End-to-End Slicing via O-RAN and Software Defined Optical Access

C. Centofanti[1,*], A. Marotta[1], D. Cassioli[1], F. Graziosi[1], V. Gudepu[2], K. Kondepu[2]

[1] *University of L'Aquila, DISIM Department, L'Aquila, Italy and CNIT National Laboratoy of Advanced Optical Fibers for Photonics, L'Aquila, Italy*
[2] *Indian Institute of Technology Dharwad, CSE Department, Dharwad, India*
[*]*carlo.centofanti1@univaq.it*

Abstract:
We propose an end-to-end slice management strategy which exploits the programmability of O-RAN and software defined optical access. Advantages in terms of user experienced latency and packet loss are experimentally evaluated.

1. Introduction

The huge increase of data traffic calls for the densification of high capacity and heterogeneous mobile and optical network infrastructures. To cope with such a heterogeneity, a vendor-agnostic and software-defined approach for the management and control of the network is essential. Open-RAN (O-RAN) architecture represents a key enabler to achieve vendor interoperability and to implement flexible and intelligent control of mobile networks, which is crucial to meet the ambitious requirements of next generation services [1].

The *openness* of O-RAN paves the way towards the cooperation between RAN and optical transport network infrastructures, which has been shown to offer advantages in terms of network efficiency and savings in cost of ownership for operators [2]. In particular, Passive Optical Networks (PONs), due to their large capillarity, represent a good transport technology to support mobile networks densification. Furthermore, the introduction of Software Defined Optical Access Networks (SDOANs) allows to implement flexible bandwidth allocation strategies in a dynamic and programmable way [3].

Network slicing has emerged as a key technology to offer dedicated virtual networks with performance tailored to specific service or customer over a common network infrastructure. Slice deployment requires the accommodation of *computational resources*, which are represented by Virtual Network Functions (VNFs), and *network resources* as well. In order to offer the targeted performance to the user belonging to a specific slice, all the involved network segments (e.g., RAN, optical access and metro networks, mobile core network) have to be configured accordingly. Recent efforts addressed resource allocation strategies for sliced PONs [4] and SDN-based slice management via NETCONF in PONs [5] but, to the best of our knowledge, no study investigated so far the utilization of O-RAN and SDOAN for end-to-end slicing.

This paper aims to present an end-to-end slice management strategy which leverages the software controllability of O-RAN and SDOAN. A slice manager element interacts with a purposely built xApp in the O-RAN and dynamically adapts resource allocation in the PON to offer low latency and guaranteed bandwidth for specific slices at the backhaul level via NETCONF. The adaptation of reserved bandwidth at the PON avoids both the over-provisioning and under-provisioning of resources in the optical access and increases the network efficiency.

2. System Model

We consider the network architecture shown in Fig. 1 where two slices are deployed in a TDM-PON backhauled O-RAN, with *Slice 1* requiring low-latency for connected users and *Slice 2* with best-effort latency requirement. The O-RAN is represented by a single O-gNB controlled by a Near Real Time (Near-RT) Radio Intelligent Controller (RIC) which is responsible to apply decisions and perform monitoring on a millisecond time scale via the standardized E2 interface. The Near-RT RIC interacts with a Non-RT RIC which provides policies with a time scale ≥ 1 s. On top of the near-RT controller, custom-built applications called *xApps* implement the control logic desired by the network operator via Application Programming Interfaces (APIs) exposed by the near-RT RIC.

A network orchestrator is responsible to deploy core network functions composing the slices at core level, i.e. the Access Management Function (AMF), Network Repository Function (NRF), Session Management Function (SMF), and User Plane Function (UPF). A Slice Manager element is responsible to map requirements of different slices into ad-hoc configurations of the network segments and to coordinate the network controllers and the network orchestrator throughout the slice life-cycle.

A PON controlled via SDN and NETCONF is adopted as backhaul infrastructure for the O-RAN; a VLAN-based traffic differentiation mechanism is implemented to identify backhaul traffic belonging to different slices and

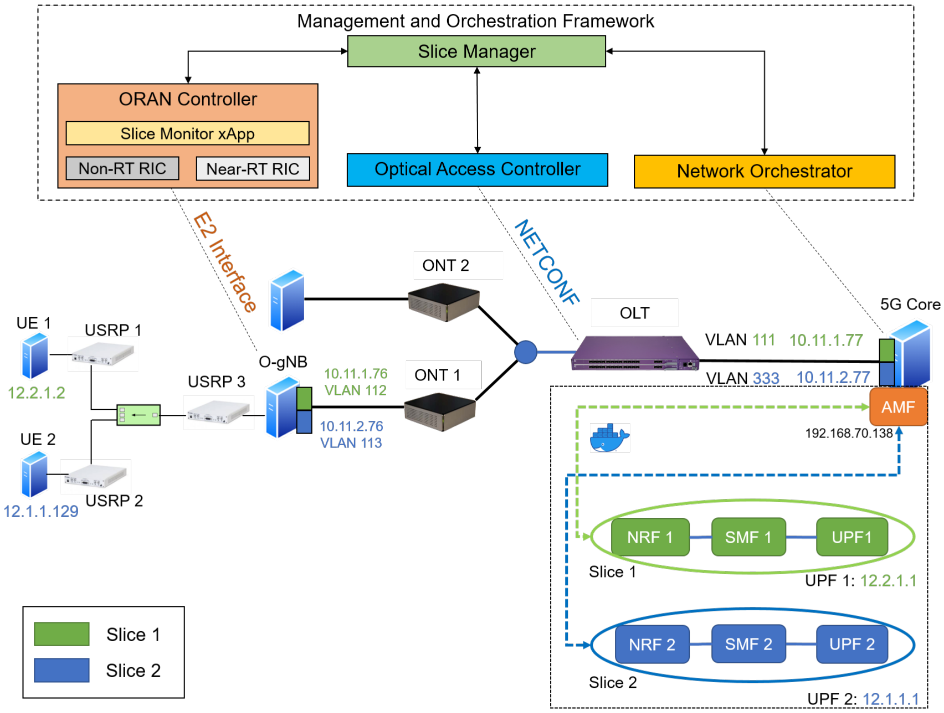

Fig. 1: Network Architecture.

adopt bandwidth allocation strategies basing on targeted slice performance. We consider the following bandwidth allocation strategies: (i) *Expedited Forwarding* which gives to a particular slice the possibility to utilize a reserved amount of bandwidth in a Grant-free fashion in order to reduce experienced latency (adopted for *Slice 1* traffic); (ii) *Request-Grant based access* which leverages Dynamic Bandwidth Assignment (DBA) procedures to exploit statistical multiplexing, thus increasing the network efficiency at the price of a larger experienced latency (adopted for *Slice 2* traffic).

Although the Expedited Forwarding offers low latency, it requires dedicated bandwidth which may lead to inefficient resources usage when slice traffic varies. To preserve resource efficiency while offering low latency, we design a slice manager element which leverages O-RAN and SDOAN programmability to dynamically adapt the amount of reserved low-latency optical access resources. As shown by the sequence diagram in Fig. 2a, the slice manager element is responsible to collect slice traffic statistics from the *Slice Monitor xApp* which runs at the near-RT RIC and obtains the instantaneous slice traffic load at PDCP layer. The slice manager then interacts with the optical access controller to adapt the reserved bandwidth basing on the effective slice conditions with pre-defined thresholds.

3. Experimental Setup and Results

The considered experimental setup primarily consists of three different parts that are the O-RAN, SDOAN and mobile core network, as shown in Fig. 1.

The O-RAN is deployed by using OpenAirInterface (OAI) open-source software stack [6]. We deployed two OAI-based UEs (*UE 1* and *UE 2*) and an O-gNB using National Instrument X310 Universal Software Radio Peripherals (USRPs). The OAI O-gNB is configured with Single – Network Slice Selection Assistance Information (S-NSSAI) list to support two different Slice/Service Type (SST) and Slice Differentiator (SD) values to evaluate the O-RAN slicing scenario. Without loss of generality, we assume that *UE 1* sends traffic over *Slice 1*, i.e. low-latency slice, and *UE 2* sends traffic over *Slice 2*. We adopt flexRIC [7] software as the near-RT RIC. At the mobile core network side, we deploy different sets of functions corresponding to the different slices.

The SDOAN is realized by utilizing a commercial Calix Axos E7-2 XGS-PON that supports NETCONF protocol. Traffic differentiation is based on VLAN tags. The O-gNB applies different VLAN tags to each packet going towards the core network by utilizing virtual ethernet interfaces. *Slice 1* and *Slice 2* are associated with VLANs 112 and 113, and VLANs 111 and 333 at ONT and OLT side, respectively. *Slice 1* is forwarded at the PON using the Expedited Forwarding mechanism while *Slice 2* applies Best Effort. Slice Manager and Slice Monitor xAPP are realized using the *C* language, whereas the Optical Access Controller is realized using *Python* language and is invoked from the Slice Manager to dynamically change the PON bandwidth based on the RAN traffic.

Fig. 2b shows a comparison in terms of average experienced latency measured between the UEs and core

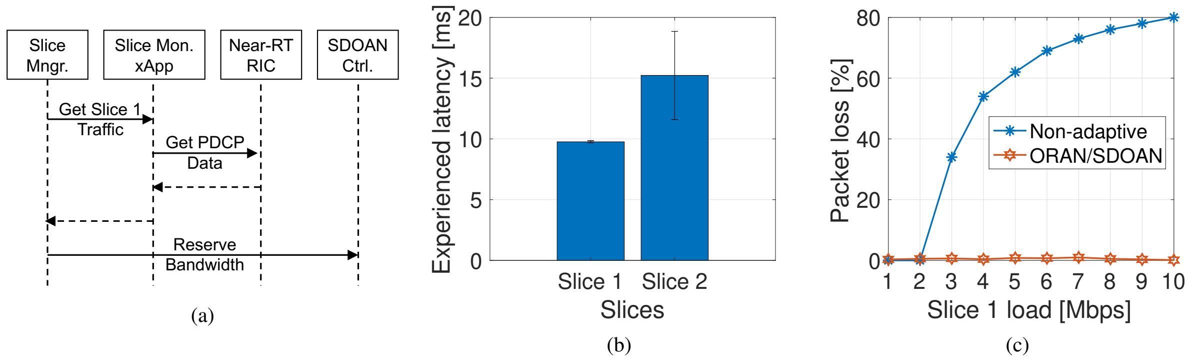

Fig. 2: (a) Slice management sequence diagram; (b) Slice experienced latency; (c) Packet loss comparison.

network. The end-to-end delay measured for *Slice 1* is significantly reduced when compared to *Slice 2*, thanks to the bandwidth reservation and traffic prioritization applied at the PON. It can be noticed that, as an effect of expedited forwarding in PON, delay variation is also reduced. It is also worth mentioning that such delays include the contribution of the O-RAN (which is impacted by OAI software implementation performance). Fig. 2c shows a comparison in terms of packet loss between the proposed slice management approach (exploiting O-RAN and SDOAN programmability) and a non-adaptive strategy, where bandwidth reservation at the PON is configured *a priori*. For the non-adaptive case, we assume 1 Mbps reserved bandwidth in the PON for *Slice 1*. We vary the traffic of *Slice 1* and measure packet loss by running *iperf* between *UE 1* and *UPF 1*. As described in Sec. 2, to offer low latency to *Slice 1*, bandwidth has to be reserved. When the traffic required by the slice exceeds the bandwidth reserved in the PON, packet loss is experienced since packets are dropped at the O-gNB and ONT without being forwarded. To this purpose, the proposed O-RAN/SDOAN based slice manager dynamically adapts the reserved bandwidth in the PON. As shown in Fig. 2c, this allows the network to keep packet loss limited while offering low-latency to a specific slice compared to a non-adaptive approach where the packet loss increases with the traffic load of the slice.

4. Conclusion

In this paper we propose an end-to-end slice management strategy which exploits O-RAN to retrieve information on radio access network load for a specific slice and dynamically adapts the bandwidth at the optical access network. Results show that the proposed mechanism allows the network to offer low end-to-end latency while reducing experienced packet loss and preserving optical access network efficiency.

Acknowledgment

This work has been partially supported by the Italian Government through projects FIRST, SICURA and INCIPICT, by the European Commission through the project H2020-MSCA-RISE-2019 OPTIMIST, and by DST SERB Startup Research Grant (SRG-2021-001522).

References

1. M. Polese, L. Bonati, S. D'Oro, S. Basagni, and T. Melodia, "Understanding O-RAN: Architecture, interfaces, algorithms, security, and research challenges," arXiv preprint arXiv:2202.01032 (2022).
2. S. Mondal and M. Ruffini, "Optical Front/Mid-haul with Open Access-Edge Server Deployment Framework for Sliced O-RAN," IEEE Trans. on Netw. Serv. Manag. (2022).
3. A. Marotta, K. Kondepu, D. Cassioli, C. Antonelli, L. M. Correia, and L. Valcarenghi, "Software defined 5G converged access as a viable techno-economic solution," in *2018 Optical Fiber Communications Conference and Exposition (OFC),* (IEEE, 2018), pp. 1–3.
4. S. Das and M. Ruffini, "Optimal virtual PON slicing to support ultra-low latency mesh traffic pattern in MEC-based Cloud-RAN," in *2021 International Conference on Optical Network Design and Modeling (ONDM),* (IEEE, 2021).
5. C. Centofanti, A. Marotta, D. Cassioli, F. Graziosi, N. Sambo, L. Valcarenghi, C. Bernard, and H. Roberts, "Slice Management in SDN PON Supporting Low-Latency Services," in *2022 European Conference on Optical Communications (ECOC),* (IEEE, 2022).
6. OAI, "OpenAirInterface: A Flexible Platform for 5G Research," (2022).
7. R. Schmidt, M. Irazabal, and N. Nikaein, "FlexRIC: an SDK for next-generation SD-RANs," in *Proceedings of the 17th International Conference on emerging Networking EXperiments and Technologies,* (2021), pp. 411–425.