Fridolin Siegmund, Philip Jonas Franz, Bjoern Nagel, Maik Rüder, Lisa Wernet, Matthias Hollick, Ralf Kundel. *State-aware Subscriber Steering in Fiber Access Networks for Improved Resilience.* To appear in the Proceedings of IEEE Conference on Network Function Virtualization and Software Defined Networks 2024, IEEE, 2024.

State-aware Subscriber Steering in Fiber Access Networks for Improved Resilience

Fridolin Siegmund*, Philip Jonas Franz*, Bjoern Nagel‡, Maik Rüder‡,
Lisa Wernet*, Matthias Hollick*, Ralf Kundel*
*Technical University of Darmstadt
fridolin.siegmund@tu-darmstadt.de
‡Deutsche Telekom Technik GmbH
{nagelb, maik.rueder}@telekom.de

***Abstract*—Network softwarization is currently emerging in home and enterprise Internet access networks. Still, the advantages of programmable hardware nodes and virtualized network functions in the access network as resilience enablers should be researched more. Access networks are typically built in a two-layer tree topology with sparse redundancy and limited adaptability. To overcome this, this work investigates how softwarization with redundancy can achieve resilience improvements in such networks. We present an architecture and prototypical implementation of a fiber-based access network and subscriber termination, with a redundant and adaptive design throughout the whole path from subscribers to the Internet. By proactively migrating states and steering subscriber flows, we reached total average downtimes of around 156 milliseconds to restore Internet access after disruption in the case of one subscriber. In comparison, it takes hours to restore Internet access in case of an outage with zero redundancy. Our solution is cost-effective by exploiting the architectural advances of Software Defined Networking, allowing the use of virtualized network functions and hardware-based functions simultaneously.**

***Index Terms*—PON, Resilience, NFV, SDN, Subscriber Steering, Hardware Acceleration, Access Networks, Redundancy**

I. INTRODUCTION

Besides ever-increasing amounts of data flowing through today's Internet, society is also increasingly dependent on an assured service quality. Passive Optical Network (PON) based fiber access networks offer a multi-gigabit bandwidth while being cost-efficient in deploying and upgrading in the field, fulfilling the needs of customers and service providers. While services such as telephony, including emergency calls, are life-threatening if disrupted, other services, such as media streaming, require stable Internet access and, therefore, result in unsatisfied customers in case of disruption. Also, enterprises demand resilient and high bandwidth networks where each second of disrupted customer services can create significant financial damage. Additionally, in other types of networks, such as the 5G backhaul network, PONs may be deployed for a cost-efficient connection between the core and radio access network. The wide deployment of PONs for this kind of service enhances the need for resilience even more. In general, enhancing the resilience of the access network improves the availability of all services built on it.

Traditional access networks for home and business subscribers are today mostly built with proprietary black box hardware, limiting, besides other factors, configurability, flexibility, and upgradeability in multiple ways. Novel white box hardware approaches aim to improve those factors by implementing the Software Defined Networking (SDN) paradigm with open Application Programming Interfaces (APIs) and separating the data and control plane [1], [2]. Besides advances for operators, white box hardware, in combination with open-source solutions, opens the possibility of research in this field.

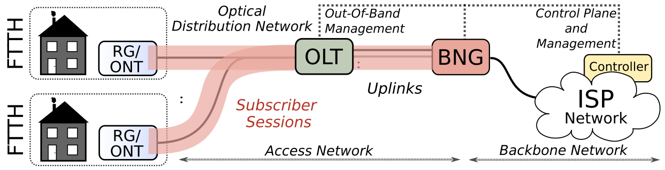

Fig. 1: PON access network providing services to subscribers.

In PONs, the Optical Line Terminal (OLT) provides fiber access to homes and businesses, as shown in Figure 1. The OLT forwards the subscriber session flows via one or multiple uplinks to one Broadband Network Gateway (BNG), terminating all sessions individually. The BNG is the gateway between the Internet Service Provider (ISP) services, such as Internet access, and the Residential Gateway (RG) at the subscriber side. Between the RG, and the BNG, the subscriber session tunnel is established. Existing proprietary OLT software stacks in multiple uplink deployments aim at static load balancing and link aggregation; however, resilience is mostly not considered. The dynamic establishment of subscriber paths would enhance resilience by responding to link failures and facilitates adaptive link load balancing.

While redundant paths in core networks are manageable in several ways, the access network has state-aware flows, with each flow mapping a subscriber session. Each subscriber is connected to the network through an RG, which utilizes a tunneling protocol, e.g., PPP over Ethernet (PPPoE), for the end-to-end subscriber session tunnel to the BNG. Steering these flows through different paths in the access network requires the nodes to be equipped with session states for flow path management and session termination.

Terminating subscriber sessions, consisting of encapsulation and decapsulation of network packets, is one fundamental and resource heavy funcionality of the BNG. Software-based implementations benefit from scalability and low energy consumption at low load, while hardware-based solutions provide higher performance. The awareness of subscriber states in the network allows for utilizing both hardware and software BNGs. Previous research has shown that BNG functionality, such as the PPPoE session termination, is achievable in hardware [3]. In this work, we investigate the combination of hardware-based and software-based network function instances as an extension to the aforementioned subscriber steering approach.

Our contributions to enable resilient and load-adaptive access networks are as follows:

- An investigation of redundancy capabilities in access networks.
- Integration of multi-uplink functionality in an open-source-based OLT software stack.
- Design of a subscriber flow steering strategy and an active state management for BNG redundancy.
- As a proof-of-concept, we demonstrate and evaluate a hybrid hardware and softwarized BNG deployment.

II. BACKGROUND

In the following, we provide an overview of access networks in general, as our proposed prototype affects multiple points.

Access networks connect the subscribers to the backbone network of ISPs and, by that, to the Internet. Access networks are typically built in a tree topology, connecting hundreds to thousands of subscribers through one access node. The access node forwards the subscriber sessions to the termination node. For PON networks, the access node is called OLT, and the termination node is called BNG. Figure 2 depicts a PON where multiple subscriber Optical Network Terminators (ONTs) are connected to one fiber tree, receiving packets as broadcasts downstream while utilizing Time-division Multiplexing (TDM) for sending packets upstream. Considering a Gigabit Passive Optical Network (GPON), the OLT as access node typically comes with, e.g., 32 or 64 PON ports, each supplying up to 256 subscribers. While different deployments are possible, the OLT forwards thousands of subscriber flows, typically through one or two uplinks to the BNG.

The open-source software stack Virtual OLT Hardware Abstraction (VOLTHA) is the only available open-source software stack for white box OLTs, with several vendors offering certified OLTs on the market. VOLTHA offers the abstraction of OLT access nodes in the access network [2], [4]. Figure 3 shows the microservice-based architecture of VOLTHA, typically deployed in a Kubernetes cluster. The abstracted API to the OLTs enables SDN controllers, such as ONOS*, managing the OLTs as logical switches with subscribers and uplinks as interfaces [5]. An SDN controller can, e.g., receive port-up events and install and remove flow rules in the logical switch to establish subscriber access to the BNG. The VOLTHA core translates the operations into device-specific messages for the OLT through modular adapters (OpenOLT Adapter). On the white box OLT, the OpenOLT Agent is executed on the x86 processor running Linux, exchanging the messages with the OpenOLT Adapter via Google Remote Procedure Call (gRPC). Based on the messages exchanged, the Agent configures the OLT hardware and writes, e.g., flow rules into the PON and switching application-specific integrated circuits (ASICs). Also, ONT management-related messages such as ONT Management and Control Interface (OMCI) are exchanged from the PON chip through the Agent with the VOLTHA control plane stack. Though VOLTHA enables SDN for OLTs, our analysis has shown the limitation of the current release 2.12. The current release supports only one uplink interface to the BNG, while at least two uplinks would be required for redundancy.

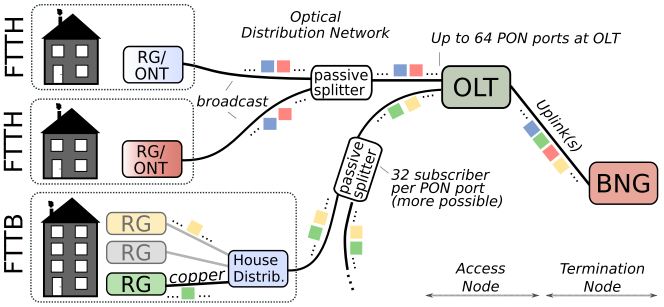

Fig. 2: Passive Optical Network.

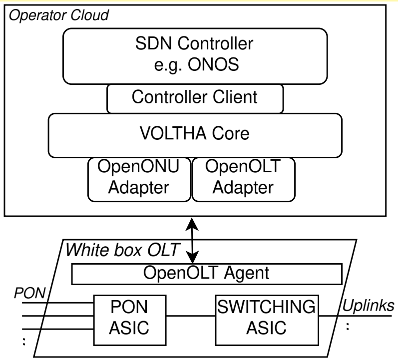

Fig. 3: Simplified VOLTHA software stack with OpenOLT Agent running on white box OLT, derived from [4].

In GPON deployments, the termination node is called BNG, which offers a standardized set of functions [6]. The termination node has the main functionality as the first L3 connection point in the network, terminating the subscriber sessions [7]. Besides authentication, IP address assignment, and billing, other functionality such as Quality of Service (QoS) enforcement is offered. The PPPoE protocol can be used for end-to-end subscriber sessions, while other deployments are possible. For one major German ISP, in 2016, there were around 900 BNG locations in operation, serving about 28 million households [8], [9].

III. DESIGN

To create a resilient access network, each part of the access network must be considered for resilience improvements.

---

*https://opennetworking.org/onos/

Fridolin Siegmund, Philip Jonas Franz, Bjoern Nagel, Maik Rüder, Lisa Wernet, Matthias Hollick, Ralf Kundel. *State-aware Subscriber Steering in Fiber Access Networks for Improved Resilience.* To appear in the Proceedings of IEEE Conference on Network Function Virtualization and Software Defined Networks 2024, IEEE, 2024.

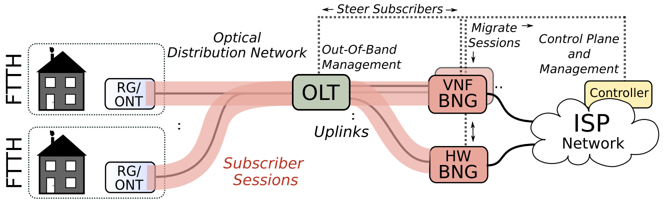

Fig. 4: Improved resilience in PON access networks with subscriber steering and hybrid BNG architecture.

While for disaggregated Virtual Network Functions (VNFs), the resilience of each network function improves the overall resilience [10], the same assumption can be applied to hardware-based nodes. For virtualized environments, the resilience improvement is mostly done by replacing functions during runtime [11]. Hardware-based approaches such as PON access networks rely entirely on hardware for data plane forwarding, where redundancy is more challenging than software instances. Although not all hardware nodes can be replaced easily, the links between hardware-based functions can be built redundant. Also, the combination of hardware-based functions and Network Functions Virtualization (NFV) enables a hybrid approach with resilience advantages of both worlds. For instance, in case of failure, a hardware-based function is replaced by a VNF to restore operation but with reduced QoS. The costs for redundant VNFs are noticeably reduced compared to multiple hardware appliances, especially when the redundant instances are only used as a fallback option.

PON-based access networks reach subscribers in 3 phases: Starting at the subscribers, the data packets of the sessions are forwarded by the PON to OLT. Following, the OLT directs the packets to the BNG. Last, the BNG as the first IP point in the network terminates the subscriber sessions and forwards to and from the ISPs network, which provides Internet access.

In this work, we propose an architecture focusing on phases two and three while we consider phase one in Section VI as part of related work. We accomplish the resilience improvements by steering the subscriber flows through different uplinks at the OLT to different BNGs. Steering events are triggered reactively upon the detection of link failures and proactively for purposes such as scheduled maintenance of links and BNGs. Also, steering events are caused to load balance subscriber flows between uplinks and BNGs.

Figure 4 shows the proposed architecture in the access network. A centralized controller in the ISP network is responsible for the overall operation, including handling node-specific control planes and instance managing. The modular controller has a core module and submodules for abstracting communication to the OLT and BNG nodes.

For the control plane functionality, the control plane module is responsible for:

- Performing PON control plane tasks
- Performing PPPoE control plane tasks
- Flow rules configuration (add, delete, modify)

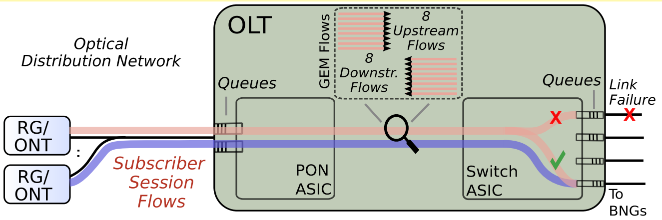

Fig. 5: Upstream and Downstream subscriber flows in OLT.

At the management plane, the management module handles:

- Surveillance of node status
- Monitoring of link load and status
- VNF instance managing

For example, upon link failure detection at the OLT by the management module, the core module can reactively steer subscriber flows to a different uplink interface. To steer a subscriber flow within the OLT, the flow rules in the node are modified through the control plane module. If a different BNG terminates the new path, the core module migrates the subscriber sessions to the new BNG. How long the failure impacts the subscriber access until access is restored depends on the degree of state redundancy. For instance, the controller could proactively mirror all the subscriber sessions to all connected BNGs before link failure occurs to reduce the downtimes.

For scheduled events, the controller can, e.g., migrate all subscriber flows to a different uplink and/or BNG to allow for planned maintenance of both. Also, if the monitoring module reports high link utilization, the subscriber flows could be migrated to another link by the controller to prevent reduced service quality. Additionally, the controller can instantiate new BNG VNFs to distribute the load between instances. However, as our architecture is capable of path adaption in the OLT and session migration between BNGs, the combination of softwarized and hardware-accelerated BNGs opens up. On top of the multiple BNG VNFs, a programmable hardware BNG completes the access network in high-load scenarios. Besides offering excellent QoS, the hardware BNG improves the heterogeneity, which reduces the risk of bug-caused network disruptions due to homogeneous implementations.

IV. IMPLEMENTATION

The designed controller is partly implemented in the proposed prototypes. The prototypes consist of the OLT, respective software stack implementations, and BNG implementations with respective controller functionality.

A. Subscriber Steering in the OLT

Considering the OLT, we modified the VOLTHA OLT software stack to circumvent the limitation of one usable uplink interface, allowing the use of all uplink interfaces. Therefore, we modified the code for the VOLTHA SDN controller (ONOS) and VOLTHA core components and components on the white box OLT itself to include the uplink interface specification in all functionality.

Fridolin Siegmund, Philip Jonas Franz, Bjoern Nagel, Maik Rüder, Lisa Wernet, Matthias Hollick, Ralf Kundel. *State-aware Subscriber Steering in Fiber Access Networks for Improved Resilience.* To appear in the Proceedings of IEEE Conference on Network Function Virtualization and Software Defined Networks 2024, IEEE, 2024.

As Figure 5 shows, each provisioned subscriber receives connectivity by establishing datapath flows through the OLT. The OLT consists of a PON ASIC for subscriber access with PON and a switching ASIC that connects the PON chip to the uplink interfaces. Logically, the datapath connectivity is established by one upstream and one downstream flow. Each flow group subflows individually to the GPON configuration and GPON Encapsulation Method (GEM) port mapping in hardware. Our OLT deployment consists of 8 GEM ports per flow direction. This low-layer configuration is abstracted to the SDN controller by grouping logical upstream and downstream flows. The logical uplink flow rules match the ingressing subscriber port and contain the egress uplink interface as an action. For the downstream flow rules, no uplink interface specification is required.

Consequentially, steering subscriber flows to a different uplink interface requires only the eight uplink flows per subscriber to be modified. Our investigation has shown that the switching chips in current OLTs do not support flow modifications. Following the SDN paradigm, the VOLTHA control plane stack manages all OLT forwarding decisions and configuration control. Before subscriber flows are added, traffic schedulers and queues for the configured uplink interface and subscriber port are created. Therefore, to steer subscriber flows, the old flow, traffic queues, and schedulers are removed and added again with a new uplink interface configuration, causing a short downtime for the subscriber datapath. The traffic schedulers and flow operations are extended by the uplink interface configuration in our prototype.

The analysis of our VOLTHA deployment with ONOS as an SDN controller shows a downtime of more than one second while steering a subscriber flow, which is caused by the multiple abstraction layers and disaggregation. To mitigate the communication delay in the VOLTHA stack, we implemented a custom controller for subscriber steering directly connected to the OLT. As shown in Figure 6, the VOLTHA stack deploys and handles subscribers as designed, but the custom controller is also connected via gRPC to the OpenOLT Agent. This way, the custom controller can send a dedicated flow modification message to the OLT. The flow modification message triggers a custom function in the OLT that executes all flow modification steps in one batch instead of each step triggered separately, as executed with ONOS before. Our analysis shows that batching reduces the communication from 12 Round-Trip Times (RTTs) communication delay down to one RTT.

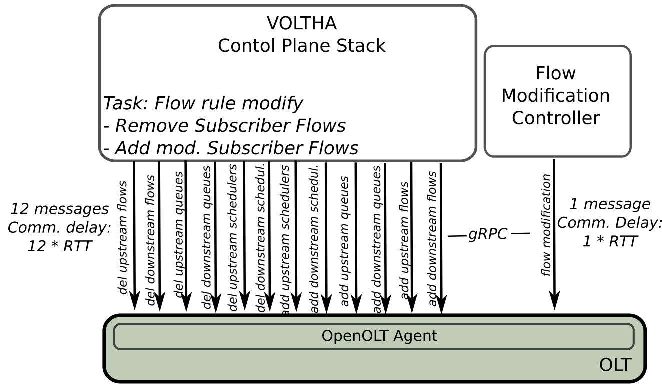

Fig. 6: Custom controller for flow rule modification, mitigating communication delay in VOLTHA control plane stack.

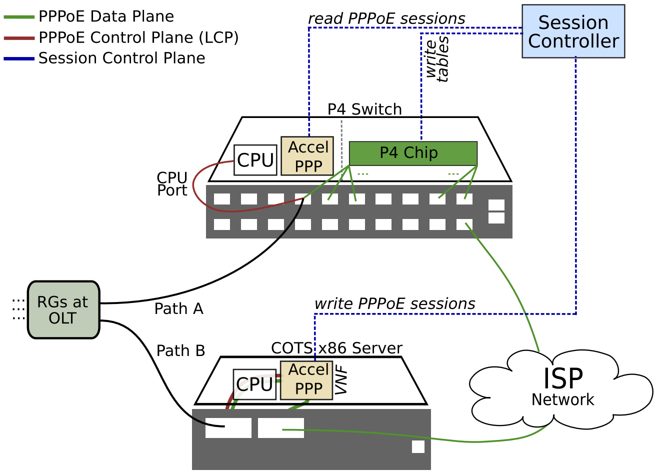

Fig. 7: Hardware-accelerated and x86-based BNG functionality in subscriber-steered access network.

B. *Redundant and hybrid BNG*

Subsequently, we enabled the steering of subscribers between BNG instances, which are the x86 VNF PPPoE server accel-ppp[†] and a hardware-accelerated solution on P4 programmable switches (derived from [3]). As shown in Figure 7, a session controller is aware of established sessions in the BNG instances. Internally, the P4 switch-based BNG handles the session establishment on the x86 CPU while the programmable forwarding chip accelerates the data plane. The BNG VNFs on the commercial off-the-shelf (COTS) x86 server handles the packet processing and session establishment purely in software. The session controller migrates the subscriber session between all instances, enabling redundancy in the hybrid BNG deployment.

V. EVALUATION

To evaluate our prototypes, we analyze the downtime for subscribers during a steering event. Both OLT and BNG implementations are measured separately to observe the downtime characteristics of each node.

A. *Subscriber flow steering in the OLT*

Subscriber flows in the OLT can be distinguished by downstream and upstream flows. Due to only upstream flow rules matching on the uplink interface, there is no compulsory downtime at the data plane when changing the uplink interface. For this reason, we only measure the downtime for subscriber flows in the upstream direction.

Figure 8 shows the measurement setup built with the P4STA [12] framework, which allows to timestamp the subscriber

---

[†] https://accel-ppp.org/

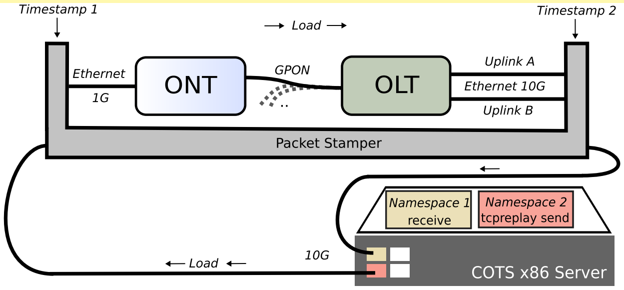

Fig. 8: Setup with P4STA [12] stamper to measure subscriber flow steering under load in OLT.

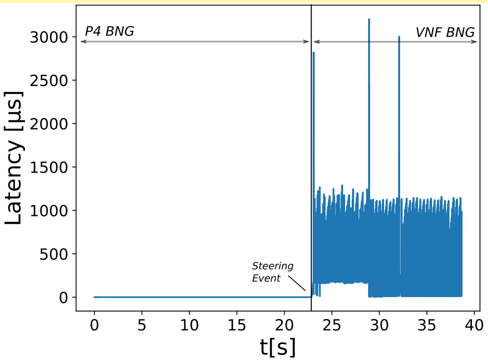

Fig. 10: Latency measurement of BNG data path in case of steering event from programmable hardware to VNF BNG.

packets with nanosecond precision. In the upstream direction, the software tcpreplay[‡] is utilized to replay 1454 byte packets into the ONT. To prevent packet bursts by tcpreplay influencing the measurement, the stamper shapes the traffic according to the desired subscriber flow test. Before ingressing into the ONT, the stamper inserts the first timestamp into the packets. As RG/ONT a *AVM Fritz!Box 5590 Fiber* configured with 1 Gbps down- and upstream bandwidth profiles to prevent influences in the measurement by, e.g., traffic shaping in the OLT, is used. After egressing the OLT at one of the both connected uplinks, the packet is timestamped a second time. As OLT, the *RLT-1600X* from Radisys is used.

The resulting timestamps retrieved from the P4STA framework allow calculating QoS metrics such as latency and jitter. Also, the time delta between two consecutive packets is calculated, so gaps in the transmissions during flow path steering are measured.

Figure 9 shows the distribution of measured downtimes of 200 flow path changes. The average downtime is 155.6 ms for the 100 Mbps flow, with a standard deviation of 19.8 ms showing variance for the downtimes. As minimum downtime, we measured 133.2 ms, ranging up to 216.7 ms as maximum downtime. We observed the same downtime characteristics for the 500 Mbps and 900 Mbps flows, with a difference of less than 1 ms for the average downtime and the respective standard deviation compared to the 100 Mbps flow.

The analysis of our prototype shows that the hardware limits the subscriber steering to sequential processing of flow rule modifications without the possibility of the parallel steering of subscribers. Therefore, in case of reacting to the failure of one uplink, all subscribers connected through the uplink are steered sequentially. In the sequential flow rule modification, the last subscriber has a downtime of about N * 155 ms for N subscribers. The throughput of the subscriber flow does not influence the downtime in case of a steering event. For load balancing or planned maintenance, the downtime is equally, on average, 155 ms for each steered subscriber.

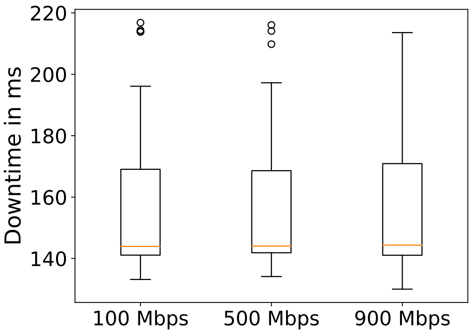

Fig. 9: Distribution of downtimes for 200 subscriber steering events with one subscriber.

B. Evaluation of redundant and hybrid BNG architecture

To solely measure the BNGs when steering subscriber flows, the P4STA framework [12] is used again but modified to steer the flows between the BNGs. The session controller proactively migrates the established PPPPoE session from the hardware-accelerated PPPoE server to the VNF. The software PPPoE client is capable of generating around 7 Gbps of traffic, magnitudes below the limitations of the hardware BNG. Figure 10 shows the latency of the subscriber flow through the BNG while switching from hardware BNG to the VNF. After the steering event, the latency spikes from below 1 μs up to more than 3 ms while the throughput is reduced to around 4.3 Gbps. Due to the proactively migrated sessions, no packet loss is observed.

C. Analysis

The evaluation shows that only the OLT induces downtime (packet loss) for subscriber flows when steered. For the hybrid BNG setup, our prototype has shown zero downtime in the best case with proactively migrated sessions. Therefore, the overall downtime is 155.6 ms per subscriber when the flow is steered to a different BNG through a different uplink.

Except in the best case, in real-world scenarios such as link failure, the downtimes are subject to several external factors. These factors are, among others, controller processing and communication delay, failure detection time, and implementation specifics. In relation to the measured downtime for flow steering with more than 100 ms, the additional external delay factors are below 50 ms for one subscriber.

[‡] http://tcpreplay.appneta.com/

Fridolin Siegmund, Philip Jonas Franz, Bjoern Nagel, Maik Rüder, Lisa Wernet, Matthias Hollick, Ralf Kundel. *State-aware Subscriber Steering in Fiber Access Networks for Improved Resilience.* To appear in the Proceedings of IEEE Conference on Network Function Virtualization and Software Defined Networks 2024, IEEE, 2024.

VI. RELATED WORK

In PON, multiple designs consider resilience in the fiber access network. Improved resilience can be achieved by creating redundancy in the fiber links as well as PON ports at the OLT [13]. Also, monitoring faulty devices or attackers in the PON could increase the resilience [14].

With an abstracted and disaggregated access network, the work-in-progress specification WT-474 presents an architecture for subscriber steering in access networks [15], [16]. However, implementations still need to be found in academia and industry. Different network functions such as the Traffic Steering Function (TSF) and User Plane Selection Function (UPSF) steer the subscriber flows through the network. Assuming OLTs, the TSF would be placed in the OLT to steer flows to specific uplink interfaces, while the UPSF represents the controller, similarly to our proposed solution.

While [11] researched the virtualization of BNG functionality to enhance reliability by redundancy and state migration, [3], [17] showed that the data plane traffic could be offloaded to programmable hardware. Our research has shown how both approaches can be combined with the subscriber steering architecture to create path redundancy and load adaption.

VII. CONCLUSION AND FUTURE WORK

In this work, we investigated the capabilities of resilience improvements in fiber-based access networks by subscriber steering. We proposed a resilient access network, starting from the OLT through the BNG to the data network. By eliminating single points of failure with redundant links and different redundant BNG instances, resilience and performance metrics are significantly improved with the proposed load and failure adapting architecture. The presented prototype shows the steering possibilities and also limitations of current white box OLT hardware on the market.

The performed improvements on the VOLTHA software stack have partially already been merged back into the open-source community or are already in the contributions process.

The analysis of VOLTHA-supported OLTs on the market has shown limitations in SDN-enabled flow operations such as flow modifications. To enable more flexibility in the access network data plane, improved switching silicons, programmable forwarding chips, or Field Programmable Gate Arrays (FPGAs) could be combined with the PON silicon in OLTs. For example, the P4 programmable Tofino silicon with high-throughput optimized queueing has shown zero packet loss for flow steering with changing egress ports in our experiments.

Further, to reduce reaction time to uplink failures in the OLT, the controller could be disaggregated and partly run on the x86 CPU of the OLT. The *onboard* OLT steering function could react to link failures by immediately steering subscriber flows to a different uplink and subsequently synchronizing the state with the global SDN controller. This approach may partly mitigate control plane communication disruptions.

ACKNOWLEDGEMENT

This work has been supported by Deutsche Telekom through the Dynamic Networks 11 and 12 projects and partly supported by the LOEWE initiative (Hesse, Germany) within the emergenCITY center [LOEWE/1/12/519/03/05.001(0016)/72]. Furthermore, we thank our colleagues and reviewers for their valuable input.

REFERENCES

[1] M. Jarschel, T. Zinner, T. Hoßfeld, P. Tran-Gia, and W. Kellerer, “Interfaces, Attributes, and Use Cases: A Compass for SDN,” *IEEE Communications Magazine*, vol. 52, no. 6, pp. 210–217, 2014.

[2] J. Montalvo, J. Torrijos, D. Cortes, R. Chundury, and M. S. Peter, “Journey toward software-defined passive optical networks with multi-PON technology: an industry view,” *Journal of Optical Communications and Networking*, vol. 13, no. 8, pp. D22–D31, 2021.

[3] R. Kundel, L. Nobach, J. Blendin, W. Maas, A. Zimber, H.-J. Kolbe, G. Schyguda, V. Gurevich, R. Hark, B. Koldehofe, and R. Steinmetz, “OpenBNG: Central office network functions on programmable data plane hardware,” *International Journal of Network Management*, vol. 31, no. 1, pp. 1–25, 01 2021.

[4] O. N. Foundation, *Virtual OLT Hardware Abstraction v2.12*, 2024. [Online]. Available: https://docs.voltha.org/master/index.html

[5] L. Peterson, C. Cascone, and B. Davie, *Software-Defined Networks: A Systems Approach*. Systems Approach, LLC, 2021.

[6] K. Wan and N. S. Turaiyur, “Multi-Service Disaggregated BNG with CUPS. Reference Architecture, Deployment Models, interface, and Protocol Specifications,” Broadband Forum, Tech. Rep. TR-459i2, 04 2023, https://www.broadband-forum.org/technical/download/TR-145.pdf.

[7] A. Cui and Y. Hertoghs, “Multi-service Broadband Network Functional Modules and Architecture,” Broadband Forum, Tech. Rep. TR-145, 11 2012, https://www.broadband-forum.org/technical/download/TR-145.pdf.

[8] K. Hackbarth, G. Kulenkampff, T. Plückebaum, and K. Zoz, *Analytisches Kostenmodell für das Breitbandnetz*. Bundesnetzagentur, Germany, 2016.

[9] D. T. AG, “Das Geschäftsjahr 2016,” Online, 2017.

[10] 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.

[11] T. Dietz, R. Bifulco, F. Manco, J. Martins, H.-J. Kolbe, and F. Huici, “Enhancing the BRAS through virtualization,” in *Proceedings of the 2015 1st IEEE Conference on Network Softwarization (NETSOFT)*. IEEE, 2015, pp. 1–5.

[12] R. Kundel, F. Siegmund, J. Blendin, A. Rizk, and B. Koldehofe, “P4STA: High Performance Packet Timestamping with Programmable Packet Processors,” in *Proceedings of the IEEE/IFIP Network Operations and Management Symposium (NOMS)*. IEEE, 04 2020, pp. 1–9.

[13] C. Christodoulou and G. Ellinas, “Resilient architecture for optical access networks,” *Photonic Network Communications*, vol. 41, pp. 1–16, 2021.

[14] V. Oujezsky, T. Horvath, and M. Holik, “Security Incident Response Automation for xPON Networks,” *Journal of Communications Software and Systems*, vol. 18, no. 2, pp. 144–152, 2022.

[15] Broadbandforum, “WT-474: Introducing Subscriber Session Steering and dynamic subscriber placement,” Online, 2022, https://www.broadband-forum.org/wt-474-introducing-subscriber-session-steering-and-dynamic-subscriber-placement.

[16] J. Newton and P. Mannion, “Transforming Fixed Access Using WT-474 Traffic Steering on a Cloud-Native Architecture,” Tech. Rep., 8 2023, https://cdrdv2-public.intel.com/780786/cloud-native-WT474-UseCase_v3.pdf.

[17] R. Kundel, L. Nobach, J. Blendin, H.-J. Kolbe, G. Schyguda, V. Gurevich, B. Koldehofe, and R. Steinmetz, “P4-BNG: Central Office Network Functions on Programmable Packet Pipelines,” in *Proceedings of the 15th International Conference on Network and Service Management (CNSM)*. IEEE, 10 2019, pp. 1–9.