Better Safe Than Sorry: Distributed Testbed for Performance Evaluation of Private Networks

Christian Arendt, Stefan Böcker, Caner Bektas, Christian Wietfeld TU Dortmund University, Germany Communication Networks Institute (CNI) Email: {christian.arendt, stefan.boecker, caner.bektas, christian.wietfeld}@tu-dortmund.de

Abstract—The operation of local private cellular radio networks in licensed frequency bands constitutes one of the core innovations of current 5G and future 6G networks. The predicted performance of an exclusively utilizable and thus interferencefree frequency range of, for example, private 5G networks, so-called campus networks, are generally of great interest to industrial companies. However, the integration of 5G network infrastructures into existing brownfield environments has to overcome major technological and administrative challenges. In contrast to very rarely encountered greenfield scenarios, the potential of 5G can only be realized in practice if significant performance advantages over existing wireless network infrastructures (e.g., Wi-Fi) can be demonstrated while guaranteeing seamless integration into the process landscape.

To this end, this paper presents an agile system for continuous, cross-network monitoring of end-to-end guarantees in terms of throughput, latency and reliability. While single pointwise measurements during network deployment often indicate the expected performance peaks, this contribution specifically explores the potential of a spatially distributed stress test that actively monitors network quality on a continuous basis. An extensive case study is conducted to demonstrate the performance of the distributed approach for performance evaluation of multi-user and cell edge environments. In addition, it is illustrated that the distributed system can be used to estimate how mission-critical service guarantees affect the overall network performance.

Index Terms—Private 5G,Campus Networks, Distributed Testbed, Performance Evaluation

I. INTRODUCTION

The 5th Generation of mobile networks (5G) is a major topic in current research and development for a good reason. In addition to enabling more sophisticated mobile network access for many people around the world, especially the possibility to implement private networks in specific rented frequency bands and therefore empower companies and public institutions to control and configure their own network marks a fundamental change in the world of mobile networks, closing the gap between mobile networks and wireless local area networks [1]. Future manufacturing processes will become more and more modular and flexible, which poses a need for reliable wireless communications for industrial applications. These come with high requirements and therefore need deliberate testing before actually being deployed in production. This is especially true for the interdependency of various applications, as currently many proof of concept tests consist of a limited set of parallel processes, which does not reflect reality.

We therefore extended the Spatially distributed Traffic and Interference Generation (STING) framework proposed in [2] to enable technology-independent, multi-network stress testing under real world conditions with multiple distributed test devices. This enables the integration of STING into a private campus network infrastructure as depicted in Fig. 1, in order to instantly take action from network stress test results. As described in [3], this can either be done by reconfiguration of the network or by adaption of the underlying applications, e.g. a production schedule.



Fig. 1. Integration of distributed testbed in network and operating environments for bilateral system and application control.

The structure of this work is as follows: Section II provides a brief summary of relevant and prior publications. Section III discusses the proposed system architecture and a section IV depicts a stress-test methodology for experimental performance evaluation in various network infrastructures. Section V presents an exemplary test case and the results of this analysis. Section VI brings this work to a conclusion.

II. RELATED WORK

The concept of private 5G infrastructure is discussed in [1]. The author also briefly describes use cases and their requirements on the communication infrastructure. Performance testing of 5G infrastructures for industrial applications has been addressed by the 5G Alliance for Connected Industries and Automation (5G-ACIA) in [4]. The authors provide procedures and parameters for reproducible performance tests for industrial communications. The 5G-ACIA also discusses traffic models for industrial network testing, as shown in [5]. Evaluation of a networks performance goes hand in hand with on demand planning, especially considering private campus networks, where performance testing can be used to evaluate and optimize network configurations for specific use cases. This is addressed in [6]. The need for stress testing private networks in order to ensure Service Level Agreement (SLA) fulfillment is pointed towards in [7], as violations can result in high costs depending on the application scenario. This is especially true for mission-critical Machine-Type Communication (MTC) applications in future networks, as depicted in [8]. First measurement studies have been conducted by the authors of [9] and [10] which both show promising performance of private 5G campus networks for industrial applications, but also derive optimization potential especially regarding achievable latency. The basis of this work has been laid in [11], where distributed measurements in the public Long Term Evolution (LTE) network have been conducted to ensure reliable communication of smart meter devices. In [2], this concept has been adapted to private networks where we introduced the STING framework and presented a robotic test case to assess the resilience to interference as an exemplary use case.

III. ARCHITECTURE OF MULTI-TECHNOLOGY TESTBED FOR PRIVATE NETWORKS

This section gives a brief overview of the underlying architecture of the STING system for technology-independent stress-testing. As shown in Fig. 2, the STING system is made up of two primary components: a central management and control system and distributed end devices for traffic generation.

These components are set up in a modular approach, currently enabling both 5G and Wi-Fi connectivity. The central management and control unit contains a server which serves as the counterpart for the end devices regarding the generated network traffic. Additionally, the server contains a database to persist and evaluate the generated performance results during



Fig. 2. Overview of STING system architecture

testing. The STING end devices consist of an embedded PC which is specialized on network connectivity supporting different types of network interfaces. In this setup, an Intel AX200 Wi-Fi 6 card and a Quectel RM500Q 5G modem are used. The devices allow flexible deployment in various lab or production setups, providing performance and stress-test capabilities in order to evaluate the fit of both technologies for different use cases. The two technologies can both be used as test networks as well as control networks. If a 5G test is performed, Wi-Fi can act as the control channel and vice versa. This prevents the tested technology from interference by the system's own control traffic, while still allowing live performance monitoring. After conducting a test scenario, the end devices provide their performance results to the central server for evaluation. The test process also used in the case study presented later is described in the following section.



Fig. 3. Technology-independent stress-testing for private networks using Spatially Distributed Traffic and Interference Generation (STING)

IV. METHODOLOGY FOR EMPIRICAL PERFORMANCE EVALUATION

The general test procedure is depicted in Fig. 3. It uses the well-established iperf traffic generator with configurations for TCP (iperf3) and UDP (iperf2) traffic. The procedure consists of two parts, one focusing on single- and one on multi-user tests. While these concepts can also be conducted in public networks, the controlled environment of private networks allows an extensive analysis of these effects without influence of unknown users and applications as well as network configurations.

The different test cases are described in the following subsections.

A. Single-User: Receiver Sensitivity Test

The receiver sensitivity test is conducted with one end device using a fully meshed channel matrix (Mini-Circuits ZTMN-0995A-S) which allows applying attenuation to every possible path between the antennas and the end device modem, as well as a shielding box (Rohde & Schwarz CMW-Z10) to make sure that there is no leakage effects. This setup is schematically shown in Fig. 4.



Fig. 4. Schematic overview of experimental setup for receiver sensitivity tests. This setup can be combined with multi-user testbed scenarios.

In this setup, the same procedure as in the health tests is performed repeatedly, with iteratively increasing programmed attenuation. Therefore, this tests allows to analyze the performance of the User Equipment (UE) at the cell edge.

B. Single-User: Performance Test

The single user performance tests have two main purposes: Ensure that all end-devices are operational and give a first estimation of the networks capacity for dimensioning of the more sophisticated tests. For this procedure, all deployed nodes sequentially generate TCP traffic in uplink and downlink direction, followed by UDP traffic for both directions. The traffic level should be high enough to touch the channels maximum capacity for one user, indicating the maximum system capacity. This value is then used for dimensioning traffic in the following test sequences.

C. Multi-User: Long Term Stability

This procedure aims to analyze the influence of multiple users transmitting in the same network. Therefore, the overall system data rate is kept constant and is distributed over all 16 STING end devices. This allows to observe the networks ability to maintain a given network load over a defined time period. As with the health tests, the experiment is repeated for TCP and UDP as well as uplink and downlink, respectively. These different configurations are performed with the system data rate set to 20%, 50%, 80% and 95% of the maximum achievable throughput of the single user performance tests.

D. Multi-User: Scalability Limits

Similar to the previous test, this procedure aims to analyze multi-user effects, but this time with constant end device data rates, resulting in increasing system data rate with increasing number of devices. The data rates used should be dimensioned in a way that allows the system to overflow its capacity limit to observe how the system handles congested conditions.

V. CASE STUDY: MEASUREMENT RESULTS FOR 5G STANDALONE PRIVATE 5G NETWORK

The aforementioned system and methodology is used in a lab environment with a commercial Ericsson Private 5G (EP5G) Standalone (SA) system, using the 3.7 GHz frequency band which can be licensed in Germany for dedicated private networks operating in Time Division Duplex (TDD) mode. Fig. 5 gives a schematic overview of the testbed. The key element of the deployment are the distributed STING devices which are deployed in the lab environment on traverse elements. The facility used for the deployment spans an area of approx. 10 m by 15 m and is generously provisioned with two 5G radios connected to a single core and forming a combined cell, ensuring that all end devices can be seen as being located in the cell center. Additionally to the 5G infrastructure, a Wi-Fi system is deployed which serves as the control link for the distributed end devices.



Fig. 5. Experimental test setup at playground in laboratory environment at TU Dortmund University

The configuration of the network as well as key parameters for the tests and the corresponding simulation, described in section V-A, are listed in Tab. I

TABLE I PARAMETER SET FOR THE TESTBEDS 5G SYSTEM AND THE RESPECTIVE SIMULATION.

Parameter	Testbed	Simulation
Operation mode	SA	NSA
Bandwidth	$50\mathrm{MHz}$	
Center frequency	$3.775\mathrm{GHz}$	
Maximum #devices	16	
TDD pattern	DDSU	
Special slot	10:2:2	dynamic
Tx power base station	$250\mathrm{mW}$	
Base station MIMO capability	4x4	
UE MIMO capability	2x2	
Time per test configuration	$10 \min$	$300\mathrm{s}$
Identified system capacity	DL: 500 Mbit/s, UL: 90 Mbit/s	
UE data rate for scalability tests	DL: 100 l	Mbit/s, UL: 20 Mbit/s

A. Validation of Case Study using ns-3

To confirm the results obtained in this testbed, the same procedures have been reproduced using the ns-3 based simulation tool 5G-LENA [12]. While the tool currently uses Non-Standalone (NSA), performance results are comparable due to the high load situation of the UEs. This results in the small amount of layer 3 Radio Resource Control (RRC) control traffic in comparison not having a big influence whether it is transmitted within the 5G spectrum or outside of it. For simplification, the simulation used one base station configured with an omnidirectional antenna and all end devices in close line of sight proximity to the base station. Results of the simulation are shown in the multi-user results (Fig. 8 and 9) and show a good match with the testbed performance, with the exception of a slight tendency towards the downlink in the simulation. This is due to the fact that the TDD pattern in the simulation is handled in a way that does not define the special slot in a static manner.

B. Single-User Tests

Starting with the single-user tests, Fig. 6 show the performance of a node under decreasing channel conditions. Using the channel matrix as described in section IV-A, we applied additional attenuation to the transmission paths before the UEs antennas.

It can be seen that the achievable data rate decreases with higher attenuation, as one would expect. This is due to the system choosing a more robust Modulation and Coding Scheme (MCS) in the process to cope with decreasing channel quality. Upon reaching an additional attenuation of 57 dB the connection was lost. The chosen MCS in also shown in Fig. 6. Naturally, in downlink direction the system allows a higher modulation, using the 256-QAM modulation table in [13, Tab. 5.1.3.1-2] instead of the 64-QAM table [13, Tab. 6.1.4.1-1] which is used instead for the uplink modulation. This can be seen in the lower parts of Fig. 6 and is also partly responsible for the lower throughput achieved in uplink direction compared to downlink.

The second single-user test is the performance test deducted sequentially with all used end devices. Those results are depicted in Fig. 7.

Fig. 7 shows the throughput reached for both TCP and UDP protocols and uplink as well as downlink direction for all STING end devices. The peak data rate achieved are slightly above 500 Mbit/s in downlink direction and around 90 Mbit/s in uplink direction, respectively. These rates are used as a baseline for the following multi-user tests. The right part of Fig. 7 depicts the result of a ping test for all STING end devices, indicating a typical low Round Trip Time (RTT) of below 18 ms on average.

C. Multi-User Tests

After testing for individual UE performance and ensuring an operational state of all STING end devices, the multi-user tests described in section IV are conducted. Fig. 8 shows the results for the long-term stability tests. The plots at the top show the



Fig. 6. Receiver sensitivity results indicate (1) validation of our distributed testbed and (2) performance range for cell center to cell edge.



Fig. 7. Single user performance tests illustrate comparable channel conditions (cell center) for all nodes and similar performance values for both TCP and UDP

individual throughput every STING end device achieves, and the plots below show the aggregated system throughput of all 16 STING end devices. From here on, only UDP results are shown due to space restrictions, however TCP results all show a very similar behavior.

It can be seen that all devices achieve a very stable individual throughput throughout all states of target system utilization. Even approaching the overall system capacity up to 95 % does not decrease the overall performance even with a constant load over multiple minutes. The behavior of the system going beyond its maximum capacity can be observed in Fig. 9. In this figure, again on the top the individual throughput and at the bottom the aggregated system throughput is shown. The x-axis however depicts the number of active STING end devices, all with a constant user data traffic of 100 Mbit/s in downlink direction and 20 Mbit/s in uplink direction, respectively.

With two and four active STING end devices, the behavior is similar to Fig. 8, as all devices achieve their target throughput and result in an aggregated system throughput below the systems overall capacity. When increasing the number of active devices however, the system's capacity is reached and exceeded. This naturally results in a decreased throughput of the individual STINGs. Looking at the error bars, it can be seen that the throughput of all devices is equally reduced, and no device is starved completely. For Fig. 10, the concepts of the receiver sensitivity test have been combined with a multi-user test case. One STING end device is placed in the shielding box connected to the channel matrix described in section IV-A, the other 15 devices are active without additional attenuation. The channel matrix is configured with an additional attenuation of 52 dB, resulting in a cell edge situation for the attenuated device. In this configuration, all nodes are set up to produce a total target utilization of 80 % of the system capacity. As seen in Fig. 10, this results in a much higher throughput for the cell center nodes compared to the attenuated cell edge node due to its lower channel quality and therefore lower MCS. This behavior is expected in a fair, nonprioritizing system. However, especially with mission-critical applications, a prioritization of specific nodes at the celledge can be favorable. Modern 5G networks allow different approaches to accomplish this, from more simple device specific prioritization to complex network slicing approaches. In this work, a device specific prioritization is applied to the node at the cell edge in order to enforce its target throughput. As Fig. 10 illustrates, this comes at the cost of the overall system performance, due to the cell edge device needing more resources than the cell center ones to achieve the same throughput. In this example, the cell edge device can achieve its target throughput at the cost of approx. 50 % of the overall system throughput, resulting in a trade-off for the operator between specific individual device performance and overall throughput of the system.



Fig. 8. Long term UDP multi-user test procedures with all 16 distributed STING nodes in parallel result in reliable fulfillment of the target system utilization from low system utilizations (20%) up to near full system load (95%).



Fig. 9. Constant high datarate per STING node leads to stepwise approach to the system capacity. At maximum system utilization, a gradual decrease of individual performance is observed.

VI. CONCLUSION & OUTLOOK

This work presents a distributed testbed for performance evaluation of private networks. The modular architecture of the utilized STING system allows stress testing for different technologies, while this work focuses on a commercial private 5G system. A methodology for systematic testing has been proposed and executed in a case study within a lab environment of TU Dortmund University. Exemplary results of this study show a good stable operation in multi-user scenarios when the system utilization is kept below the maximum network capacity. Whenever the maximum capacity is approached and exceeded, performance degradation is evenly split over all end devices. Multi-user tests with a dedicated cell edge user depict a trade-off between the cell edge device's performance and the overall network performance, where the latter is severely decreased when the cell edge device is prioritized.

Future enhancements of the test bed include the integration of the highly performant millimeter wave technology introduced with 5G. This extension comes with research challenges with regards to spatial diversity and beamforming measures. Additionally, the direct integration of real application traces to enable realistic test scenarios is under current development. This enables fine grained testing for very specific private network needs, especially when comparing private 5G with industrial Wi-Fi technologies.



Fig. 10. Heterogeneous deployment scenarios indicate a trade-off between prioritizing edge nodes and keeping overall performance

ACKNOWLEDGMENT

This work has been partly funded by the Ministry of Economic Affairs, Industry, Climate Protection and Energy of the State of North Rhine-Westphalia (MWIKE NRW) along with the 5Guarantee project under grant number 005-2008-0077 and supported by the Competence Center 5G.NRW under grant number 005-01903-0047 and the project Plan&Play under the grant number 005-2008-0047, as well as the German Federal Ministry of Education and Research (BMBF) in the course of the 6GEM research hub under grant number 16KISK038 and 6G-ANNA under grant number 16KISK101.

REFERENCES

- A. Aijaz, "Private 5G: The Future of Industrial Wireless," *IEEE Industrial Electronics Magazine*, vol. 14, no. 4, pp. 136–145, Dec. 2020.
- [2] C. Arendt, M. Patchou, S. Böcker, J. Tiemann, and C. Wietfeld, "Pushing the limits: Resilience testing for mission-critical machine-type communication," in 2021 IEEE 94th Vehicular Technology Conference (VTC2021-Fall), 2021, pp. 01–06.
- [3] M. Müller, J.-M. Knorr, D. Behnke, C. Arendt, S. Böcker, C. Bektas, and C. Wietfeld, "Enhancing Reliability by Combining Manufacturing Processes and Private 5G Networks," 2022 IEEE 20th International Conference on Industrial Informatics (INDIN), p. 8, Sep. 2022.
- [4] 5G-ACIA, "Performance Testing of 5G Systems for Industrial Automation," 5G Alliance for Connected Industries and Automation, 5G-ACIA White Paper, Feb. 2021.
- [5] —, "A 5G Traffic Model for Industrial Use Cases," 5G Alliance for Connected Industries and Automation (5G-ACIA), 5G-ACIA White Paper, Nov. 2019.

- [6] C. Bektas, C. Schüler, R. Falkenberg, P. Gorczak, S. Böcker, and C. Wietfeld, "On the benefits of demand-based planning and configuration of private 5g networks," in 2021 IEEE Vehicular Networking Conference (VNC), 2021, pp. 158–161.
- [7] J. Prados-Garzon, P. Ameigeiras, J. Ordonez-Lucena, P. Muñoz, O. Adamuz-Hinojosa, and D. Camps-Mur, "5g non-public networks: Standardization, architectures and challenges," *IEEE Access*, vol. 9, pp. 153 893–153 908, 2021.
- [8] N. A. Mohammed, A. M. Mansoor, and R. B. Ahmad, "Mission-Critical Machine-Type Communication: An Overview and Perspectives Towards 5G," *IEEE Access*, vol. 7, pp. 127 198–127 216, 2019.
- [9] J. Ansari, C. Andersson, P. de Bruin, J. Farkas, L. Grosjean, J. Sachs, J. Torsner, B. Varga, D. Harutyunyan, N. König, and R. H. Schmitt, "Performance of 5G Trials for Industrial Automation," *Electronics*, vol. 11, no. 3, p. 412, Jan. 2022.
- [10] J. Rischke, P. Sossalla, S. Itting, F. H. P. Fitzek, and M. Reisslein, "5G Campus Networks: A First Measurement Study," *IEEE Access*, pp. 1–1, 2021.
- [11] D. Kaulbars, F. Schweikowski, and C. Wietfeld, "Spatially Distributed Traffic Generation for Stress Testing the Robustness of Mission-Critical Smart Grid Communication," in 2015 IEEE Globecom Workshops (GC Wkshps), Dec. 2015, pp. 1–6.
- [12] N. Patriciello, S. Lagen, B. Bojovic, and L. Giupponi, "An E2E simulator for 5G NR networks," *Simulation Modelling Practice and Theory*, vol. 96, p. 101933, Nov. 2019.
- [13] 3GPP, "NR; Physical layer procedures for data," 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 38.214, 03 2022. [Online]. Available: https://portal.3gpp.org/desktopmodules/ Specifications/SpecificationDetails.aspx?specificationId=3216