Experimental Demonstration of a Packet-based Protection for Seamlessly Recovering from a Multi-layer Metro Network Fronthaul Failure

K. Kondepu[1], S. Ramanathan[2], M. Tacca[2], M. Razo[2], B. Mirkhanzadeh[2],
F. Giannone[1], L. Valcarenghi[1], and A. Fumagalli[2]
[1]Scuola Superiore Sant'Anna, Pisa, Italy; Email:k.kondepu@sssup.it
[2]Open Networking Advanced Research (OpNeAR) Lab, The University of Texas at Dallas

***Abstract*—Packet loss in the fronthaul adversely affects radio services, especially the low-layer functional split options between Distributed Unit (DU) and Central Unit (CU). A 1+1 protection scheme is implemented in a Software Defined Networking (SDN) packet-over-optical transport test-bed, achieving seamless recovery of the 7-1 split option from an optical link failure in the fronthaul.**

***Index Terms*—5G, Resiliency, Functional split, SDN, Protection, GENI.**

I. Introduction

A number of functional split options is defined in 3GPP TR 38.801 [1], where Next Generation NodeB (gNB) functions are physically separated and hosted by two distinct entities, i.e., the Distributed Unit (DU), deployed next to the radio antenna, and the Central Unit (CU), deployed at a centralized location . The DU and CU are connected through the fronthaul, which provides both control and data plane communications. DU, CU, and fronthaul form the so called 5G next generation radio access network (NG-RAN), also referred to as centralized/cloud RAN (C-RAN).

The TR 38.801 specifications define the fronthaul one-way delay and capacity requirements. However, they do not address reliability and packet loss. Methods have been proposed to handle different types of failure in C-RAN, accounting for both hardware (e.g., fronthaul network link failure) and software (e.g., virtual machine crash). Through simulation, partial protection of the access cloud network (ACN) is shown to require only 8% additional network resources by accepting a certain level of degraded services [2]. A two-step recovery scheme makes use of both lightpath transmission adaptation and gNB functional split reconfiguration to handle degradation of transmission quality in the fronthaul optics [3]. This scheme maintains the connectivity between virtualized DU (vDU) and virtualized CU (vCU), while accounting for the fronthaul capacity constraints. Fronthaul network failures may also be overcome by leveraging traditional underlying network recovery schemes [4], [5].

Most proposed schemes do not guarantee a seamless recovery of the DU-CU connection. In fact, it is not uncommon to experience some packet loss during a network recovery procedure, as the network takes time to both detect the link outage and reroute the disrupted connection. Such packet loss is likely to cause DU-CU loss of synchronization, which in some instances may lead to disconnection [1]. Reconnecting DU and CU is a lengthy procedure. Described next, a packet-based 1+1 protection scheme is shown to achieve lossless recovery of the DU-CU connection in an SDN packet-over-optical fronthaul test-bed.

II. Upgrading the PROnet SDN Test-bed with a Packet-based 1+1 Protection Scheme

The Ethernet-over-WDM PROnet SDN test-bed [12] is used to carry out the experiment. Specifically, a newly designed kernel software module is added to implement the 1+1 protection scheme at the Ethernet layer, as this function is not available in openflow/Open V-Switch (OVS) [6], [7]. The correct functionality of the 1+1 protection scheme is tested using the Intra-PHY split (or split option 7-1) C-RAN while failing one of the test-bed fiber links.

Fig. 1 shows an example that illustrates how the 1+1 protection scheme works. Two fiber link-disjoint optical circuits (i.e., primary (1) and secondary (2)) are computed and provisioned by the PROnet SDN resource orchestrator [12]. The ingress packet flow (e.g., the flow that is coming from the transmitting DU[1]) is duplicated and transmitted over the two optical circuit interfaces by the ingress OVS using the *all group table function* [6]. At the egress OVS, the two copies of the flow arriving from the two optical circuits are processed in real-time to ensure that only one copy of each packet is forwarded to the receiving CU.

The egress OVS processes the packets from the two incoming flows using netfilter pre-routing hook [8]. Each time a packet is captured by the hook, the following actions are performed: (i) a unique signature is created for the packet using some of the radio link control (RLC) packet field(s), like frame, subframe, and symbol fields; (ii) if the packet signature is found in the check list, the packet is discarded and the signature value is removed from the list[2]; (iii) else the packet is passed onto the switch, which performs the necessary forwarding, and the packet signature is stored in the check list. To prevent check list overflow, the software also removes the

[1]The same technique is applied to the packet flow in the opposite direction of propagation, i.e., from CU to DU.

[2]The module assumes that a packet can only be duplicated once and no packet will be received more than twice.

oldest signatures from the check list when the list exceeds a predefined number of entries (e.g., 100).

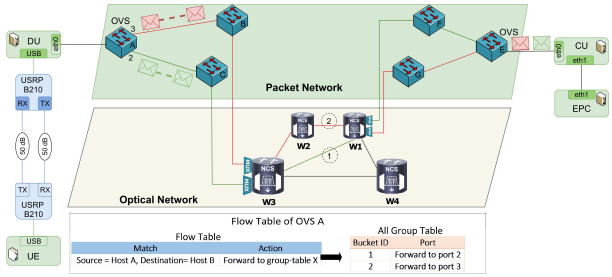

Fig. 1: 1+1 protection scheme implemented in the PROnet testbed

III. C-RAN Implementation

The C-RAN part of the test-bed is implemented using the following components. The open source OpenAirInterface (OAI) platform [9] is utilized as mobile network software. OAI provides an implementation of few NG-RAN functional splits as defined in 3GPP TR 38.801 [1], including the functional split Option 7-1 used in this experiment. As shown in Fig. 1 two Ettus B210 radio front-end boards and four Intel-i7 servers (implementing User Equipment (UE), DU, CU, and Evolved Packet Core (EPC), respectively) host the OAI software modules. OAI EPC implements the following network elements: the Serving Gateway (S-GW), the PDN Gateway (PDN GW), the Mobility Management Entity (MME) and the Home Subscriber Server (HSS). All these OAI core elements can be deployed as individual virtualized elements or can also be deployed as bundle virtualized EPC (vEPC). In the demonstration, the bundle vEPC is utilized.

The four Intel-i7 servers are equipped with 1GE NICs, which suffice to support the (single) UE traffic demand over a 5MHz radio channel bandwidth. Both DU and CU servers are connected to a distinct OVS, equipped with multiple 1GE NICs and controlled through Openflow 1.3. Each OVS is connected to a distinct Dell N2048 switch, which in turn provides a 10GE connection to a Cisco optical muxsponder. The DWDM layer consists of four Cisco NCS 2000 ROADM nodes that are controlled and dynamically configured through RESTconf [10] to access YANG data. The entire test-bed is controlled by the PROnet SDN orchestrator, which ensures coordination of resource reservation across the two network layers [11], [12].

IV. Experimental Results

The experiment is conducted with the newly implemented 1+1 protection scheme. As shown in the Fig. 1, to obtain the initial experimental results all of the OAI RAN and core components are deployed in PROnet test-bed only. Fig. 2 shows the Wireshark output at the CU. The CU server (IP address 192.168.0.134) is set to transmit periodic echo request ICMP packets to the DU server (IP address 192.168.0.136) over the 1+1 protected flow, at intervals of about 200ms. Before the optical fiber failure is induced, the 1+1 kernel module at the egress OVS is instructed to drop duplicate packets of the C-RAN protocol, while duplicated ICMP (ping) packets are purposely not removed. As a result, when pinging the DU from the CU server, both echo request and reply packets are duplicated, for a total of 4 echo reply packets generated (i.e., sequence number 23) per ping request, as shown in Fig. 2.

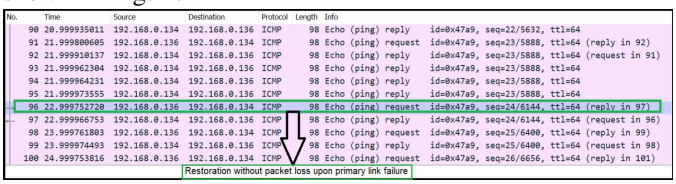

Fig. 2: 1+1 protection scheme output captured by Wireshark packet sniffer at the server hosting CU

Once the fiber failure is induced and one of the two optical circuits is disrupted, both ICMP request and reply packets are received only once, as only one copy of each packet can make it through the partially failed DU-CU connection. By inspecting the Wireshark output, it is easy to identify the exact moment when the optical circuit is disrupted. The contiguous sequence numbers reported in the trace reveal that none of the C-RAN packets is lost during the experiment, including the moment when the optical circuit is disrupted.

V. Conclusions

The contribution of the presented experiment is twofold. First, it was determined that DU-CU connection recovery time and packet loss highly matter when supporting lower-layer functional split (e.g., split option 7-1). Second, a newly implemented kernel module provided a packet-based 1+1 protection scheme in a Ethernet-over-DWDM SDN test-bed. The implemented scheme is able to overcome a single fiber failure in the fronthaul without causing any packet loss, thus ensuring that the DU and CU modules remain connected during the recovery phase.

Acknowledgment

This work has been partially funded by the EC H2020 “5G-Transformer” Project (grant no. 761536) and NSF grants no. CNS-1405405, CNS-1409849, ACI-1541461, and CNS-1531039.

References

[1] 3GPP TR 38.801, “Study on new radio access technology; radio access architecture and interfaces”, V2.0.0 (2017-03).
[2] C. Colman-Meixner et al., “Resilient cloud network mapping with virtualized BBU placement for cloud-RAN”, in Proc. of ANTS (2016).
[3] K. Kondepu et al., “Orchestrating Lightpath Recovery and Flexible Functional Split to Preserve Virtualized RAN Connectivity”, J. Opt. Commun. Netw. 10, 843-851 (2018).
[4] L. Valcarenghi and A. Fumagalli, “IP restoration vs. WDM protection: Is there an optimal choice?”, IEEE Network, vol. 14, no 6 (2000).
[5] A. Giorgetti et al., “Segment routing for effective recovery and multi-domain traffic engineering”, J. Opt. Commun. Netw. 9 (2017).
[6] “OpenFlowSwitch Specification”, www.opennetworking.org/wp-content/uploads/2013/04/openflow-spec-8.
[7] “Open vSwitch”, www.openvswitch.org/.
[8] “Netfilter.org Project”, www.netfilter.org.
[9] “OpenAirInterface”, gitlab.eurecom.fr/oai/openairinterface5g/wikis/home.
[10] M. Bjorklund et al., “RESTCONF Extensions to Support the Network Management Datastore Architecture”, https://tools.ietf.org/html/rfc8527.
[11] D. Hicks et al., “PROnet: A programmable optical network prototype”, in Proc. of ICTON (2016).
[12] B. Mirkhanzadeh et al., “An SDN-enabled multi-layer protection and restoration mechanism”, Optical Switching and Networking (2018).

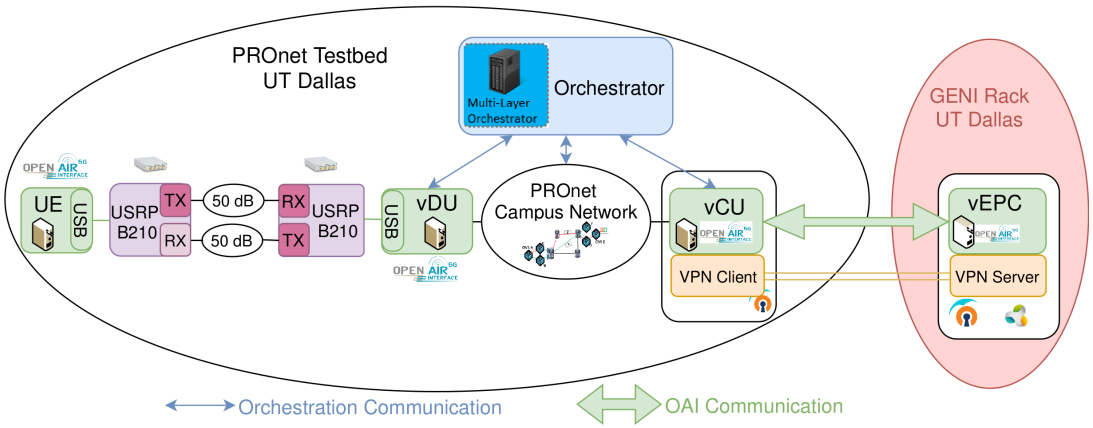

Fig. 3: Demonstration Setup

Description of Live Demonstration

The fronthaul reliability is demonstrated live by using the C-RAN configuration as shown in Fig. 3 with the following network connections: (i) the connection between the virtualized DU (vDU) and the virtualized CU (vCU) is implemented in the PROnet Etherenet-over-DWDM test-bed, and (ii) the connection between the vCU (hosted in the PROnet test-bed) and the virtualized EPC (vEPC), which is hosted by the UT-Dallas GENI rack, is implemented using a Virtual Private Network (VPN) tunnel. The vDU and vCU are deployed using Docker Container virtualization technology, while the vEPC is deployed on a reserved virtual machine (VM) with public IPv4 (XEN VM). The demonstrated vEPC platform may help Mobile Network Operators (MNOs) add capacity and flexibility to their mobile network infrastructure more easily.

The PROnet test-bed resources can be accessed from the reserved GENI XEN VM, through the central authority (iMinds) tool such as jFed account. jFed is a useful tool to configure experiments that require interconnection of resources from multiple test-beds, reserve and access their resources. As shown in Fig. 3, all of the OAI virtualized RAN and core components are used in the live demonstration. An OpenVPN server is configured at the reserved XEN VM node, which hosts the OAI vEPC module. On the other hand, an OpenVPN client is configured at the vCU hosted in the PROnet test-bed. Consequently, the OAI components in the PROnet test-bed and the PROnet SDN orchestrator can be accessed through Secure Shell (SSH).

Shell scripts provide a third party researcher (the user) the opportunity to conduct three distinct experiments from the XEN VM terminals.

Exp1: the user can prepare the PROnet test-bed and OAI C-RAN modules to run the experiment, followed by the disruption of one of the optical circuits connecting the DU to the CU server. A collection of simple shell scripts is used to provision the network resources from the PROnet SDN orchestrator, and to start the C-RAN modules (e.g., vDU, vCU, and the UE) from the reserved XEN VM terminal. Once the UE is connected, a fault in injected by running the shell script — delete lightpath. This script automates the deletion of one of the optical circuits between the optical nodes by making use of RESTful API calls to the optical network L1 service.

Exp2: the user can activate the 1:1 protection scheme, which is implemented using the *fast failover table* function, already available in OpenFlow and OVS. As shown in Fig. 1, the PROnet orchestrator is instructed to provision two fiber link-disjoint network paths (primary path — red, and secondary path — green color). Upon failure of the primary path, the flow is switched from the primary path to the secondary path at the Ethernet layer. Here, the OpenFlow enabled Ethernet switch detects the link failure and reroutes the flows to the other link based on the rules configured in the *fast failover table*. As soon as the fault is induced, the C-RAN stops working and the "disconnect" UE message is shown. This is due to the strict network recovery time and packet loss requirements that must be guaranteed between the vCU and vDU. Moreover, the required recovery time (in the 100ms range) of the fast failover table protection scheme does not guarantee the delivery of all the transmitted packets over the fronthaul. This is due the fact that the switch takes time to detect the link failure and reroute the flow over the secondary path of the 1:1 protection scheme.

Exp3: the user can activate the 1+1 protection scheme (described in Sec. II) based on the newly developed kernel module implementation. With this experiment the user can see that the fronthaul connection is recovered without any packet loss. In addition, the time required to compute unique signatures and carry out comparison with the signatures already stored in the check list can be assessed during the demo, revealing that the proposed duplicate packet method does not contribute significant extra time to the fronthaul latency. The Zabbix tool is used to show failure occurrence and service recovery from the vCU interface viewpoint, along with Wireshark traces.