Lisa Wernet, Laura-Marie Spang, Fridolin Siegmund, Tobias Meuser. *Resilient User Plane Traffic Redirection in Cellular Networks*.
To appear in the Proceedings of IEEE Conference on Network Function Virtualization and Software Defined Networks 2024, IEEE, 2024.

Resilient User Plane Traffic Redirection in Cellular Networks

Lisa Wernet*, Laura-Marie Spang*, Fridolin Siegmund, Tobias Meuser
Multimedia Communications Lab, Technical University of Darmstadt
{firstname.lastname}@tu-darmstadt.de

***Abstract*—The rising usage of cellular networks for various devices and applications increases the requirements for the availability and reliability of those networks. Yet, a failure of components on the data path inevitably leads to connectivity loss for users. The standardized restoration procedures after failures, as well as mechanisms for session continuity or redundancy on the data path either have a significant overhead in the control plane or duplicate the entire user traffic.**

In this paper, we propose a resource-efficient session management mechanism that proactively mirrors the session states of the central core function on the data path, the User Plane Function (UPF), to another instance. Our design is agnostic to UPF and gNB realizations and transparent to the user devices. Following a UPF failure detection, we redirect data plane traffic to the redundant UPF. Additionally, proactive user session mirroring reduces recovery times for the user plane of 5G and beyond. We demonstrate the feasibility of the proposed design with a prototypical implementation, and our experimental results emphasize its scalability and performance.

***Index Terms*—6G, Resilience, Redundancy, VNF, SDN**

I. Introduction

Current advancements in mobile telecommunications have introduced ground-breaking changes to the architecture of mobile core networks. 5G and beyond networks are picking up recent trends in Software-Defined Networking (SDN) [1] and cloud computing, among others. Progressive design decisions create an adaptive, evolvable, scalable, and programmable networked system that allows operators to diversify their deployed hardware components independently of the softwarized network functionality. Despite these efforts, ensuring zero downtime for all components of a cellular network remains a significant challenge. Failures like power outages or cyber attacks may disrupt cellular networks just like any other system composed of hardware and software. A failure of critical infrastructure like mobile networks leads to considerable disruptions of public safety.

Hence, the resilience of the entire cellular network is a recurring topic in research and has been examined in numerous works. We define resilience as “the ability of the network to provide and maintain an acceptable level of service in the face of various faults and challenges to normal operation” [2]. Following the separation of control and user plane in the 5G core network, we concentrate on the resilience of the data path, i.e., the user plane, in this work. The user plane provides session-based connectivity for users by transmitting user data via the next generation NodeB (gNB) and the User Plane Function (UPF) in the core network to the data network, e.g., the Internet. Section II further introduces the network components. In this work, we present a design and prototype implementation for 5G user session migration to realize session continuity in the presence of localized UPF failures. We show that the resilience of the user plane can be increased significantly through hot standby UPF redundancy and traffic redirection. We extend a network function of the control plane, the Session Management Function (SMF), with a session management approach that distributes identical session information to a designated primary UPF and its failover instance. Upon session establishment, modification, or release, all messages are mirrored to the redundant UPF. We use the 3GPP standardized heartbeat procedure [3] to detect UPF node failures and trigger a path switch to the replacement UPF. An inserted OpenFlow (OF) switch at the N3 interface between gNB and UPF instances enables a data path switch between UPF endpoints on OSI layer 2, which is agnostic to the GPRS Tunneling Protocol - User Plane (GTP-U) tunnels and their endpoints. The SMF triggers the operation automatically upon detecting a UPF node failure. Session migration is entirely transparent to the affected User Equipments (UEs), gNB, and N3 tunnels. It is also agnostic to the used UPF implementations, i.e., hardware, software, on-premises, or in the cloud. With a prototype implementation using free5GC and Open vSwitch (OVS), we demonstrate the feasibility and performance of session migration.

*Authors contributed equally to the paper.

II. Background

5G cellular networks are composed of three main elements: User Equipment (UE), next generation NodeB (gNB), and core network (see Figure 1). In this work, we focus on the 5G core network, which is divided into control plane and user plane. The control plane consists of around 30 network functions [4], of which only the Session Management Function (SMF) and Access and Mobility Management Function (AMF) are relevant for this work as they play a central role in session management. In the user plane, the User Plane Function (UPF) forwards user traffic in both ways between the gNB and the data network, e.g., the Internet. User traffic is sent over unidirectional tunnels on the N3 interface between the UPF and the gNB. Tunnel Endpoint Identifiers (TEIDs) inside GPRS Tunneling Protocol - User Plane (GTP-U) headers map transmitted user data to user sessions. The SMF installs and modifies user sessions and forwarding rules over the N4 interface to the UPF. The allocation and mapping of TEIDs and UE IP addresses are done either by the UPF or the SMF. Towards the Internet on the N6 interface, the UPF acts as PDU Session Anchor (PSA). For our simple network depicted in Figure 1, this means that any communication partners in the Internet know the address of the single PSA per session. The N9 interface connects two UPFs, i.e., an intermediate UPF (i-UPF) and a PSA UPF.

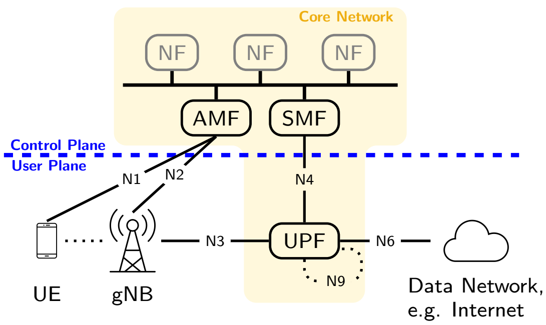

Fig. 1: Overview of the 5G network including the core network with network functions (NFs) and the data path between UE, gNB, UPF and the Internet. Inspired by [4].

A. Failure Detection and Restoration Procedures

The 5G specification includes several conceptual measures for the detection and mitigation of various failures in the 5G system, including UPF failures [3]. For the detection of node failures and restarts, the SMF and UPF exchange heartbeat messages over the N4 interface. Also, GTP-U endpoints may use GTP-U Echo Request and Echo Response messages for failure detection. A UPF failure always means an interruption of the user data path. After a restart of the faulty UPF, there are two possibilities: If the UPF is responsible for allocating TEIDs and UE IP addresses, it may not reuse previous TEIDs or addresses. UEs must re-establish sessions to obtain new session parameters. If the SMF is responsible for the allocation, the UPF accepts the allocated parameters as long as the relevant network resources are available. In this case, UE IP addresses and TEIDs may also not be preserved. In both cases, the data path is interrupted until the UPF restarts and the subsequent session restoration mechanism is completed.

B. User Plane Redundancy

Redundant paths or nodes in network topologies are a typical means to increase the overall resilience of a network [5] [2]. The 3GPP standard includes topologies for using multiple PSA UPFs per user session [4]. An i-UPF serves as a branching point and terminates the N3 interface towards the gNB. This i-UPF is then connected to several PSA UPFs over N9 interfaces. In this topology, the UE has a single IP address. Multi-homing for user sessions follows a similar approach but is only eligible for IPv6 addressing where one user session is associated with multiple IPv6 prefixes, one for each PSA UPF. Although these topologies could overcome the downtime during session restoration after a failure of a PSA UPF, a failure of the i-UPF would still interrupt the connection to the gNB.

Additionally, the 5G standard contains a section on redundant transmission in the context of Ultra Reliable and Low Latency Communications (URLLC) [4]. Yoshizawa et al. briefly summarize three potential methods for redundant user plane paths: dual connectivity, redundant transmission on N3/N9 interfaces and redundant transmission at the transport layer [6]. All three methods replicate user data and send them over all available data paths. For dual connectivity, a UE is connected with two redundant user plane paths comprising two gNBs and two UPFs that are connected to the same data network. Also, each UPF is connected to its own SMF, and both SMFs are connected to one AMF. The other two methods are restricted to the redundant user plane paths between the gNB and the PSA UPF. Redundant transmission on N3/N9 interfaces means that each GTP-U tunnel is duplicated. The third method refers to the redundant transmission at transport layer without tunnel duplication. All those measures for traffic duplication are resource-intensive and inflexible and may induce new challenges, such as dropping the redundant traffic flows.

Standby modes refer to the operational state of redundant Virtualized Network Functions (VNFs) that are meant to replace primary instances in the case of failures [7]. A VNF in cold standby is inactive and it takes some time to reach an operational state. In hot standby, states are constantly synchronized and the secondary instance can take over at any time. Warm standby means that the redundant VNF is active but running only few tasks.

C. Optional Modes for Session and Service Continuity (SSC)

The 3GPP also specifies procedures to change the PSA UPF of an active session in the context of SSC [4]. If the UE supports SSC mode 2 or 3 (optional), the PSA can be relocated to another UPF, but only at the cost of breaking the active session and establishing a new one. In mode 2, the 5G core network informs all UEs that their sessions have been released from the old PSA UPF and that they can directly establish new sessions as the new PSA UPF is already available. In mode 3, UEs are informed before the connectivity to the old PSA is released, and additionally, the 5G network guarantees the availability of the old PSA UPF for some time while the new PSA UPF is already available. However, these procedures are not applicable in the case of unexpected failures.

III. Related Work

Various publications studied the relocation of GPRS Tunneling Protocol - User Plane (GTP-U) tunnel endpoints in data planes of cellular networks. To reduce data plane latency, Heinonen et al. [8] suggest to dynamically switch the GTP-U endpoint of the 4G Evolved Packet Core (EPC) to a more potent hardware instance [8]. In the context of latency reduction for Mobile Edge Computing (MEC), Fondo-Ferreiro et al. propose an SDN-based solution to dynamically change the endpoint of the GTP-U tunnel for a particular User

Lisa Wernet, Laura-Marie Spang, Fridolin Siegmund, Tobias Meuser. *Resilient User Plane Traffic Redirection in Cellular Networks*.
To appear in the Proceedings of IEEE Conference on Network Function Virtualization and Software Defined Networks 2024, IEEE, 2024.

Equipment (UE) to a closer 4G termination point [9]. In both works, the authors do not evaluate session continuity and are limited to on-demand adaptations for individual user sessions. Said et al. demonstrate the feasibility of recovering from data plane failures in 4G networks by realizing the data plane of the 4G EPC and the evolved NodeB (eNB) using OpenFlow (OF) switches [10]. An Software-Defined Networking (SDN) controller detects switch failures and installs the necessary session information on an instance in cold standby. They show that with their approach TCP throughput drops for 2 s after a failure until it fully recovers. Despite significantly reducing the mean time to recover and packet loss compared to session re-establishment, this approach is limited to OF-capable data plane nodes, including the eNB. Peters and Khan propose a solution for Quality of Service- and latency-aware next generation NodeB (gNB) handover situations [11]. Using predictions of user mobility, they anticipate the optimal placement of an intermediate UPF (i-UPF) in the data path between the source and target gNB. This i-UPF is preconfigured with new tunnel information for individual user sessions to accelerate user plane reconfiguration during the handover event.

Additional related work investigates the state transfer between User Plane Function (UPF) instances, omitting the need for traffic redirection. One general approach is the outsourcing of all session information to an external component like a database [12] or a SDN controller [9]. Other solutions consider the migration of containers or virtual machines [13] [14], which are limited to software UPFs and induce migration times in the range of seconds. Lastly, [15] introduce an Edge Synchronization Protocol to synchronize subscriber state between edge data centers in 4G. In a hot standby failover scenario, the researchers demonstrate that an affected UE notices no significant increase in delay of the performed 4G procedures. However, the usage of such a protocol directly between UPF instances contradicts the separation of control and user plane and may risk the performance of user traffic forwarding.

Disaggregated Radio Access Network (RAN) solutions like those standardized by the O-RAN consortium allow the implementation of some components on commodity servers instead of using monolithic black boxes based on specialized hardware. Regarding fault tolerance of such virtualized RANs, Lazarev et al. [16] observe variations in existing downlink traffic on the physical layer to recognize failures in layer-1 processes. In case of a failure, they use two middleboxes to redirect the traffic to a running, redundant stateless process. In [17], an existing controller for the RAN is extended to monitor heartbeat messages from components handling user traffic. If a heartbeat message is missed, the controller triggers a session migration to a redundant component. Although this failover component is initialized and running to speed up failure mitigation, relevant states related to user sessions are not synchronized, and thus user sessions are disrupted.

IV. Resilient User Plane: Design and Implementation in the 5G core Network

Our solution for a resilient data path in cellular core networks using redundant User Plane Functions (UPFs) consists of the following steps:

1) Proactive state replication between the primary UPF and the secondary UPF instance
2) Failure detection in UPF instances
3) Redirection of user traffic to the secondary UPF in the failure case

In the following, we discuss these steps in detail. We concentrate on failures in PDU Session Anchor (PSA) UPFs which are, from now on, referred to as UPFs.

A. Proactive UPF State Replication

We assume that a UPF failure can occur at any time in an isolated manner. To quickly regain availability of communication, we propose a hot standby secondary UPF instance, immediately synchronizing changes in session states to assure that this instance is operable at any time.

The UPF session state consists of the tunnel information (identifier and IP address) for the downlink and uplink tunnels and the corresponding session rules. Additionally, the UPF keeps PFCP session states from the N4 interface with the Session Management Function (SMF) for each session to track local and remote Session Endpoint Identifiers (SEIDs). This information is created, modified, and deleted when sessions are established, modified, or released. The 3GPP specification allows that Tunnel Endpoint Identifiers (TEIDs) and IP addresses for User Equipments (UEs) are assigned by either the UPF or the SMF (see Section II). Here, we require that only the SMF creates these values or that the SMF can modify these values in UPF instances.

We adapt the SMF functionality to support the synchronization between two UPF instances. To validate that both UPF instances are in sync, the SMF always waits for both PFCP responses until it resumes by sending a response to the Access and Mobility Management Function (AMF). Although session modification procedures during normal operation may be slightly slower, the effects on the network performance are negligible. Note that although the implementation of the SMF needs to be adapted, the N4 interface between the SMF and one UPF instance is not modified. Also, the operation of UPFs as primary or secondary instances requires no changes.

B. Failure Detection in UPF Instances

For failure detection, we rely on the existing mechanism of heartbeat messages sent between the SMF and associated UPF instances over the N4 interface, mentioned in Section II. The average Round-Trip Times (RTTs) for heartbeat request and response messages are highly system-dependent. Thus, the interval between subsequent heartbeat requests needs to be configurable to fit the requirements of the system. Additionally, system dependent variance in heartbeat RTTs may lead to unnecessary failovers. To limit the influence of delayed heartbeat messages, we require that it is possible to ignore

Lisa Wernet, Laura-Marie Spang, Fridolin Siegmund, Tobias Meuser. *Resilient User Plane Traffic Redirection in Cellular Networks*.
To appear in the Proceedings of IEEE Conference on Network Function Virtualization and Software Defined Networks 2024, IEEE, 2024.

a configurable number of subsequent heartbeat misses. This reduces the number of unnecessary failovers. If a failure in one UPF instance is detected, the SMF first checks whether the faulty UPF instance is the primary one. Failures of the passive instance are noted but ignored.
After confirming that the primary UPF is faulty, the SMF verifies the association state of the secondary UPF. If the verification is successful, the user traffic is redirected to the secondary UPF.

C. Traffic Steering between UPF Instances

Backed by the amount of SDN-related research on 5G networks and its adoption in production networks, we propose the topology depicted in Figure 2. We insert an SDN-capable network switch into the N3 link between the next generation NodeB (gNB) and two redundant UPFs and into the N6 link between the UPFs and the Internet. The controller to install and modify flow rules for this network switch is placed in the SMF. Upon failure of the primary UPF, the controller switches the path to the secondary UPF by modifying a single flow rule for the N3 and N6 links, respectively. This path switch is a single OSI layer-2 operation that is agnostic to the upper layers. Neither the GPRS Tunneling Protocol - User Plane (GTP-U) header nor the encapsulated user traffic need to be parsed. Both UPF instances, primary and secondary, must either be configured with the same MAC and N3 and N6 IP addresses, or the respective fields in the outer packet headers must be adapted through OpenFlow. As GTP-U is sent over UDP/IP, changing the MAC and IP falls within the default set of actions available in OpenFlow.

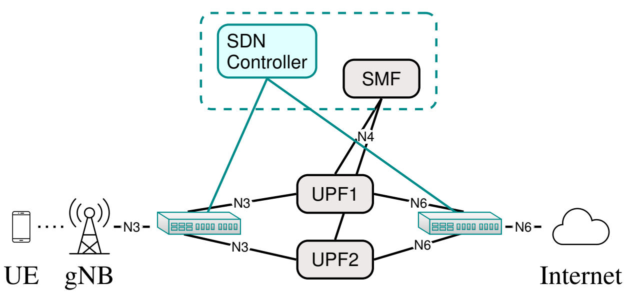

Fig. 2: The modified SMF manages two redundant UPF instances. The SDN controller triggers a path switch to the second UPF by modifying the current OpenFlow rules.

D. Implementation

To showcase the feasibility and performance of our design, we integrated the proposed mechanisms into the open-source 5G core implementation free5GC (v3.4.0) [18]. Our prototype realizes dual UPF session management and the corresponding session state handling, a UPF failure detection mechanism based on configurable N4 PFCP heartbeats and the failover mechanism using Software-Defined Networking (SDN). The SMF is extended with OpenFlow capabilities to directly communicate with a virtualized switch (Open vSwitch (OVS) [19]) connected to two free5GC PSA UPFs in the data plane. Supporting OpenFlow allows us to realize a concise setup without limiting the generalizability to specific deployments, including hardware accelerated as well as entirely software based UPF solutions.

E. Summary

Combining the previous steps, we can describe the resulting procedure as follows: After startup, both UPF instances are associated with the SMF and sessions are established. During normal operation of the primary UPF, the SMF synchronizes session states between both UPF instances and regularly exchanges heartbeat messages with both instances. Upon noticing a missing heartbeat response from the primary UPF, the SMF checks if the secondary UPF is available and ready. Then, the SMF triggers the path switch to redirect the traffic.
In summary, our resilience mechanism for isolated UPF failures in cellular networks respects the following requirements:

Transparency towards UEs:
Our mitigation strategy does not include communication between the Control Plane (CP) of the core network and individual UEs. In contrast to the specified restoration strategies described in Section II-A, we ensure that existing sessions are continued using the same TEIDs and UE IP addresses. In comparison with optional Session and Service Continuity (SSC) modes 2 and 3, our design is independent of the capabilities of individual UEs (see Section II-C).

Transparency towards gNBs:
During failover, existing GTP-U tunnel endpoints at the gNB do not change. Lost packets due to the temporary interruption of the user plane are handled by inner transport protocols of the encapsulated user traffic.

Independence of UPF implementation:
All interfaces towards other network nodes, as well as the functionality of the UPF, remain unchanged.

Minimal control plane overhead:
Although the dual UPF management increases the processing load on the SMF and also increases the network load on the N4 interface, the resulting control plane overhead involves no further signaling towards other 5G network components or individual UEs.

Resource efficiency in the user plane:
Compared with known approaches involving link redundancy (Section II-B), our hot standby approach consumes fewer resources during normal operation. We omit traffic duplication in the user plane as only one UPF instance is active at a time, i.e., forwarding user plane traffic.

V. Evaluation

To evaluate our session migration design, we include our implementation in an end-to-end 5G testbed. The hardware and installed operating systems are listed in Table I. The server hosts the 5G core consisting of the free5GC control plane (v3.4.0), including our modified version of the Session Management Function (SMF) and two free5GC User Plane Functions (UPFs). We place each UPF within its own Linux network namespace. The server runs an Open vSwitch (OVS)

Lisa Wernet, Laura-Marie Spang, Fridolin Siegmund, Tobias Meuser. *Resilient User Plane Traffic Redirection in Cellular Networks*.
To appear in the Proceedings of IEEE Conference on Network Function Virtualization and Software Defined Networks 2024, IEEE, 2024.

that later operates the path switch. The SMF can control this switch using the inbuilt `ovs-ofctl` tool. The VM runs UERANSIM [20], a widely used software simulator for next generation NodeBs (gNBs) and User Equipment (UEs). We simulate one gNB connected to different numbers of UEs. We generate traffic using iPerf3 or tcpreplay. The external N3 and N6 interfaces of the free5GC server and the UERANSIM VM are connected over a programmable switch running P4STA, an open-source framework for packet timestamping and measurements [21], as shown in Figure 3. P4STA stamps user plane packets at two points: at *s1*, which indicates incoming uplink traffic on the N3 link from the gNB, and at *s2*, which denotes outgoing uplink traffic on the N6 link, which was forwarded by the active UPF instance. The server's external interfaces are further connected to the OVS and forwarded into the UPF namespaces. A `SIGKILL` signal triggers UPF failures. We focus on demonstrating the feasibility of our approach and its transparency to the involved UEs, gNB, and UPFs, and its scalability in terms of the number of connected UEs.

Server	Intel Xeon D-1541 2.10 GHz, 32 GB RAM, 10 GbE NIC	Ubuntu 20.04 LTS, kernel 5.4.0
VM	Intel Xeon ES-1410 v2 2.8 GHz, 32 GB RAM, 10 GbE NIC	Ubuntu 20.04 LTS, kernel 5.4.0
P4 Switch	APS Networks BF6064X-T	External Host: Ubuntu 20.04 LTS, kernel 5.4.0

TABLE I: Testbed hardware and operating systems.

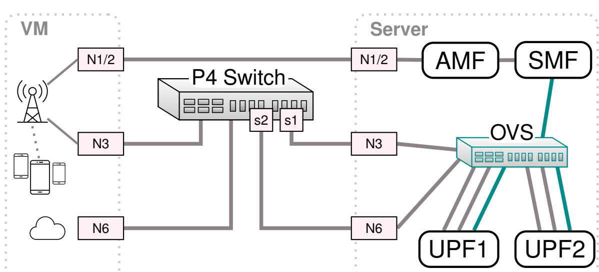

Fig. 3: Topology of the evaluation testbed.

A. Session Continuity

First, we show that our approach maintains session continuity during a failure of UPF1. We connect one UE and use iPerf3 to generate TCP traffic for $10\ s$. We measure the resulting throughput in different scenarios: normal operation without failure, UPF failure without session continuity (baseline scenario) and UPF failure with dual session management and traffic redirection. Figure 4 shows the throughput in the uplink direction. In the baseline scenario, the UE is directly disconnected after the failure. As a result, the iPerf3 client crashes. Additionally, it takes about $5\ s$ until we can successfully restart UPF1. The gray area indicates that the restarted UPF1 is available but still has no active user sessions. Then, the UE may request a new session and a new pair of GTP-U tunnels between the gNB and UPF1 may be established. Eventually, a new iPerf3 client process may be started. This observation matches similar results for recovery times after failures in the 5G data plane that are in the range of several seconds, e.g., $6.2\ s$ in [16]. However, our approach preserves the session parameters, and the iPerf3 client keeps working. The Packet Forwarding Control Protocol (PFCP) heartbeat interval is set to $2\ ms$ and the SMF will trigger a failover after missing 5 consecutive replies from UPF1. This results in a detection time of approximately $10\ ms$. The throughput drops to zero for $221\ ms$ but recovers directly afterward. Note that packets already received at the N3 interface in the namespace of UPF1 but have yet to be processed or transmitted by this failed UPF are lost. The same is true for unprocessed downlink packets from the iPerf3 server. After the TCP stack retransmits those lost packets, the throughput recovers and resembles the throughput without UPF failure. Our approach transparently redirects the traffic and does not interrupt existing user sessions.

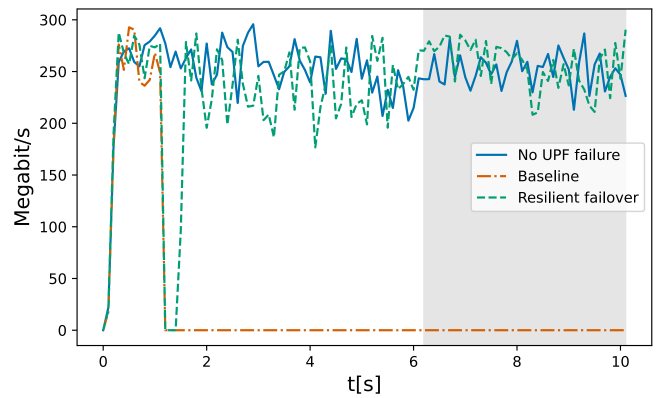

Fig. 4: TCP throughput for different scenarios.

B. Scalability

We want to evaluate the influence of the number of connected UEs on failure recovery. As previously reported by Lando et al., the free5GC control plane suffers from scalability problems with increasing numbers of parallel session establishments [22]. In the baseline scenario described in Section V-A, all UEs try to reconnect to the network in parallel after a failure. In the best case, without re-transmissions, the control plane traffic for session reestablishment scales linearly with the number of UEs.

Our approach omits this scalability issue as it does not release user sessions during failover. Although it induces some overhead for the state synchronization between both UPF instances during regular operation, the traffic redirection during UPF failures involves no control plane messages between UEs and the core network. We determine the downtime of the user plane using P4STA: the second timestamp of the first packet processed by UPF2 subtracted by the second timestamp of the last packet processed by UPF1. First, we establish sessions for the desired number of UEs. We use the control interface for the GTP-U kernel module to extract the corresponding Tunnel Endpoint Identifiers (TEIDs) and IP addresses. Then, we use this information to craft a pcap file containing 200 GTP-U packets where the number of packets belonging to different

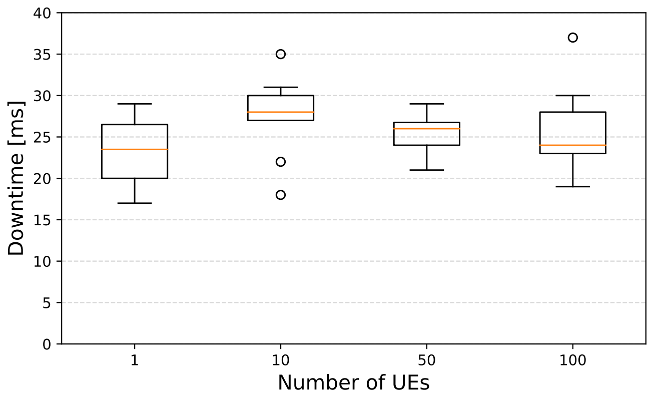

Fig. 5: Distribution of downtimes during failover for different numbers of UEs.

UEs is equal. The payload of the GTP-U packets consists of TCP packets with 1320 Bytes random payload. Then, we inject the packets directly into the N3 interface using tcpreplay in a loop for $5\ s$. After $2\ s$, we terminate UPF1. We compare the downtimes for different numbers of connected UEs, i.e., $1, 10, 50$, and 100 UE(s), as shown in Figure 5. All measured downtime values range from $17ms$ to $37ms$. The mean value for one UE is $23.5\ ms$, and for 100 UEs, it is $24\ ms$. We conclude that with our approach, the resulting downtime is independent of the number of connected UEs.

VI. Conclusion and Future Work

A resilient user plane is essential for reliable connections over 5G and beyond networks. In this work, we showed that proactive state mirroring between two User Plane Function (UPF) instances and traffic redirection significantly increase the availability of the user plane in failure events. Our solution preserves existing user sessions and transparently resolves UPF failures. Compared to other solutions regarding UPF redundancy, our approach is more resource efficient as we do not rely on user traffic duplication. Our results confirm that the number of connected devices does not influence the recovery time.

To further decrease resource consumption, one UPF instance could act as secondary UPF for several primary UPFs. Also, given that we currently only consider UPF node failures or N4 link failures, additional information is required to detect N3 link failures.

Acknowledgement

This work has been co-funded by the German Research Foundation (DFG) within the Collaborative Research Center (CRC) 1053 MAKI. The authors acknowledge the financial support by the Federal Ministry of Education and Research of Germany in the project “Open6GHub” (grant number: 16KISK014).

References

[1] Open Networking Foundation (ONF), “Software-defined networking: The new norm for networks,” 2012, available at https://opennetworking.org/wp-content/uploads/2011/09/wp-sdn-newnorm.pdf, last accessed 2024-07-15.

[2] J. P. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller, and P. Smith, “Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines,” *Computer Networks*, vol. 54, no. 8, pp. 1245–1265, 2010.

[3] 3rd Generation Partnership Project (3GPP), “5G System; Restoration Procedures,” Tech. Spec. TS 23.527 v17.6.0 Release 17, 2022.

[4] 3rd Generation Partnership Project (3GPP), “System architecture for the 5G System (5GS),” Tech. Spec. TS 23.501 v17.8.0 Release 17, 2023.

[5] A. Khalil, B. Becker, L. Wernet, R. Kundel, B. Richerzhagen, T. Meuser, and R. Steinmetz, “Dependability: Enablers in 5G Campus Networks for Industry 4.0,” in *2023 19th International Conference on the Design of Reliable Communication Networks (DRCN)*. IEEE, Apr 2023.

[6] T. Yoshizawa, S. B. M. Baskaran, and A. Kunz, “Overview of 5G URLLC system and security aspects in 3GPP,” in *2019 IEEE Conference on Standards for Communications and Networking (CSCN)*. IEEE, 2019.

[7] H. Chantre and N. L. S. d. Fonseca, “Reliable broadcasting in 5g nfv-based networks,” *IEEE Communications Magazine*, vol. 56, no. 3, pp. 218–224, 2018.

[8] J. Heinonen, T. Partti, M. Kallio, K. Lappalainen, H. Flinck, and J. Hillo, “Dynamic tunnel switching for SDN-based cellular core networks,” in *Proceedings of the 4th workshop on all things cellular: operations, applications, & challenges*. ACM, 2014.

[9] P. Fondo-Ferreiro, F. Gil-Castineira, F. J. Gonzalez-Castano, and D. Candal-Ventureira, “A software-defined networking solution for transparent session and service continuity in dynamic multi-access edge computing,” *IEEE Transactions on Network and Service Management*, vol. 18, no. 2, pp. 1401–1414, 2021.

[10] S. B. H. Said, B. Cousin, and S. Lahoud, “Software defined networking (SDN) for reliable user connectivity in 5G networks,” in *2017 IEEE Conference on Network Softwarization (NetSoft)*. IEEE, 2017.

[11] S. Peters and M. A. Khan, “Anticipatory user plane management for 5G,” *2018 IEEE 8th International Symposium on Cloud and Service Computing (SC2)*, pp. 9–15, 2018.

[12] 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*. ACM, 2021.

[13] M. Gundall, J. Stegmann, C. Huber, and H. D. Schotten, “Towards organic 6G networks: Virtualization and live migration of core network functions,” in *Mobile Communication - Technologies and Applications; 25th ITG-Symposium*, 2021, pp. 1–6.

[14] A. Leiter, A. Hegyi, N. Galambosi, E. Lami, and P. Fazekas, “Automatic failover of 5G container-based user plane function by ONAP closed-loop orchestration,” in *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*. IEEE, 2022.

[15] E. Cau, M. Corici, P. Bellavista, L. Foschini, G. Carella, A. Edmonds, and T. M. Bohnert, “Efficient exploitation of mobile edge computing for virtualized 5G in EPC architectures,” in *2016 4th IEEE International Conference on Mobile Cloud Computing, Services, and Engineering (MobileCloud)*. IEEE, 2016.

[16] N. Lazarev, T. Ji, A. Kalia, D. Kim, I. Marinos, F. Y. Yan, C. Delimitrou, Z. Zhang, and A. Akella, “Resilient Baseband Processing in Virtualized RANs with Slingshot,” in *Proceedings of the ACM SIGCOMM 2023 Conference*, 2023, p. 654–667.

[17] L. Wernet, F. Siegmund, R. Kundel, T. Meuser, and R. Steinmetz, “Fault-Tolerance in Future RAN: Enabling Redundancy for Efficient Recovery of RAN Network Functions,” in *2023 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*. IEEE, 2023, pp. 105–110.

[18] free5GC, “free5GC GitHub repository,” available at https://github.com/free5gc/free5gc, last accessed 2024-07-15.

[19] Open vSwitch, “Open vSwitch GitHub repository,” available at https://github.com/openvswitch/ovs, last accessed 2024-07-15.

[20] A. Güngör, “UERANSIM GitHub repository,” available at https://github.com/aligungr/UERANSIM, last accessed 2024-07-15.

[21] R. Kundel, F. Siegmund, J. Blendin, A. Rizk, and B. Koldehofe, “P4STA: High Performance Packet Timestamping with Programmable Packet Processors,” in *NOMS 2020 - 2020 IEEE/IFIP Network Operations and Management Symposium*, 2020, pp. 1–9.

[22] G. Lando, L. A. F. Schierholt, M. P. Milesi, and J. A. Wickboldt, “Evaluating the performance of open source software implementations of the 5G network core,” in *NOMS 2023-2023 IEEE/IFIP Network Operations and Management Symposium*. IEEE, 2023.