A Converged Fixed-Wireless TDMA-based Infrastructure Exploiting QoS-Aware End-to-End Slicing

C. Matrakidis(1), E. Kosmatos(1), A. Stavdas(1), P. Kostopoulos(2), D. Uzunidis(1), S. Horlitz(3), Th. Pfeiffer(3), A. Lord(4)

(1) OpenLightComm Ltd, Adastral Park, Ipswich, IP5 3RE, UK; cmatraki@openlightcomm.uk
(2) Dpt of Informatics and Telecoms, University of Peloponnese, Tripolis, 22100, Greece
(3) Nokia Bell Labs, Lorenzstrasse 10, 70435 Stuttgart, Germany
(4) Applied Research, BT, Polaris House, Adastral Park, Ipswich, IP5 3RE, United Kingdom

Abstract *We report a converged fixed wireless testbed spanning Metro-Access, built with no dedicated or mobile technology. An overarching intelligent orchestrator is implemented to tailor, in real-time, the fixed-line end-to-end QoS-aware bandwidth slices to the wireless system, taking into account the status of the entire network.*

Introduction

The formation of the "Networked Intelligence" (NI)[1] would be a milestone for Industry 4.0 (I4.0). The NI needs to capitalize on a cost-effective, pervasive and ubiquitous 5G/6G Converged Fixed-Wireless (CFW) infrastructure but today, even the modelling of the advanced mobile wireless systems becomes a very hard, if not impossible, task[2]. In turn, this makes it challenging and costly to plan and operate the associated fixed-line connectivity infrastructure based on a paradigm that relies on the over-provisioning of dedicated connectivity resources.

However, since a significant fraction of I4.0 use-cases concern still, or slowly moving, end-users (like humans, IoT terminals, robots, smart factories etc), we consider here an alternative CFW infrastructure. Such a heterogeneous infrastructure consists of a WiFi Access Point (AP) in conjunction with the fixed connectivity platform of[3] that exploits the data-planes of an XGS-PON and GPON system in Metro and Access, respectively.

In this work, we report the details and the performance of such a CFW testbed with an intelligent SDN-enabled orchestrator to create and manage QoS-aware end-to-end (e2e) network slices. Via this intelligent orchestrator we tailor, in real-time, the fixed-line QoS-aware E2E bandwidth slices to the specifics of a wireless application, taking into account the status of the entire network.

A real-time E2E connectivity testbed

The experimental testbed is schematically depicted in fig.1 and is employing no dedicated connectivity resources. Instead, a Nokia ISAM 7360 XGS-PON is used for the data-plane in Metro, a Nokia ISAM 7330 GPON system serves as a fixed access system and a MikroTik SXT2 WiFi AP serves as the wireless access technology.

The deployment of PONs in the fixed-line part ensures: a) the statistical use of the available bandwidth; b) the traffic flows are traversing an E2E path in a lossless way; c) there is a significant reduction in the number of transceivers needed in this segment - half compared to the case of dedicated point-to-point (P2P) connections. In the context of I4.0, we aim to remotely control a robotic entity. The robot exchanges data via the WiFi AP which is attached to ONU_1 (fig.1). The robot's operator is interfaced to the testbed via ONU_3 that sends traffic to Metro Node B via path (1) [3]. The GPON flows are aggregated in the L2 switch that forwards them to the transmitter (Tx) of the XGS-PON's ONU at point (2). Metro Nodes A and B

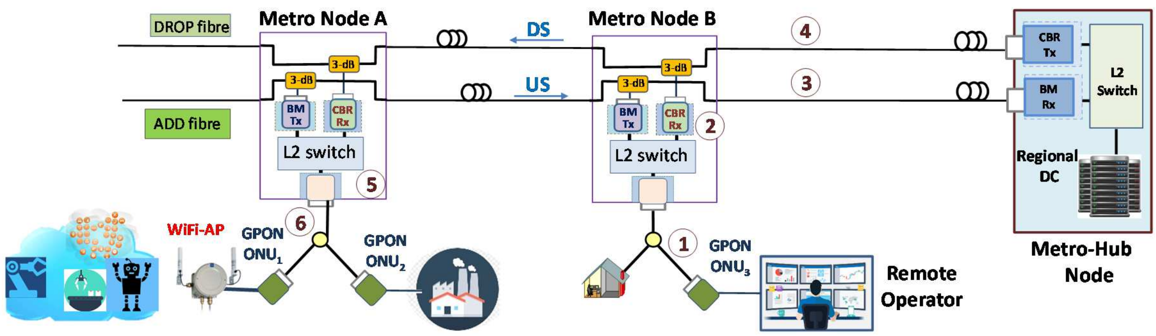

Fig. 1: The real-time e2e end connectivity testbed

978-1-6654-3868-1/21/$31.00 ©2021 IEEE

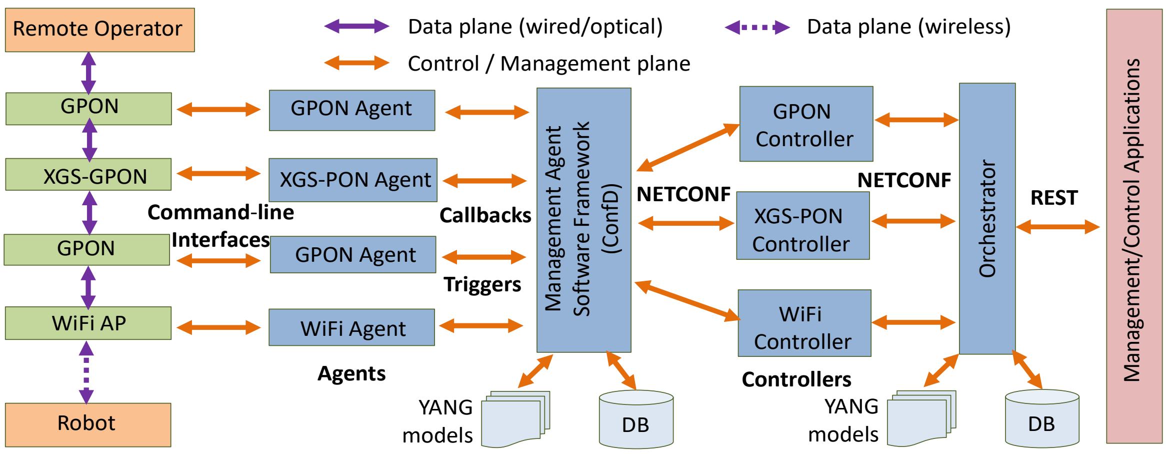

Fig. 2: The control-plane architecture for the testbed

operate burst-mode in the upstream (US) direction and they TDMA-share the capacity of the "ADD" link by means of a filterless architecture that is based on 3dB couplers arranged in a bus-topology, with one coupler per Metro node[4]. The L2 switch at the Metro-Hub node aggregates the US flows and directs them towards either a Regional Datacentre (DC) or to the DROP link. Metro Nodes A and B TDM-share the downstream (DS) capacity in the DROP link (4) that operates in constant-bit-rate (CBR) mode. Then, the flows following the symmetric paths (5) and (6) terminate to the ONU$_1$. Packets from the robot follow a symmetrical reverse path.

QoS-aware E2E slicing and latency

The purpose of an intelligent orchestrator for this multi-purpose testbed is: a) to provide E2E connectivity slices tailored to the per service QoS requirements; b) to maximize the utilization of the entire pool of available connectivity resources.

This is made feasible with PON technology, where the connectivity object is identified as a four-state entity using T-CONT containers[8], while in P2P CBR connections the state of a connection appears as a binary object[7] We demonstrate that the selection of the appropriate T-CONT is key for bounded latency performance during the service set-up phase, as well as, for the minimisation of the operational latency (after connection set-up is completed). The operational latency can be further improved using new burstyfication techniques[5],[6], however these techniques were not implemented here.

An SDN-enabled control-plane for the testbed

The intelligent orchestrator architecture in fig.2 is based on the standardized Netconf, OpenFlow and REST protocols to ensure interoperability. The configuration of the testbed elements is realised using a set of agents, as instructed by a set of controllers for the heterogeneous GPON, XGS-PON and WiFi AP systems. The agents configure the PONs via their CLIs, and the WiFi AP using the RouterOS[9] interface of the AP.

All agents are managed by the Management Agent Software Framework (MASF) based on the ConfD[10] platform. The MASF is responsible to maintain stability across the whole infrastructure configuration by retaining a set of call-backs and creating appropriate triggers on the agents. The MASF exposes the abstracted XGS-PON and GPON schemes to the controllers and acts as a Netconf Server enabling the configuration using the Netconf protocol. The adopted abstraction scheme[11] represents OLTs and ONUs as legacy Netconf L2 switches to which the GPON QoS queue models[12] are incorporated. This approach enables both the use of standard protocols and the abstraction of vendor specific details. In MASF, the appropriate YANG models for the GPON, XGS-PON and WiFi AP are derived and exposed to the controllers. The controllers communicate with the MASF using Netconf, while the testbed is E2E-managed by a higher-hierarchy-controller which has the role of orchestrator. The orchestrator instructs the controllers using Netconf, while it enables the automated configuration of the whole network through the northbound interface's REST API.

Experimental validation

The intelligent CFW overarching orchestrator of fig.2 is employed to create QoS-aware E2E slices on the experimental testbed of fig.1. To measure the latency for setting up a new slice, a 100 Mb/s TCP flow is generated between the WiFi-AP and the remote-control server (ONU$_3$) in fig.1 while the corresponding elements on the CFW testbed are concurrently set to the requested configuration. The corresponding e2e service (control-plane dependent) set-up latency is reported in Table I. The total latency is 4.4s on average from which 2.2s is for realising the network reconfiguration via the three agents and

2.2s for the realisation of the communication across the control layer (from the control application request down to the network elements and back). This response time is half that of[3] due to further improvements on the authentication procedures.

Table I: End-to-end slice set-up latency

Control plane latency		Orchestrated configuration of GPON, XGS-PON and WiFi		
		AVG	MIN	MAX
GPON Agent	Latency in GPON Agent for GPON reconfiguration (ms)	1344	942	1804
GPON Agent	Configuration preparation and processing (ms)	175	171	182
GPON Agent	Total latency in GPON Agent (ms)	**1519**	**1116**	**1986**
XGS-PON Agent	Latency in XGS-PON Agent for XGS-PON reconfiguration (ms)	598	570	624
XGS-PON Agent	Configuration preparation and processing (ms)	102	28	217
XGS-PON Agent	Total latency in XGS-PON Agent (ms)	**700**	**652**	**809**
WiFi	Total latency in WiFi Agent (ms)	**2**	**2**	**2**
Total	Controller-agent communication latency, MASF processing latency (ms)	2205	1953	2436
Total	Latency of whole reconfiguration cycle (app-controllers-agents-controllers-app)	**4426**	**3913**	**4782**

To study the impact of the QoS-aware slices on the E2E latency performance during operation, we launched two flows: Flow-I (flow of interest) between the wireless robot connected to the WiFi AP (ONU_1) and the Operator (ONU_3) to transport the high priority control messages of a robotic platform, as in fig.1, and Flow II between ONU_2 and a DataCentre located at the Metro-Hub node (emulating non critical traffic). Initially, only Flow II of 500Mb/s is active, while at t=3s, Flow I of 100Mb/s is activated for 10s. In scenario A, both flows are transported in the fixed-line segment using T-CONT-4 (best-effort). In fig.3a it is observed that the CFW platform fails to reach the requested performance for Flow I. Then, a 100 Mb/s QoS-aware slice was concurrently set

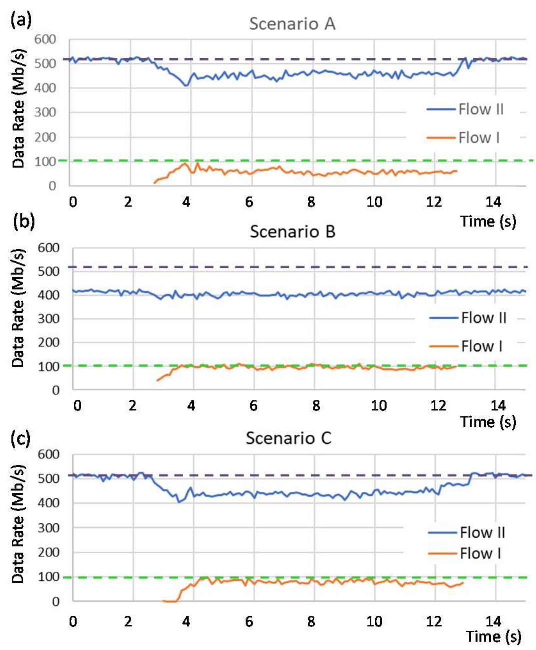

Fig. 3: Impact of the QoS-aware slices on latency

across the CFW testbed using the T-CONT-1 and T-CONT-2 options to serve Flow-I across the GPON/XGS-PON systems in Scenarios B and C, respectively. We observe in fig.3b that in Scenario B, the T-CONT-1 based E2E slice ensures the requested performance for Flow I at the expense of Flow II's bandwidth. In fig.3c (using T-CONT-2) the corresponding performance is slightly lower with a slower connection set-up time.

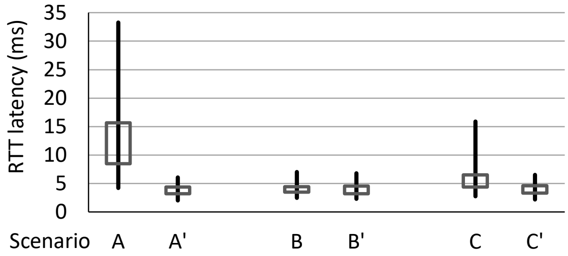

Fig. 4: Round-trip-time latency performance

Finally, we measured the round-trip-time (RTT) latency performance of the CFW testbed during operation using the ping tool. Two sets of values were obtained by having Flow II active or inactive (inactive denoted with accented letters for each scenario), while measuring Flow I. The results are shown in fig.4, a "box and whisker" plot with the whiskers showing the full range of measured values while the boxes show the range containing 50% of the measurements around the average.

Scenario A, where we employed T-CONT-4, is the one with the highest latency when Flow II is active. This value drops considerably when Flow II is inactive (A'). On the other hand, Scenarios B and C are showing a similar and substantially better performance compared to A even in the presence of Flow-II. Finally, we observe that the impact of using T-CONT-2 instead of T-CONT-1 is not negligible as the former gives slightly worse latency results on average and with higher variance. Using T-CONT-1, a RTT latency of ~ 4 ms is measured. To this end, this CFW testbed offers significantly more flexibility to allocate the necessary bandwidth using the appropriate T-CONT set, taking into account the overall QoS performance requests across the network.

Conclusions

We report a CFW Metro-Access testbed consisting of WiFi AP, GPON and XGS-PON heterogeneous systems. An overarching intelligent orchestrator is needed to provide the appropriate E2E QoS-aware bandwidth slices exploiting the T-CONT repertoire to match the WiFi AP's service status. We report a service set-up time ~4.4 seconds and a RTT latency of ~4 ms using T-CONT-1.

References

[1] A. Stavdas; “5G as a Catalyst for a Wider Technological Fusion that Enables the Fourth Industrial Revolution”, 7th Conference on Competitive Advantage in the Digital Economy (CADE), Warwick, 2021

[2] K. Letaief et al.; "The Roadmap to 6G: AI Empowered Wireless Networks", IEEE Communications Magazine, Vol.57 (8), pp.84-90, 2019

[3] E. Kosmatos, C. Matrakidis, D. Uzunidis, P. Kostopoulos, A. Stavdas, A. Lord; “An End-to-End Real-Time Connectivity Testbed Exploiting Standardized Dynamic Data-Plane PON Technologies”, ECOC 2020, Brussels, Belgium

[4] D.Uzunidis, C. Matrakidis, A. Stavdas, A.Lord;"DuFiNet:Architectural Considerations and Physical Layer Studies of an Agile and Cost-Effective Metropolitan Area Network", OSA/IEEE Journal of Lightwave Technology, Vol.37, issue 3, pp.808 - 814, 2019

[5] Th. Pfeiffer; “Considerations on transport latency in passive optical networks”, ECOC 2019, Ireland

[6] S. Bidkar, R. Bonk, Th. Pfeiffer; "Low-Latency TDM-PON for 5G Xhaul", ICTON 2020, Italy

[7] https://opennetworking.org/;TR-547-TAPI-v2.1.3-Reference-Implementation-Agreement

[8] Rec. ITU-T G.984.3 and Rec. ITU-T G.9807.1

[9] RouterOS, https://mikrotik.com/software

[10] ConfD platform, http://www.tail-f.com/management-agent/

[11] E. Kosmatos, C. Matrakidis, A. Stavdas, T. Orfanoudakis; "An SDN Architecture for PON Networks Enabling Unified Management using Abstractions"; ECOC 2018, Rome, Italy.

[12] "ITU-T PON YANG Modules", BroadBand Forum, WT-385, May 2017