedicated network slices. Due to the novelty of this technology, apart from initial field trials such as in [3] [4] or proof of concepts [5], operators are only now implementing the technology into their live networks [6]. As a result, the commercial deployment of network slicing is still being tested, but focusing on the general capability to prioritize data traffic according to the respective network slice in RAN and core. The technology itself will be relevant for the future of communication network (6G) focusing on current trends towards AI-enabled networks [7] and is also seen as promising solution for Smart Grid control [8]. As mentioned in the introduction, the key value of network slicing is highlighted in scenarios where the cell is heavily loaded and resources are not sufficiently available. In contrast to previous publications, this work presents, to our knowledge a first-of-its-kind field trial in Germany that tests the ability of 5G to transmit grid monitoring data over public cellular infrastructure via network slicing. We demonstrate the concept's viability to support energy grid uses cases such as monitoring and control.

SCENARIO DESCRIPTION AND METHODOLOGY

The importance of network slicing as an equivalent alternative to dedicated networks for reliable operation of energy systems is highlighted during periods of high load on the communication networks, i.e., when the amount of available resources is exhausted by the number of connected UEs and thus prioritization is required. This section describes the scenario chosen to provide such an environment in a live communication network, as well as the methodology used to capture the system behavior. Within this work, we compare the performance of a regional network slice in the city of Dortmund, Germany, located close to the sixth-largest soccer stadium in Europe, during a high-load scenario. This network slice is operated by the local utility and thus used in the context of energy management. The measurements were performed during a soccer match with approximately 81,000 visitors in the stadium and more in the stadium area, inducing high loads on the communication network. The experimental setup is highlighted in Fig. 2, which shows the location and hardware setup for the measurement campaign. The cell under test is located close to the stadium and is additionally stressed by the arriving visitor streams. The focus of this work is on the location closer to the base station. The experimental setup is placed in line of sight (LOS) to the base station to ensure a stable baseline in terms of the radio environment. Channel conditions were continuously recorded and averaged -72 dBm Reference Signal Received Power (RSRP), resulting in an average Modulation and Coding Scheme (MCS) index of 22, corresponding to good channel conditions.

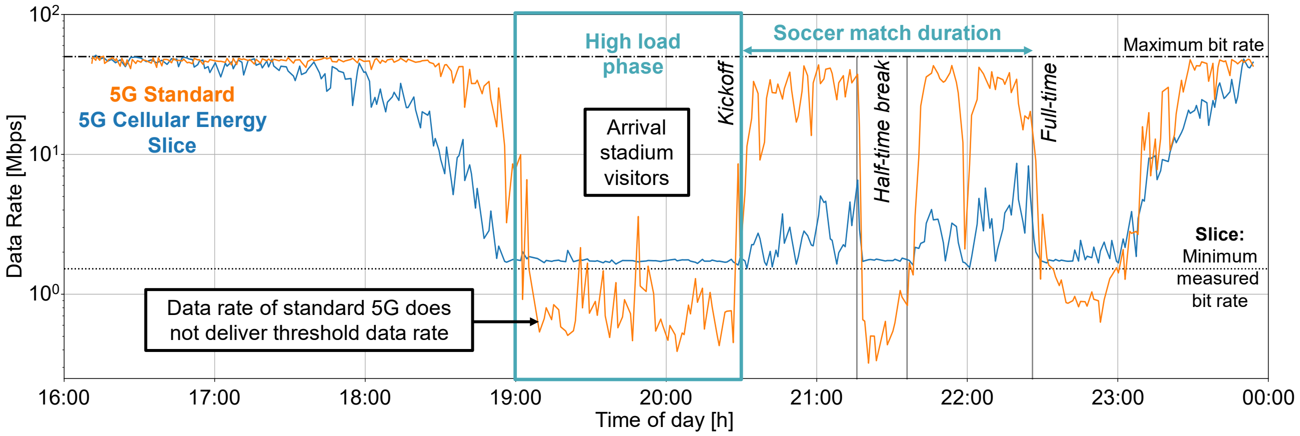

Figure 3: Results demonstrate compliance with service quality guarantees for the device within the 5G slice (blue) in comparison to another in the same cell but without slicing (orange) which violates service guarantees.

To enable a comparison between a slice and a non-slice deployment, we performed the measurements on two devices, one with a standard Subscriber Identity Module (SIM) card from the same operator and the other with a slice-enabled SIM card. The key performance indicators (KPIs) latency and throughput were measured over an eight-hour period, including pre- and post-game visitor flows that stress communications capacity. The cellular energy slice was configured to provide a constant number of physical resource blocks to the connected device. At the same time the standard UE is handled based on a best-effort basis, but with the possibility to switch cells (5G non-standalone with LTE anchor). The measurements were performed on two LattePanda 3 Delta single board computers, each equipped with Quectel RM520N modems, as shown in Fig. 2 on the right. The board computers were powered by mobile battery packs to ensure continuous monitoring of the communication network performance. As previously described, the standard UE is capable of connecting to 5G non-standalone cells, potentially delaying throughput degradation by switching to less congested cells with LTE anchors. However, the modem does not use other fallback technologies such as standard LTE and 2G. The cellular energy slice UE is set up to connect only to 5G standalone, meaning that the only cell to connect to, is the one outlined on the map in the center of Fig. 2, where the slice technology is available. The device is connected in Single Input Single Output (SISO) mode. An adjustable attenuator is used to test the effect of channel degradation on the sliced UE. The network slice is configured to provide at least 1 % of the cell capacity to the connected tenants of the slice. The cell contains a bandwidth of 70 MHz, this results in 189 available Physical Resource Blocks (PRBs). Therefore, at least 1 PRB is reserved for the slice in each time slot (1 ms). The cell is configured to operate in Time-Division Duplex (TDD), which is a standard 5G configuration for most mobile network operators (MNOs) in the operated frequency band. During the measurements both directions (downlink and uplink) were tested. The methodology and measurement sequence are summarized in Fig. 4. We focus on the KPIs latency and throughput. The former relies on ping measurements based on Internet Control Message Protocol (ICMP) packets and Two-Way Active Measurement Protocol (TWAMP). The latter is measured using iperf3 v3.11. To minimize the impact of load balancing of public Domain Name System (DNS) servers, both KPIs are also measured against a server at the backend of the TU Dortmund University network. Additionally, a server close to the MNO's peering node is used as endpoint, thus the latency measurements mirror an optimal location close to the core network. We focus the findings on these servers. Each measurement cycle varies in its length to complete as waiting times are included for finishing processes before starting the next one to mitigate the impact of unfinished processes in the background. The end-to-end latency measurements concentrate on the results achieved with the TWAMP protocol, as it allows for more realistic packet behavior, i.e., closer to real application traffic regarding payloads. The measurements comprised different payload sizes, however this work concentrates on the results achieved with 60 Byte payloads. The throughput measurements are based on the User Datagram Protocol (UDP). The test cycles are performed simultaneously on both devices (slice and non-slice), so the experienced cell load is equal for both settings.

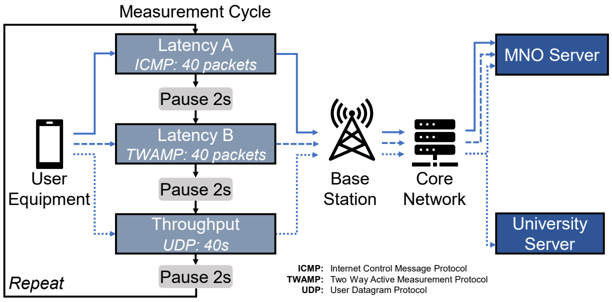

Figure 4: The measurement methodology employed in the live network utilizes looping test runs and includes different communication protocols as well as server endpoints.

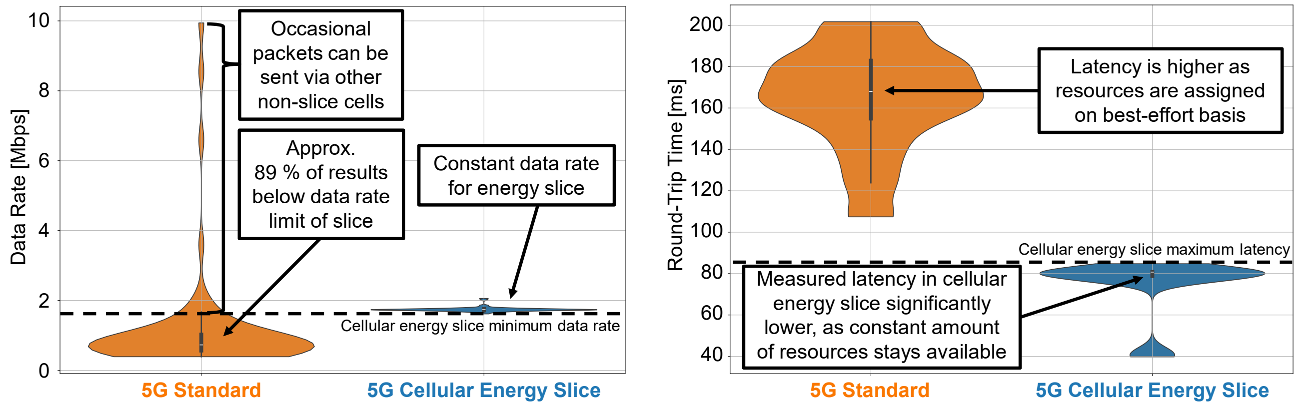

Figure 5: Achieved data rates (left) and latency (right) for the cellular energy slice and a non-slice device with focus on the pre-match high-load phase between 19:00-20:30h (c.f. blue box in Fig. 3).

MEASURED RESULTS AND DISCUSSION

This section presents the results of the measurement campaign, focusing on the behavior of the communication network during the high-load scenario. However, the focus of the measurement is only indirectly on the data rate or latency. Rather, the focus is on the consistency of the transmission in order to guarantee stable transmission of necessary data. Accordingly, a minimum threshold is defined which must not be undercut. Fig. 3 shows the achieved data rates for the standard UE (orange) and the cellular energy slice UE (blue) over an eight-hour period, highlighting the impact of visitor flows on the cell capacity. The data rate is plotted on a logarithmic scale to highlight the underperformance of the standard UE during periods of high load in comparison to the sliced UE. Both UEs are limited to a maximum throughput of 50 Mbps. Due to the single cell connection of the cellular slice and the resulting lack of ability to switch to less congested cells or connection technologies (such as LTE), the UE experiences throughput degradation earlier than the standard UE. However, during the high load phase just before the soccer match starts, between 19:00h and 20:30h, the cellular energy slice remains at a steady level with at least 1.61 Mbps. Therefore, the throughput is sufficient, e.g., to provide reliable transmission of monitoring data. In contrast, the standard UE experiences drops below this threshold, resulting in packet loss and possible violation of QoS guarantees. The mean data rate is at 1.22 Mbps with a minimum of 0.39 Mbps. As a result, the standard UE is not only able to provide such guarantees during this phase. It only recovers during the game, when most visitors typically do not load the cell with data applications. This behavior can also be observed during the half-time break, when the cellular energy slice remains at the guaranteed bitrate level, while the standard UE fails to deliver. The resulting KPIs are shown and summarized in Fig. 5, focusing on the high-load phase, where the communication network cell is highly congested. On the left-hand side, the figure focuses the achieved data rates for the cellular energy slice (blue) and standard UE (orange).

As foreshadowed in Fig. 3., the slice UE on the right-hand side delivers a mean data rate of approximately 1.74 Mbps over this time period, thus providing sufficient throughput for monitoring and controlling the connected energy devices. While the standard UE on the left can occasionally achieve a higher data rate, it remains below the average data rate of the cellular energy slice UE in 89 % of the measured result points. This results in unreliable transmission, e.g., in the form of wait times induced by delayed resource allocation This is also highlighted on the right-hand side of Fig. 5, where the latency is focused. The cellular energy slice is able to deliver the packets within an acceptable time, while the standard UE experiences latency that is twice as high which can result in important messages being received too late. This behavior can be explained by the constant amount of resources available to the sliced UE, while the standard UE is assigned on a best effort basis; if packets are not scheduled, the latency peaks due to increased waiting times. The cellular energy slice achieves a mean end-to-end latency of 74 ms, and a maximum at 84 ms, which is marked as reference for the comparison with the standard UE. Minimum latency is achieved, when the cellular energy slice UE is able to use more than the one guaranteed PRB to transmit. Contrary, the standard UE experiences higher end-to-end latency with a mean at 165 ms and a peak of 201 ms, which is accompanied by larger variations in the arrival time, decreasing the overall QoS of the connection. It is shown that the cellular energy slice UE outperforms the standard UE in terms of constant throughput and low jitter. However, the static allocation can be improved by utilizing methods of machine learning to dynamically allocate resources to network slices according to dynamic environments. These requirements may vary due to higher throughput demands or challenging and varying channel conditions. The importance of the latter is also highlighted by initial tests performed with the adjustable attenuator. By manually reducing the channel quality, the performance of the sliced UE decreased correspondingly, as the minimum PRB has to be modulated more robustly, thus reducing the available bits.

CONCLUSION AND OUTLOOK

This work highlights the benefits of 5G network slicing in the context of a real-world deployment, demonstrating reliable connectivity even during extreme communication network loads. Slicing is shown to be a flexible and cost-efficient approach for monitoring grid devices over existing public telecommunication infrastructures. The presented measurements were conducted in the context of a research project with, among others, utilities and network operators aiming for increased efficiency and reliability in grid control by utilizing machine learning methods (cf. [9]). The results show that the cellular energy slice guarantees a minimum bit rate for connected devices, which a non-slice communication network does not provide. Additionally, the latency remains stable compared to the non-slice UE, which enables expectable transmission rates. However, dynamic allocation of resources corresponding to current requirements of each network slice can be pursued and is currently being researched (cf. [10]). Furthermore, future work will concentrate on the simulation-based evaluation of predictive allocation schemes to provide dynamic network slices for energy use cases in high-load scenarios derived from real-world behavior.

ACKNOWLEDGEMENT

This work has been partly funded by the German Federal Ministry of Education and Research (BMBF) in the course of the 6GEM Research Hub under grant no. 16KISK038, and is supported by the Federal Ministry for Economic Affairs and Climate Action (BMWK) via the project ReCoDE under funding reference 03EI6093B as well as by the Ministry of Economic Affairs, Industry, Climate Action, and Energy of the State of North Rhine-Westphalia (MWIKE NRW) along with the Competence Center 5G.NRW under grant no. 005-01903-0047.

REFERENCES

[1] 3GPP, „System Architecture for the 5G System (5GS),“ 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 23.501, Jul 2022, Version 17.5.0.

[2] X. Foukas, G. Patounas, A. Elmokashfi und M. K. Marina, „Network Slicing in 5G: Survey and Challenges,“ in *IEEE Communications Magazine*, May 2017, pp. 94-100.

[3] P. Rost, M. Breitbach, H. Roreger, B. Erman, C. Mannweiler, R. Miller und I. Viering, „Customized Industrial Networks: Network Slicing Trial at Hamburg Seaport,“ in *IEEE Wireless Communications*, 2018, pp. 48-55.

[4] Y. Luo, M. Jiang, D. Zhang und F. Effenberger, „Field Trial of Network Slicing in 5G and PON-Enabled Industrial Networks,“ in *IEEE Wireless Communications*, 2023, pp. 78-85.

[5] Ericsson, „Ericsson and Telefónica redefine consumer experiences with on-demand network slicing,“ 23.02.2024 [Online]. Available: https://www.ericsson.com/en/press-releases/3/2024/ericsson-and-telefonica-redefine-consumer-experiences-with-on-demand-network-slicing. [Accessed on 07.06.2024].

[6] D. Jones, „Jio boasts 5G network slices, considers 5G private network uses,“ 21.03.2024. [Online]. Available: https://www.fierce-network.com/5g/jio-has-5g-network-slices-considering-5g-private-network-uses. [Accessed on 07.06.2024].

[7] ITU-R, „Framework and overall objectives of the future development of IMT for 2030 and beyond,“ International Telecommunication Union - Radiocommunication Sector, Tech. Rep. Recommendation ITU-R M.2160-0, 2023.

[8] M. O. Nait Belaid, V. Audebert und B. Deneuville, „Designing a 5G based Smart Distribution Grid Protection System,“ in *CIRED 2021 - The 26th International Conference and Exhibition on Electricity Distribution*, Online Conference, 2021, pp. 1241-1245.

[9] D. Overbeck, F. Kurtz, S. Böcker und C. Wietfeld, „Design of a 5G Network Slicing Architecture for Mixed-Critical Services in Cellular Energy Systems,“ in *IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, Singapore, 2022, pp. 90-95.

[10] D. Overbeck, N. A. Wagner, R. Wiebusch, J. Püttschneider, T. Faulwasser und C. Wietfeld, „Data-Driven Proactive Uplink Slicing enabling Real-Time Control within an Open RAN Testbed,“ in *IEEE Infocom Wkshps: International Workshop on Computer and Networking Experimental Research using Testbeds (CNERT)*, Vancouver, 2024.