Towards Stateless Core Networks: Measuring State Access Patterns

Malte Höweler*, Zuo Xiang*, Franz Höpfner*, Giang T. Nguyen‡†, Frank H. P. Fitzek*†
*Deutsche Telekom Chair of Communication Networks, Technische Universität Dresden, Germany
‡Haptic Communication Systems, Technische Universität Dresden, Germany
†Centre for Tactile Internet with Human-in-the-Loop (CeTI), Technische Universität Dresden, Germany
Email: {forename}.{surname}@tu-dresden.de

***Abstract*—Future mobile communication networks, such as 5G and beyond, can benefit from Virtualized Network Functions (VNFs) when deployed on cloud infrastructures to achieve elasticity and scalability. However, new challenges arise as to managing states of Network Functions (NFs). Especially control plane VNFs, which are mainly found in cellular core networks like the 5G Core (5GC), received little attention since the shift towards virtualizing NFs. Most existing solutions for these core networks are often complex, intrusive, and are seldom compliant with the standard. With the emergence of 5G campus networks, UEs will be mainly machine-type devices. These devices communicate more deterministically, bringing new opportunities for elaborated state management.**

This work presents an emulation environment to perform rigorous measurements on state access patterns. The emulation comes with a fully parameterized Markov model for the UE to examine a wide variety of different devices. These measurements can then be used as a solid base for designing an efficient, simple, and standard conform state management solution that brings us further towards stateless core networks.

***Index Terms*—5GC, VNF, SBA, Measurements, MTC**

I. INTRODUCTION

With the emerge of 5G campus networks in vertical industries, the era of Industry 4.0 is approaching faster than ever before. Human-machine-interaction and especially machine-to-machine (M2M) communication are characterizing our daily life. Backbone of all this is the 5G System (5GS) that is fully virtualized according to Software Defined Networking (SDN) and Network Function Virtualization (NFV) paradigms and comes with a state-of-the-art Service Based Architecture (SBA).

This brings all prerequisites for excellent fail-over and scaling capabilities. However, these capabilities do not come inherently by using SDN and NFV, instead the NF must be carefully designed to meet them. Designing the state management is a difficult problem in NFV; a lot of attention has been payed towards user plane NFs over the past years. Nevertheless, since LTE, the control plane has also increasingly faced virtualization. These so called cellular NFs have different demands towards their state management. They are deployed in the cloud, and their state is composed out of complex objects including, for instance, timers and state machines. Accessing it requires high performance and availability. Some papers have already looked at LTE's EPC, and recently the first studies on the 5GC have been published. However, these works try to solve the problem in a very generic way and thus have to cover many corner cases, which in turn results in an immense complexity and design overhead. Constraining the application of the state management framework to campus networks, the target device group is no longer composed only of mobile phones, instead a large proportion is now accounted for by machine type devices. This is also the device group for which former mentioned capabilities are the most important. Machine type communication (MTC) inherits different characteristics, offering new opportunities for elaborated state management.

At the current state, this work provides a handy and flexible emulation setup that can be run on a single machine and emulate UEs with various communication characteristics. Rigorous measurements and analysis of state access patterns of 5GC NFs are ongoing work and will be published subsequently. These measurements and findings will later on provide a solid base for building an efficient and easy-to-use state management framework for cellular NFs.

II. BACKGROUND

This section first gives an overview of the 5GS, especially its control plane, and then shortly elaborates on 5G campus networks. In the second part, we work out the state management problem and discuss related works.

A. 5G System Architecture and Campus Networks

Figure 1 depicts the 5GS. The 5GS, or especially the 5GC, is designed in a cloud-native way. Functions are implemented as microservices that communicate with each other to facilitate independent and distributed deployment. User and control plane are separated in an SDN-fashion. The user plane is simply responsible for forwarding packets according to rules defined by the control plane. The control plane is, next to user plane configuration, responsible for connection, session, and mobility management, as well as for authentication and charging. From an architecture point of view, the biggest innovation of the 5GC is the SBA. NFs do not have reference point connections between each other anymore, instead they are mere service providers and consumers communicating via a stateless RESTful API over HTTP2 [1]. But still, some stateful connections remain. Specifically, the connections interfacing the world outside the 5GC control plane. For instance, the N1/N2 interface towards the Radio Access

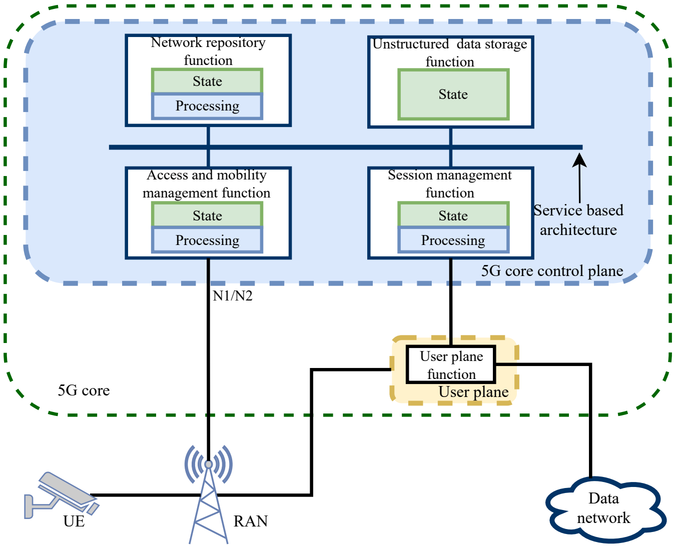

Fig. 1. Simplified 5G System architecture emphasizing the statefulness of the control plane.

Network (RAN) using NGAP over SCTP. Also, the 5GC NFs are not inherently stateless; Figure 1 illustrates how state and processing are locked into the same NF instance. The 3GPP already proposes a fully stateless NF design. To achieve such a stateless NF, it shall be split into a stateless processing instance that executes its logic on a shared data store. However, the technical reports on this data store, the Unstructured Data Storage Function (UDSF), and its interfaces are still under review [2]. Furthermore, no implementation details are given neither on deployment nor state management.

With the demand for private and secure communication networks that cover new use cases like delay sensitive machine type communication for e.g. robot control, 5G campus networks have emerged. A 5G campus network is a private 5G network only accessible for users affiliated with the campus organization that is bounded to a certain area. Peers of such a campus networks are mainly machine type devices like sensors, which are stationary and periodically report to a sink. Section III-D will provide further details on MTC and how it can be exploited for state management.

B. State Management

Stateless NFs are preferable to stateful implementations in many ways. Even though a VNF can be restarted quickly after failure, all data obtained at runtime (its state) will be lost. Similarly for the scenario of scaling the number of active instances due to increased and decreased traffic. Rerouting traffic to a new instance will not work as the new instance is not aware of the state, like forwarding rules. Generally, a lot of work in this domain has been done for user plane NFs, which maintain a rather simple state that can be mapped to key-value-stores. Much less focus has been paid to cellular NFs, which are present in today's mobile core networks like the EPC or the 5GC. These NFs maintain a complex set of states including timers, connections, and state machines. Additionally, failing is more critical as it could lead to longer down times [3]. StatelessNF [4] can be regarded as the flagship framework for user plane NFs, it considers multiple NF instances reading a shared state from an extra data store layer. However, their solution does not cover strong transactional semantics, like high availability and serialization. SCALE [5] and ECHO [6] are both designed for the EPC. ECHO however is only concerned with reliability, not scalability. Furthermore, it modifies the eNB, which is rather intrusive. SCALE on the other hand focuses solely on scalability. To the best of our knowledge are rVNF [3] and the work of Kulkarni et al. [7] the only works that are concerned with the 5GC control plane. rVNF does not pay special attention to state access semantics, thus everything is covered with a rather complex procedure. Additionally, it relies on a custom load balancer that overrules the load balancing of the network repository function. rVNF has not been tested with a 5GC implementation yet. In [7], state is shared between 5GC NFs, but similar to ECHO only reliability is considered. They additionally compare the delay costs of different backup granularities and transports used for transactions.

III. METHODOLOGY

In this section, we provide a detailed overview of our measurement approach, used metrics during measurements, as well as our measurement setup. This is then followed by a discussion of our used MTC traffic model and its configuration for the actual measurements.

A. Measurement Approach

To measure the state access patterns of the 5GC NFs we deploy a 5GC together with a gNB and an UE in an emulation. The UE connects/disconnects and sends/receives traffic according to a Markov model (compare Section III-D) that mimics, defined by its parameters, different types of UEs. This setup enables us to perform measurements on a variety of different UE types and to exactly observe the impact of each UE action on state accesses. This will later allow us to identify crucial and possibly omittable parts of the state that the framework needs to handle.

B. Metrics

Generally, NF state can be divided into static and dynamic state. The static state gets initialized once at boot time and stays unchanged thereafter, thus it is not critical. The dynamic state is created and continuously modified at runtime. It can be further divided into instance internal and network state. Through our measurements we want to identify how the dynamic state is handled exactly. This includes currently: i) what data is write once read many, ii) how ephemeral is the data, iii) what storage strategies would be most efficient, and iv) can the data be reduced by the determinism of MTC or would stateless devices from the viewpoint of the core/NF even be feasible. These questions can be answered by looking at memory allocations, writes/re-allocations, reads

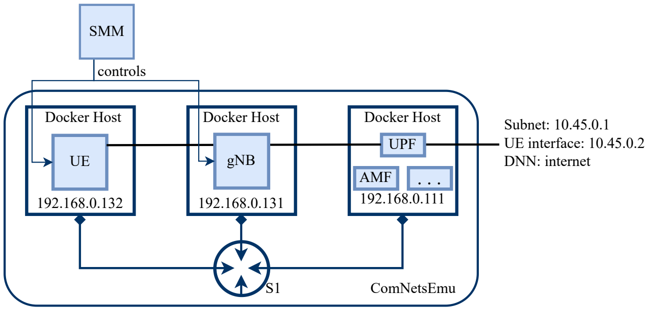

Fig. 2. Emulation setup in ComNetsEmu

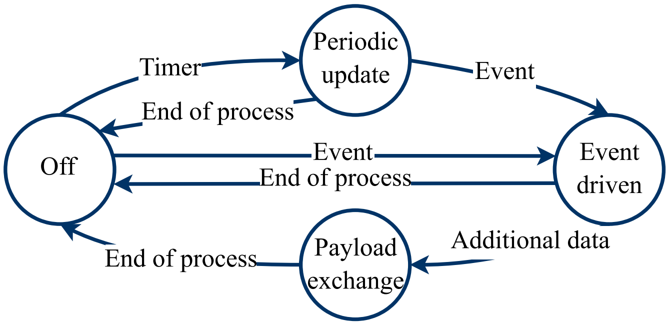

Fig. 3. Semi-Markov model for MTC traffic emulation

and frees, which we will obtain during runtime through code instrumentation. The measurements are justified by the fact that the 3GPP does not provide any implementation details. Also, it is not feasible to read reliable information out of either the specifications or existing implementations since the sheer amount of text is not graspable by human beings.

C. Emulation Setup

Measurements will be obtained in an emulation. Since we are only interested in the NF's reactions to certain UE events, for instances changes in its connection state, a physical testbed would neither provide any benefits nor different results. Furthermore, emulation ensures reproducibility. All measurements can be performed on a single laptop with the ComNetsEmu network emulator [8]. ComNetsEmu supports versatile computing applications in arbitrary network topologies. It extends Mininet [9] by deploying Docker containers [10] as hosts to enhance isolation and emulation capabilities. Figure 2 illustrates the emulation setup. UE, gNB and 5GC are all deployed on their own Docker host and are connected via the switch S1. The solid black line resembles the data path, which is established via TUN interfaces. Our Markov model (compare Section III-D) runs outside the ComNetsEmu in a Python script and controls the UE and gNB hosts through the Docker API. UE and gNB are based on UERANSIM v3.2.6 [11], as 5GC we are using Open5gs v2.3.2 [12] (we made the setup scrips available under [13]). Open5gs proves to be a viable choice for the 5GC, it is based on Release 16 and comes, compared to other open source 5GCs, with the broadest feature set [14]. Traffic is emulated using iPerf [15]. The iPerf server runs on the TUN interface of the UPF, the client connects from the TUN interface of the UE. This also accounts for the up link to down link ratio of MTC. For the actual measurements of the state access patterns, the code of open5gs has been instrumented to record the metrics discussed in Section III-B. The modified code and header-only helpers to quickly instrument further NFs are as well available under [16].

D. Traffic Model

As we aim to measure the specific reactions of NFs to UE behaviour, we will use a source model to control the MTC traffic for different device classes. Source models offer a broad parameter set to mime the communication behaviour of a single device and thus provide a good view of the fine-grained reactions of the network [17]. Shafiq et al. [18] found that MTC traffic strongly differs among device classes, this in turn results in different models proposed by the literature. Devices found in campus networks, such as metering devices and modems, inhibit strong deterministic traffic patterns and are stationary, which can be well accounted by deterministic source models.

Nikaein et al. [19] narrowed MTC traffic down to three elementary patterns; real world traffic is then a combination of these. (i) Periodic Updates (PU), where a status report is transmitted at fixed time intervals, (ii) Event Detection (ED) that maps to traffic triggered by a certain event, and (iii) Payload Exchange (PE) that covers the case when more data between device and server needs to be exchanged, for instance to provide additional information after an ED message. Figure 3 shows how these three patterns are integrated into a Markov chain that is extended by an additional OFF state. OFF resembles a traffic type where no traffic is sent (idle or disconnected). Each state of this Markov chain can then be added with further information to define, for instance, QoS rules, packet departure times, and sojourn times; each with either constant values or stochastic distributions. State transitions are defined as a Semi-Markov Model (SMM) [20], thus the transition probability p to the current state is given with $p_{s,s} = 0$. SMMs are a good fit for modeling MTC as they allow to capture a variety of traffic characteristics and are extendable (number of states/parameters) to allow for a good fitting [17].

E. Measurements

The model in Figure 3 reflects a device with traditional periodic updates superimposed by event driven communication, like a classical IoT sensor. Transition from OFF to PU is timer based. If an event occurs before the timer fires, the device moves to ED, from there, according to the defined transition probabilities, it either moves back to OFF (with $p_{ED,OFF} = 0.3$) or to PE to send additional data (with $p_{ED,PE} = 0.7$). The transition probability into ED is rather low [17] with $p_{OFF,ED} = 0.1$. Traffic in PU is only a few bytes with a low data rate, packet sizes in ED and PE are distributed following the IMIX genome [21] and are sent with higher data rates. Running over the course of the emulation, this model stimulates the connection state of the UE and triggers its iPerf client to send traffic. Generally, we expect that traffic amount and data rate will have little, if any, influence on the state. Whereas changing connection state is expected to have major impact.

IV. SUMMARY AND FUTURE WORK

In this paper, we propose to tackle the state management problem for cellular NFs with a bottom-up approach. We present an emulation environment to perform rigorous measurements on state access patterns, which will be the first of their kind.

The fully parameterized Markov model for the UE empowers us to emulate various devices and can be extended to control a group of multiple UEs. However, if multiple UEs are emulated, a further model has to be applied on top to account for possible correlations between the single UEs.

Performing and analyzing these measurements is currently ongoing work. The results will, later on, provide a solid base for building a simple and efficient state management framework for cellular NFs. We plan to leverage the measurements to design and implement a framework for 5G campus networks as future work.

ACKNOWLEDGMENT

This work was supported by the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) as part of Germany's Excellence Strategy—EXC 2050/1—Cluster of Excellence "Centre for Tactile Internet with Human-in-the-Loop" (CeTI) of Technische Universität Dresden under Project ID 390696704 and the Federal Ministry of Education and Research of Germany in the program of "Souverän. Digital. Vernetzt.". Joint projects 6G-life (Project ID: 16KISK001K) and 5G Campus++ (Project ID: F-011522-521-000-1121002).

REFERENCES

[1] 3GPP, "System architecture for the 5g system (5gs)," 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 23.501, 03 2022, version 17.4.0.

[2] ——, "Unstructured data storage services," 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 29.598, 03 2022, version 17.5.0.

[5] A. Banerjee, R. Mahindra, K. Sundaresan, S. Kasera, K. Van der Merwe, and S. Rangarajan, "Scaling the lte control-plane for future mobile access," in *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '15. New York, NY, USA: Association for Computing Machinery, 2015. [Online]. Available: https://doi.org/10.1145/2716281.2836104

[3] A. Katsarakis, Z. Tan, M. Balkwill, B. Radunovic, A. Bainbridge, A. Dragojevic, B. Grot, and Y. Zhang, "rvnf: Reliable, scalable and performant cellular vnfs in the cloud," Microsoft, Tech. Rep. MSR-TR-2021-7, April 2021. [Online]. Available: https://www.microsoft.com/en-us/research/publication/rvnf-reliable-scalable-and-performant-cellular-vnfs-in-the-cloud/

[4] M. Kablan, A. Alsudais, E. Keller, and F. Le, "Stateless network functions: Breaking the tight coupling of state and processing," in *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*. Boston, MA: USENIX Association, Mar. 2017, pp. 97–112. [Online]. Available: https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/kablan

[6] B. Nguyen, T. Zhang, B. Radunovic, R. Stutsman, T. Karagiannis, J. Kocur, and J. V. d. Merwe, "Echo: A reliable distributed cellular core network for hyper-scale public clouds," in *Mobicom 2018*. ACM, October 2018. [Online]. Available: https://www.microsoft.com/en-us/research/publication/echo-a-reliable-distributed-cellular-core-network-for-hyper-scale-public-clouds/

[7] U. Kulkarni, A. Sheoran, and S. Fahmy, "The cost of stateless network functions in 5g," in *Proceedings of the Symposium on Architectures for Networking and Communications Systems*, ser. ANCS '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 73–79. [Online]. Available: https://doi.org/10.1145/3493425.3502749

[8] Z. Xiang, S. Pandi, J. Cabrera, F. Granelli, P. Seeling, and F. H. P. Fitzek, "An open source testbed for virtualized communication networks," *IEEE Communications Magazine*, vol. 59, no. 2, pp. 77–83, 2021.

[9] M. P. Contributors. [Online]. Available: https://mininet.org

[10] D. Inc. [Online]. Available: https://www.docker.com/

[11] A. Güngör, "Ueransim source code," https://github.com/aligungr/UERANSIM, 2021, online, accessed 01.11.2021.

[12] S. Lee, "Open5gs source code," https://github.com/open5gs/open5gs, 2021, online, accessed 01.11.2021.

[13] M. Höweler and F. Höpfner, "5g emulation setup," https://github.com/Molle94/comnetsemu_open5gs, 2022, online, accessed 28.02.2022.

[14] L. B. Dominato, H. C. de Resende, C. B. Both, J. M. Marquez-Barja, B. O. Silvestre, and K. V. Cardoso, "Tutorial on communication between access networks and the 5g core," 2021. [Online]. Available: https://arxiv.org/abs/2112.04257

[15] J. Dugan, S. Elliott, B. A. Mah, J. Poskanzer, and K. Prabhu. [Online]. Available: https://iperf.fr/

[16] F. Höpfner, "open5gs state measurements," https://github.com/Molle94/open5gs, 2022, online, accessed 28.02.2022.

[17] M. Sansoni, G. Ravagnani, D. Zucchetto, C. Pielli, A. Zanella, and K. Mahmood, "Comparison of m2m traffic models against real world data sets," in *2018 IEEE 23rd International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)*, 2018, pp. 1–6.

[18] M. Z. Shafiq, L. Ji, A. X. Liu, J. Pang, and J. Wang, "A first look at cellular machine-to-machine traffic: Large scale measurement and characterization," *SIGMETRICS Perform. Eval. Rev.*, vol. 40, no. 1, p. 65–76, Jun. 2012. [Online]. Available: https://doi.org/10.1145/2318857.2254767

[19] N. Nikaein, M. Laner, K. Zhou, P. Svoboda, D. Drajic, M. Popovic, and S. Krco, "Simple traffic modeling framework for machine type communication," in *ISWCS 2013; The Tenth International Symposium on Wireless Communication Systems*, 2013, pp. 1–5.

[20] R. Nelson, *Probability, Stochastic Processes, and Queuing Theory*. Springer-Verlag, 1995, New York, NY, USA.

[21] A. Morton, "Imix genome: Specification of variable packet sizes for additional testing," Internet Requests for Comments, RFC Editor, RFC 6985, 7 2013.