Accepted for presentation in: 2024 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids: SmartGridComm, Oslo, Norway, September 2024.

Experimental 5G Platform for Managing Mixed-Critical Traffic using Network Slicing in Concentrated Solar Power Plants

Robin Wiebusch, Pascal Jörke and Christian Wietfeld
Communication Networks Institute
TU Dortmund University, Germany
Email: {Robin.Wiebusch, Pascal.Joerke, Christian.Wietfeld}@tu-dortmund.de

***Abstract*—As global energy needs shift towards sustainable sources, Concentrated Solar Power (CSP) plants are increasingly relevant due to their ability to convert sunlight into electrical energy using vast fields of motorized heliostats. However, the effective management of CSP operations presents significant communication challenges, especially in coordinating the precise and rapid adjustments needed for heliostats under varying environmental conditions and integrating data-intensive Unmanned Aerial Vehicle (UAV) systems for plant calibration and monitoring. In this paper, we implement a predictive traffic steering strategy based on Machine Learning (ML) and Network Slicing to optimize resource allocation in a mixed-critical power plant communication environment, comprising a critical Machine Type Communication (cMTC) heliostat control service and an enhanced Mobile Broadband (eMBB) high-quality calibration camera stream. Evaluated in a 5G testbed with real-world channel conditions, our results demonstrate that the developed traffic optimizations can ensure 100 % service reliability for critical applications, even during severe network congestion.**

***Index Terms*—5G Campus Networks, Scalability, Mixed-Critical, Network Slicing, Concentrated Solar Power Plants**

I. Introduction

The demand for clean energy production leads to the development and optimization of green power plants. As a promising solution for electrical power generation CSP plants have gained new attention in recent years. Using motorized mirrors, called heliostats, sunlight is reflected over an area of several km^2 to a local receiver tower, which transforms thermal energy into electrical energy [1]. Due to the natural movement of the sun, heliostats are frequently repositioned. Clouds introduce additional challenges since the partial opacity of heliostat fields disrupts the even illumination of the light receiver and produces critical local temperature differences.

Tracking heliostats requires fast and reliable communication networks between a central control system and each field unit. Currently, the transmission is performed by wired networks, which makes building new CSP plants expensive due to the groundwork for up to hundreds or thousands of individual heliostats. Additionally, new services such as automated UAV-based calibration of heliostats require wireless high data rate communication throughout the heliostat field. This paper evaluates 5G campus networks for CSP power plants as a holistic solution for field communication including heliostats and calibration UAVs. The mixed critical scenarios for massive Machine Type Communication (mMTC) and cMTC as well as eMBB applications are based on definitions from the research project 5hine, which analyzes 5G as a communication solution for solar tower plants [2]. Although the data size of the heliostat communication is small, it is critical, nevertheless since it is responsible for the safe operation of the plant (ref. Fig. 1). To ensure precise heliostat orientation, periodic calibration is required. For efficient and frequent calibration automated UAVs will be used in 5G enabled CSP plants [2]. The UAV-based calibration requires high data rate communication to transmit a high-quality video stream besides UAV control and status updates. Unlike heliostats, UAV communication is not critical since the calibration procedure can be paused at any point, provided that UAV control is still ensured.

In cellular mobile communication, several techniques exist to address the challenges of mixed-critical traffic. In this work, we propose an implementation of Network Slicing in combination with predictive traffic steering, optimizing network performance and reliability. Network Slicing allows for the segregation of the network architecture to cater to both critical and non-critical communications. Predictive Traffic Steering dynamically manages traffic flows, enhancing the responsiveness and adaptability of the network to varying

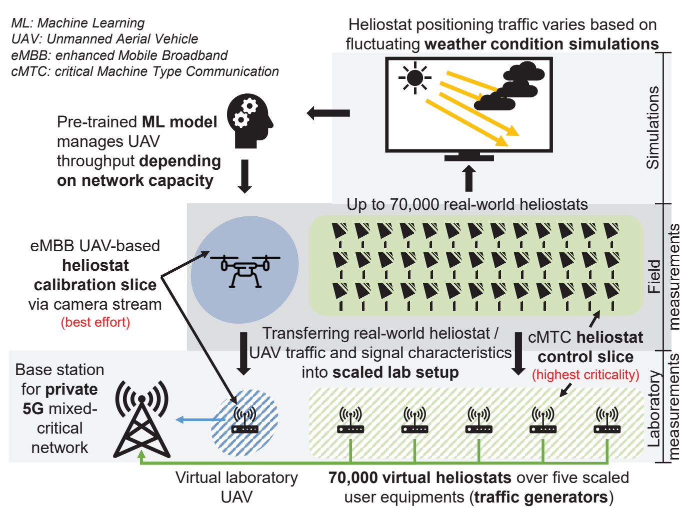

Fig. 1: Multi-dimensional evaluation setup for mixed-critical communication links in lab and field measurements.

operational demands. Together, these strategies ensure robust network performance across all user applications. In Fig. 1 the key contributions of this work are illustrated. Initially, we conducted field measurements at a real-world CSP plant to model the communication channel for heliostat and UAV devices. From a simulation environment, we then collected data on fluctuating cloud movements and the corresponding heliostat traffic to train a ML model, enabling us to predict future congestions and determine optimal times to throttle the UAV video stream. We then utilized the pre-trained ML model and the established channel model to translate the CSP cell conditions into a scaled lab setup. Here, we performed measurements using 70,000 virtual heliostats across five User Equipments (UEs), testing and evaluating the network and our traffic management methods in a mixed-critical scenario.

The sections of this work are structured as follows: First, an overview of related work is given in Sec. II. Then, Sec. III describes the considered CSP applications and the mixed-critical scenario. Afterward, in Sec. IV, the deployed technologies to manage and optimize the traffic for the given scenario are presented. In Sec. V, we then discuss the evaluation of these technologies for the measured network Key Performance Indicators (KPIs). Finally, Sec. VI closes the paper with a conclusion and an outlook on future work.

II. Related Work

Different wireless communication solutions were previously evaluated for CSP plants. In [3], a mesh-based approach was evaluated, using IEEE 802.15.4 devices for multi-hop communication in the heliostat field. In a first field test with a limited number of 93 devices, a packet error rate between 1 % and 20 % was observed due to the lack of acknowledgments, as acknowledged transmissions would overload the network. This demonstrates that the IEEE 802.15.4 mesh-based approach is limited for high-scaled heliostat fields. Another approach using IEEE 802.15.4 mesh networks is discussed in [4]. Employing simulations with 340 heliostats, a Packet Delivery Rate (PDR) of 93 % for nodes close to the base station and 88% at the edge of the field were obtained, which still offers insufficient reliability. For critical communication, the network was extended with an additional unicast communication link, using sub-1-GHz frequencies for better signal ranges and, thus, reliable emergency communication. The results demonstrate that networks lacking centralized scheduling and resource management are not well suited for reliable communication of critical heliostat fields. In [5], the performance of private high-density microcells using Narrowband Internet of Things (NB-IoT) as an alternative 5G communication solution optimized for small data transmissions is evaluated. The results demonstrate that these networks can be massively scaled to meet the requirements for monitoring and controlling heliostat fields. Still, with NB-IoT, only small data transmissions with low data rates are available, rendering UAV-based broadband communication, as required in modern CSP plants, impractical. To evaluate the utilization of Network Slicing for mixed-critical applications, the authors of [6] deploy an open 4G/5G

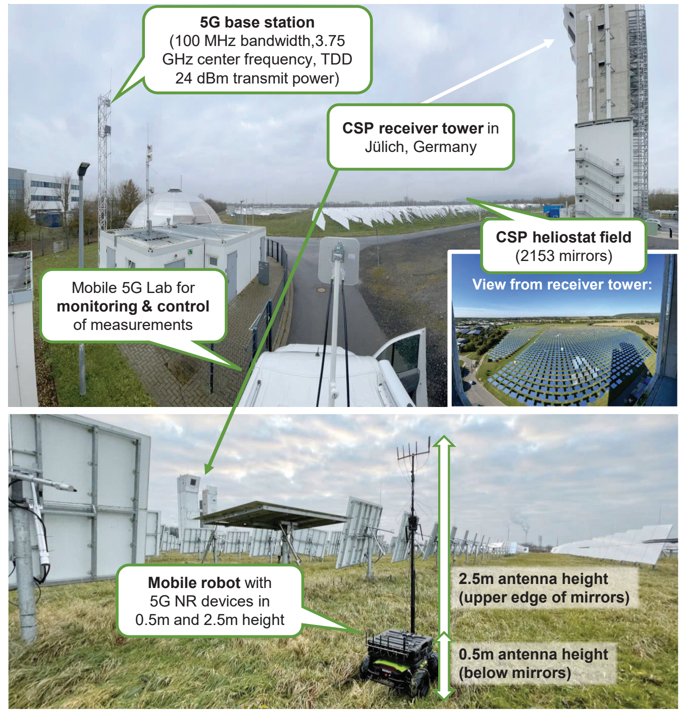

(a) Field measurement setup in CSP plant Jülich, Germany.

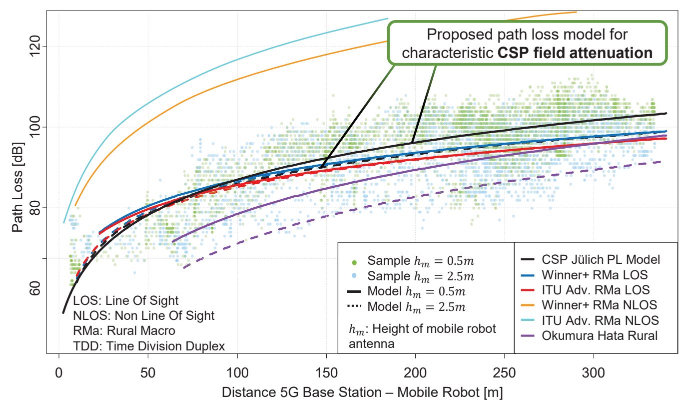

(b) Measurement results to derive channel models for heliostat fields.

Fig. 2: Real-world 5G measurements in a CSP plant for evaluating heliostat-specific attenuation and MCS characteristics.

radio testbed, integrating a modified Media Access Control (MAC) layer scheduler for low latency communication. The results demonstrate that the data rate and latency of higher-priority slices can be stabilized under network strain. However, due to the limited bandwidth of 5 MHz and low packet rates, the scalability of the setup is not further explored. An alternative to 5G communication systems is presented in [7] with OpenWiFi as a solution for flexible deployments in industrial use cases, which can be adapted to different frequency bands, especially for remote regions typical for CSP plants. In [8] and [9], the authors discuss the utilization of ML for Dynamic Adaptive Streaming over HTTP (DASH) to proactively steer network traffic for varying network conditions. Both studies effectively show that by using predictive approaches, video segments can be prefetched at the mobile edge, enhancing overall video quality. However, in both instances, the video server is either a large media server or a content delivery network, which does not account for scenarios where the video server must produce a real-time camera stream while being constrained in terms of CPU and memory capacity.

III. Mixed Critical Applications for CSP Plants

5G networks for CSP plants allow heterogeneous applications, ranging from eMBB for video streams from field calibration UAVs to mMTC/cMTC for monitoring and controlling heliostat fields. This section describes the modeled data traffic and channel conditions based on a real-world scenario.

A. Heliostat and UAV Application Models

Modern CSP plants range from hundreds of heliostats to up to 70,000 heliostats in a single CSP plant [10], which is used in this work as an evaluation scenario. Each heliostat that received a new aim point from Downlink (DL) broadcast transmissions transmits an individual status update with an application payload of 120 bytes to the central control unit. Since modern heliostats automatically track the natural sun movement, only diverging aim point switches are transmitted, for instance when clouds are shading parts of the heliostat field. For the setup in this paper, simulation results for moving cloud passages over a reference CSP plant and the associated aim point switches of individual heliostats from [2] are used, which can be assumed as a worst-case scenario due to the assumed fast movement of clouds. With more extreme weather conditions heliostats will drive in a protective position and do not require additional aim point switches. The traffic from the reference plant with 6482 heliostats is linear scaled to the 70,000 heliostats considered in this work. For this setup, it is assumed that the heliostats publish their status updates in an interval of 10 s. This results in an aggregated maximum throughput of 6.72 Mbps of the heliostat service. Missing updates are interpreted as misaligned heliostats and can result in emergency plant shutdowns preventing damage from overheating. Thus, status updates of the heliostats are critical for the plant [2] and require a reliability of 99.9 %.

In addition to the heliostat traffic, a UAVs transmits data for field calibration and control. A high-definition video stream using an H.264 codec with a data rate of 8 Mbps is used for calibration of misaligned heliostats. Since this service is less critical, the calibration is interrupted when network capacity is limited and the UAV data rate is reduced to 1 Mbps for a low-definition video stream. This renders the stream robust enough to still maintain manual control of the drone. It is assumed, that a service reliability of 99.9 % for the video stream is sufficient given a robust codec is chosen [11].

B. CSP-specific Path Loss

To reproduce the specific wireless signal attenuation in heliostat fields with their large metal frames and metal-coated mirrors we performed signal quality field measurements in a CSP plant in Jülich, Germany. A private 5G New Radio (NR) campus network with a center frequency f of 3.75 GHz was set up to provide 5G coverage to the heliostat field with 2000 heliostats as depicted in Fig. 2a. A mobile robot providing different 5G antenna heights was used for comprehensive field measurements. Fig. 2b presents the signal strength samples from the field measurements. Since none of the available Path Loss (PL) models fit the measured characteristic, a new empirical model for a CSP-specific PL is derived:

$$PL = 34.36 \log_{10}(d) + 7.83 h_m - 4.0 \log_{10}(d) h_m + 17.59 \quad (1)$$

with PL as the path loss in [dB], d as the distance between the 5G base station and the heliostat / mobile robot in [m], and h_m as the height of the mobile robot antenna in [m].

Considering a radius of 1590 m for concentric CSP plants with 70,000 heliostats [10], the PL model from Eq. 1 and the uplink Modulation and Coding Scheme (MCS) samples from the measurements in Fig. 2b we can derive a typical MCS of 9 to be used in deeper parts of large-scale heliostat fields when the base station transmit power is set to 20 dBm to limit the private cell size. Therefore, for our lab measurements, the utilized maximum MCS is 9 for a worst-case analysis.

IV. Applied 5G Technologies for Mixed-Critical Applications

This section outlines the technologies used to achieve high network reliability in mixed-critical applications. First, we employ Network Slicing to stabilize the highly critical heliostat traffic during periods of network congestion. Additionally, we deploy a predictive throttling technique as a way to proactively reduce the bitrate of the UAV camera stream, ensuring the reliable operation of the drone control.

A. Uplink Network Slicing in the Air Interface

Network Slicing stands as a pivotal technology in 5G networks, enabling the logical division of a single physical network into multiple virtual networks, each designed to meet specific service requirements. For instance, in applications that manage large-scale critical infrastructure communications (cMTC) alongside less critical data streams (eMBB), Network Slicing ensures that essential services maintain high performance and reliability, even under significant network loads. In this work, we adopt a modified Round Robin (RR)-based Network Slicing strategy, tailored to scenarios that require hard service guarantees and fixed prioritization between services. In a RR scheduling mechanism that lacks internal weighting among clients, the Next Generation Node B (gNB) scheduler allocates resources fairly among all UEs via the Downlink Control Information (DCI). As a result, all connected UEs receive an equal share of resources, leading to an implicit over-prioritization of services comprising multiple clients over those with a single client. Consequently, the scheduling priority of a service is given by $P_{\text{service}} = N_{\text{service}}/N_{\text{total}}$, where N represents the number of connected UEs. In our experimental setup we further increase the critical service priority by enforcing that resource allocation is based on the order in which UEs are registered in the 5G Core (5GC). Specifically, a UE registered earlier (UE A) will always receive all requested resources, while a subsequent UE (UE B) only receives the remaining resources. This modification ensures that Service A consistently receives priority over Service B independently of the network traffic, safeguarding critical communications.

B. Traffic Management via Predictive Bitrate Adaptation

Traffic management in the form of Adaptive Bitrate Streaming (ABR) plays a vital role in service quality assurance by dynamically adjusting the quality of video streams to match current network conditions, and optimizing bandwidth usage while maintaining the stability of critical services. Conventional ABR technologies, like DASH, work by encoding multiple streams with different bitrates simultaneously on the server side. Clients then select the optimal bitrate based on their current network conditions to ensure a consistently smooth video stream. While effective, the concurrent encoding of multiple streams requires substantial server-side CPU resources. In scenarios like ours, where UAVs must remain lightweight with limited processing power, multiple concurrent encodings are not feasible. Instead, our approach involves the host proactively adapting the bitrate of a single stream depending on current weather conditions and network congestion. Implementing proactive bitrate decisions during real-time camera streams requires external, proactive throttle commands. For this work, we designed and evaluated a ML-based approach to manage this process. A supervised learning method is employed, utilizing historical data that includes the fluctuating proportion of shaded heliostats at the power plant and the resulting share of heliostats that communicate in the uplink to maintain the stability of the power plant. The challenge is framed as a time-series forecasting problem, with the model tasked to predict the proportion of transmitting heliostats for the forthcoming control interval based on data from previous intervals. Forecasted network congestions in upcoming control intervals can then be utilized to proactively throttle the UAV's video stream. The dataset is generated using a simulation environment as discussed in [2], where a CSP with 6482 heliostats is assumed. It comprises 358 simulated data points, each representing the weather conditions and the resulting communication effort of the heliostats within a 10 s control interval. To predict the proportion of transmitting heliostats at time frame t_0, the input features include the share of shaded heliostats from t_{-w} to t_0, and the proportion of transmitting heliostats from t_{-w} to t_{-1}, with window size w.

To train and evaluate the model, the set is divided into 80 % training data and 20 % testing data. Using the scikit-learn Python library [12], we deployed and compared the following model types: Random Forest, Gradient Boosting, and Support Vector Machine. The model's hyperparameters are optimized using grid search in combination with cross-validation, optimizing the Mean Absolute Error (MAE) on the training set. The best performance is achieved using the Random Forest model with a window size of 5, 100 base estimators and a maximum tree depth of 10 as optimal hyperparameters. The final Random Forest achieves an MAE of 8.81 on the test data set.

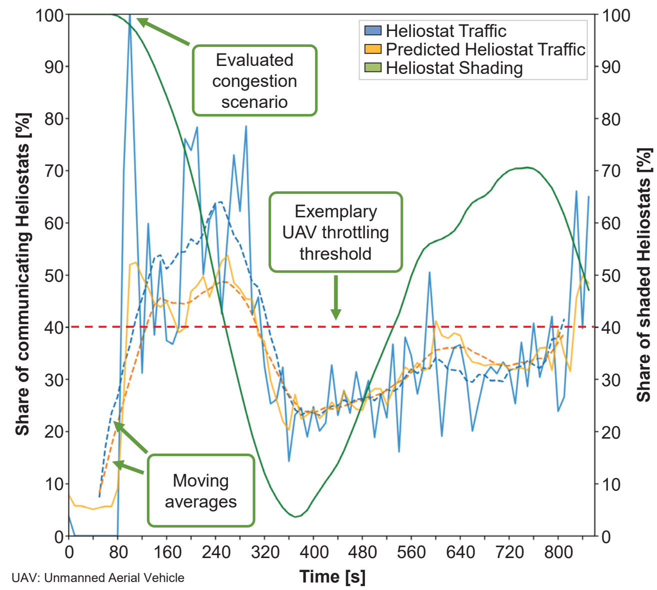

Fig. 3: Test data set to evaluate the Random Forest model.

Fig. 3 illustrates the fluctuating share of shaded heliostats in [%] (green, right axis) and the resulting heliostat traffic, measured as a proportion of the total number of heliostats in [%] (blue, left axis). The prediction of this traffic (yellow, left axis) is utilized to send proactive throttle commands to the UAV during periods of network congestion. Though the model's performance might vary with different training and/or test datasets, the significant peak of 100% within the test data provides a challenging scenario to ensure a robust evaluation. The plot demonstrates that the ML model can effectively predict the trend of heliostat uplink traffic, thereby validating its use as a trigger for proactive bitrate adaptation based on a chosen threshold as depicted by the dashed red line.

V. Laboratory Measurement Setup and Results

This section outlines the measurement setup, evaluates the baseline parameters, and assesses the performance of the optimization techniques used.

A. Experimental Setup

For evaluating the discussed traffic optimization technologies, we deploy a 5G radio network. Focusing on traffic optimization rather than maximizing network throughput or robustness via a commercial 5G solution, we utilize open-source components. This approach allows us to extend the existing resource scheduler, thereby enhancing the experimental capabilities of our testbed. The 5GC is hosted using open5GS [13], while the gNB is based on the srsRAN Project [14]. We implemented the Network Slicing by modifying the RR Scheduler in the MAC layer of srsRAN to prioritize scheduling based on the order of UE attachment, as discussed in Sec. IV-A. Both the 5GC and gNB are hosted on a shared server (AMD Ryzen 7735U, 32 GB RAM, Ubuntu 22.04.4). The radio frontend of the gNB is deployed using a Universal Software Radio Peripheral (USRP) B210 Software Defined Radio (SDR). The

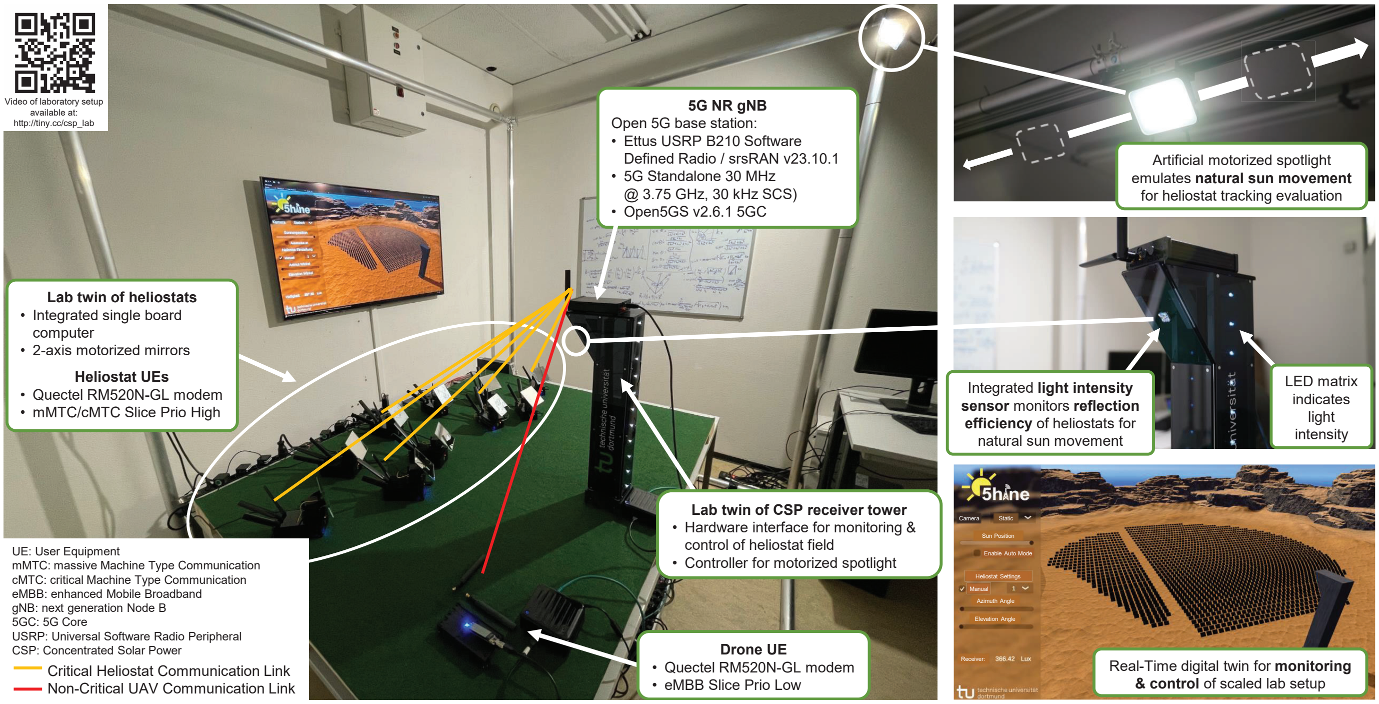

Fig. 4: Experimental setup consisting of an open 5G base station at the receiver tower, heliostats and UAV UEs.

cell operates in Time Division Duplex (TDD) mode at a center frequency of 3.75 GHz, with a 30 MHz bandwidth and 30 kHz subcarrier spacing. To emphasize uplink traffic, the following TDD pattern is utilized: DDDSUUUUUU. The UAV UE operates on equivalent hardware as the gNB, while the heliostat services are deployed on integrated single board computers (ARM Cortex-A72, 2 GB RAM, Ubuntu 22.04.3). All UEs are equipped with Quectel RM520N-GL 5G radio modems. To ensure reliable and accurate measurements, all UEs are synchronized with the gNB using chrony, leveraging the Network Time Protocol (NTP) protocol for precise time synchronization. The maximum MCS in both Uplink (UL) and DL is set to 9 to match the network conditions described in Sec. III. Fig. 4 illustrates the measurement setup, which is integrated in a CSP laboratory demonstrator. The demonstrator includes ten 2-axis motorized heliostat models equipped with 5G modems as well as a receiver tower model to visualize communication and control procedures in CSP plants. The heliostats are aligned to a light-emitting motorized spotlight (Fig. 4 top right) and reflect light to the receiver tower, which is equipped with a light intensity sensor to monitor and assess the alignment of each heliostat. The demonstrator is controlled using a digital twin as shown in Fig. 4 downright.

B. Data Generation and Performance Evaluation

To effectively evaluate the experimental scenario, we require an application capable of generating highly scalable and yet precise data rates. For this purpose, the developed application layer client is designed to scale through adjustments in packet size and the inter-arrival time of packets. The application's main loop operates dual-threaded to manage packet generation and transmission via First In First Out (FIFO) buffers. Each application data packet is structured to contain a Unix timestamp marking the time of packet generation and an identifier for the host UE. The remainder of each packet is filled with dummy data to achieve the desired packet size. Thread A is responsible for generating these application layer data packets. It operates in 100 ms batches, producing the required number of packets and storing them in the application buffer. This buffer is virtually unlimited in size, constrained only by the host's RAM capacity. Once Thread A completes the packet generation for a batch, it enters a sleep state for the remainder of the 100 ms period. Thread B continuously monitors the application layer FIFO buffer. Upon detecting new packets, it transfers them to the underlying User Datagram Protocol (UDP) buffer, where they are prepared for transmission over the srsRAN radio stack. If the UDP send buffer reaches full capacity, thread B pauses, ensuring that no packets have to be dropped on application layer. This method, however, can lead to continuously increasing packet latencies over time if the buffer remains full. At the receiver side, packets are collected and assessed post-transmission. As each new packet arrives, its arrival Unix timestamp is recorded. This data, combined with the packet's generation timestamp, enables the calculation of throughputs and delays at the application layer by calculating the difference between the generation and reception times.

C. Evaluated Key Performance Indicators

To assess the performance of the network and the effectiveness of the deployed traffic management techniques, three application layer network KPIs are measured for each service at the receiver side: achieved throughput in [Mbps], UL one-way packet delay in [ms], and service reliability in [%]. Service reliability is defined as the ratio of the number of

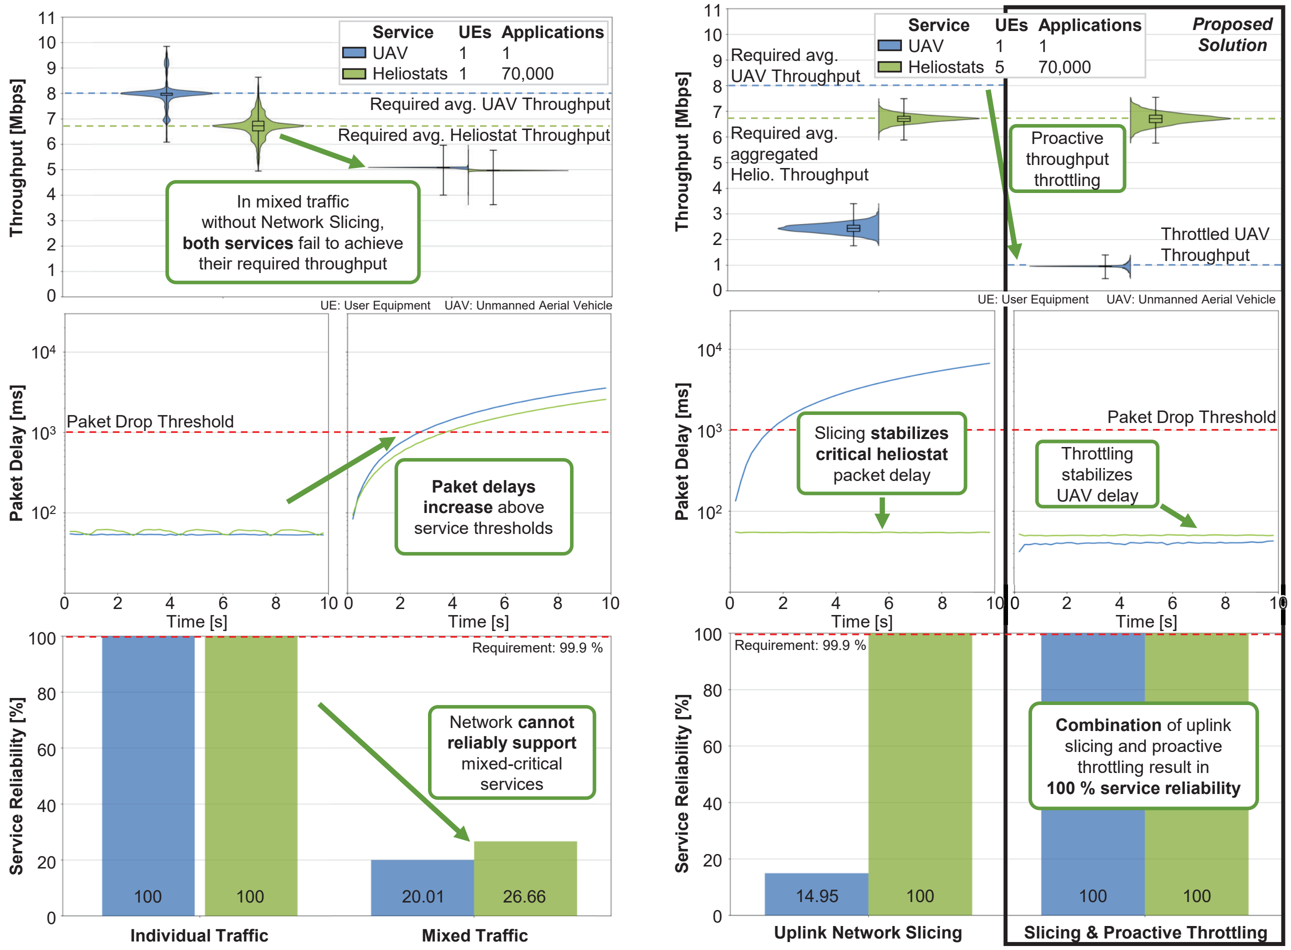

(a) Established network baseline without the use of traffic optimization technologies demonstrates that in cases of network congestion, both services become nearly unavailable.

(b) Network Slicing ensures 100 % reliability for the prioritized heliostat traffic. Proactive UAV throttling stabilizes the video stream for UAV control even under worst-case conditions.

Fig. 5: Evaluation of the deployed traffic optimization techniques compared to baseline 5G laboratory measurements.

successfully transmitted packets N_{success} to the total number of transmitted packets N_{total} in percent as expressed by the formula $Reliability = N_{success}/N_{total} \cdot 100\,\%$. A packet delivery is considered successful if it reaches the receiver's application layer with a delay of one second or less since packets with larger delays cannot be utilized by the service. While reliability is crucial for the final assessment of service quality, the other KPIs provide valuable insights into the dynamics of the network. For each evaluated configuration, 100 independent 10 s measurement runs, generating approx. 7.5 million data packets, are performed. All collected data packets are aggregated into 200 ms batches using their timestamps to calculate average throughputs, delays, and reliabilities within these windows. Furthermore, to avoid inconsistencies during the startup and shutdown of the processes, only time slots where all UEs are in active state are considered and data points from the first and last 200 ms batches of each measurement are discarded. First, a baseline network performance is measured without traffic optimization and a single UE per service. Then, the performance of a scaled scenario with 5 heliostat UEs with Network Slicing and predictive traffic management is assessed.

D. Network Baseline Performance

To establish a baseline for the network performance, we initially measure the previously defined KPIs for each service individually and then assess them in a mixed traffic scenario. The results are depicted in Fig. 5a. From top to bottom, the figure illustrates the achieved KPIs using a violin plot for the throughput, a line plot throughout a 10 s interval for the average packet delay at each time point, and a bar plot for the service reliability. In the left column, the results for the separated services are presented, whereas the right column displays the results for mixed traffic. Starting at the top, it is demonstrated that both services, when operating individually, on average achieve their required throughput of 8 Mbps and 6.72 Mbps respectively. Since occasional drops in throughput are compensated by following transmission intervals, packet delays of both services remain within a stable range at approx. 55 ms. These delays are induced by both the radio stack and computational overheads. As a result, the service reliability of both services is above their requirements. In the mixed traffic scenario, the median throughputs decline to 5.08 Mbps and 4.97 Mbps. This indicates that the RR scheduler treats both services with equal priority since only one UE is considered for the heliostats. As both services underperform relative to their throughput requirements, packet delays significantly increase above the paket drop threshold, peaking at 5.03 s and 3.7 s. Despite the throughputs reduced by less than 50 %, service reliability decreases to 20.01 % and 26.66 %, as successful packet deliveries are only possible at the start of the transmission interval and UDP send buffers quickly fill up. This demonstrates that, without Network Slicing or traffic steering, congestion compromises the availability of both critical heliostat control and UAV calibration services.

E. Results of the Applied Traffic Optimizations

Now that the baseline has been established, we proceed to investigate the performance of the optimization techniques. The same KPIs are measured across individual 10 s transmission intervals as previously described. Fig. 5b, structured similarly to before, displays these measurement results. This time, the columns represent scenarios where only Network Slicing is applied (left), and a combination of Network Slicing and UAV throttling is applied (right). Starting in the top left, the plot demonstrates that Network Slicing effectively stabilizes the higher-prioritized heliostat traffic at a median rate of 6.72 Mbps. As a result, the UAV throughput significantly underperforms relative to its requirements, reaching 2.44 Mbps. As illustrated in the delay plot, only the UAV packet latency increases above the threshold, reaching up to 7.41 s. Subsequently, employing Network Slicing alone allows the heliostat service reliability to reach 100 %, though at the expense of the UAV service, which suffers from an even lower reliability of 14.95 %. The right column examines the combined effect of slicing and throttling. Here, the UAV transmission is proactively reduced to 1 Mbps to maintain sufficient reliability for UAV control. This results in the effective median throughputs of 0.96 Mbps for the UAV and 6.72 Mbps for the aggregated heliostats. Due to proactive UAV throttling, the packet delays of both services remain stable at about 40.53 ms and 50.53 ms respectively. Ultimately, the service reliability plot confirms that the combined application of these technologies successfully maintains operations of 100 % for both services, even under severe network constraints.

VI. Conclusion

In this paper, we developed and integrated a proactive traffic steering mechanism based on 5G Network Slicing and a Random Forest ML model. To validate our approach, we established an open testbed reflecting real-world configurations and channel conditions found in a CSP plant. In a mixed-critical communication scenario, where a cMTC heliostat control service competes for resources against a less-critical UAV-based calibration camera stream, we demonstrated the necessity of our deployed optimizations by showing that during periods of high congestion, packet delays escalate up to 5.03 s, resulting in service reliabilities as low as 20.01 %. By predicting congestive intervals in advance using historical weather data, we can proactively throttle the camera stream. In combination with Network Slicing, prioritizing the highly critical heliostat control service, we managed to stabilize delays at around 45 ms, achieving 100 % reliability for the heliostat control and the UAV-based camera stream. In future work, we plan to integrate ML predictions to explore a dynamic scenario with fluctuating weather conditions, placing greater emphasis on the real-time prediction performance of the model.

Acknowledgment

This work has been carried out in the course of the projects 5hine and Competence Center 5G.NRW, funded by means of the Federal State NRW by the Ministry for Economic Affairs, Industry, Climate Action and Energy (MWIKE) under the Funding ID 005-2108-0073 resp. 005-01903-0047. The authors thank the student project group HEAT as well as David Ronschka and Selin Öz for their contributions to the experimental setup.

References

[1] M. Liu *et al.*, "Review on concentrating solar power plants and new developments in high temperature thermal energy storage technologies," *Renewable and Sustainable Energy Reviews*, 2016.

[2] P. Schwarzbözl, I. Miadowicz, D. Maldonado Quinto, J. Golembiewski, P. Jörke, T. Faulwasser, and C. Wietfeld, "5G as Communication Platform for Solar Tower Plants," in *International Conference on Concentrating Solar Power and Chemical Energy Systems, SolarPACES 2023*, 2023.

[3] S. Kubisch, M. Randt, R. Buck, A. Pfahl, and S. Unterschütz, "Wireless heliostat and control system for large self-powered heliostat fields," in *17th International Conference on Concentrating Solar Power and Chemical Energy Systems, SolarPACES 2011*, 09 2011.

[4] A. Pfahl, M. Randt, F. Meier, M. Zaschke, C. Geurts, and M. Buselmeier, "A Holistic Approach for Low Cost Heliostat Fields," *Energy Procedia*, vol. 69, pp. 178–187, 2015, international Conference on Concentrating Solar Power and Chemical Energy Systems, SolarPACES 2014. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1876610215003276

[5] P. Jörke, D. Ronschka, and C. Wietfeld, "Performance Evaluation of Random Access for Small Data Transmissions in Highly Dense Public and Private NB-IoT Networks," in *2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*, 2023, pp. 1–7.

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

[7] C. Arendt, S. Fricke, S. Böcker, and C. Wietfeld, "Distributed Performance Evaluation of 5G and Wi-Fi for Private Industrial Networks," in *IEEE Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, 2024.

[8] R. Behravesh, A. Rao, D. F. Perez-Ramirez, D. Harutyunyan, R. Riggio, and M. Boman, "Machine Learning at the Mobile Edge: The Case of Dynamic Adaptive Streaming Over HTTP (DASH)," *IEEE Transactions on Network and Service Management*, vol. 19, no. 4, 2022.

[9] J. Uriol, I. Yeregui, A. Gabilondo, R. Viola, P. Angueira, and J. Montalban, "Context-Aware Adaptive Prefetching for DASH Streaming over 5G Networks," in *2023 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*, 2023, pp. 1–6.

[10] N. R. E. Laboratory, "Noor Energy 1 / DEWA IV - 100MW tower segment CSP Project," https://solarpaces.nrel.gov/project/noor-energy-1-dewa-iv-100mw-tower-segment, accessed: April 18, 2024.

[11] 3GPP, "System architecture for the 5G System (5GS)," 3rd Generation Partnership Project, Technical Specification (TS) 23.501, Jun. 2024, version 19.0.0.

[12] F. Pedregosa *et al.*, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[13] "open5GS v2.6.1," 2024. [Online]. Available: https://open5gs.org/

[14] "srsRAN Project v23.10.1," 2024. [Online]. Available: https://docs.srsran.com/projects/project/en/latest/