QoS-Aware scheduling policies for Open Fronthaul transport networks

Fátima Khan[†], Luis Diez[†], Luis M. Contreras[¶], Óscar Gil[¶], Elena Serna[¶], David Gregoratti[‡], Ramón Agüero[†]
[†]*Universidad de Cantabria*, Santander, Spain. khanf@unican.es, {ldiez, ramon}@tlmat.unican.es
[¶]*Telefónica*, Madrid, Spain. {luismiguel.contrerasmurillo, oscar.gillucia, elena.sernasantiago}@telefonica.com
[‡]*Software Radio Systems*, Barcelona, Spain. david.gregoratti@srs.io

***Abstract*—We have recently seen a paradigm shift on the architecture for Radio Access Networks (RAN), being Open-RAN (O-RAN) architecture one of the most relevant examples for 5G and 6G networks. In this work, we combine theoretical models, based on the 5G New Radio standard, and cross-validation through real experimentation, using *srsRAN* software, to characterize Open Fronthaul traffic. Then, a realistic simulation framework based on *ns-3* is designed and implemented to assess the performance of various scheduling policies for intermediate nodes providing fronthaul connectivity, focusing on the integration of Quality-of-Service (QoS) strategies for serving fronthaul together with other services (in the form of 5G slices) on those nodes. By exploiting such methodology, we assess the benefits of prioritization schemes, highlighting the need of Hierarchical QoS solutions, particularly for congestion events. We carry out an extensive simulation campaign, and we show that using appropriate scheduling solutions yields relevant benefits to slices having tight delay requirements, compared to baseline solutions, reducing not only the average value, but also its variability.**

***Index Terms*—ORAN, Open Fronthaul, 5G, scheduling, QoS, backhaul**

I. Introduction

Together with the advancements in 5G and the forthcoming 6G technologies, Open Fronthaul Networks have recently loomed as a pioneering multi-vendor networking paradigm, initially conceived within the framework of the Centralized Radio Access Network (C-RAN) architecture. The Open Radio Access Networks (O-RAN) architecture [11], in particular split `7.2x`, separates legacy radio base stations into two different nodes: (1) Open-Radio Unit (O-RU) and (2) Open-Distributed Unit (O-DU). This innovative approach introduces strict requirements for the transport network (the so-called fronthaul) used to connect these nodes, which needs robust solutions to accommodate these complex demands.

In this sense, this new architectural solution poses novel challenges to the network planing. On the one hand, moving the O-DU farther away from the O-RU would enable a higher centralization degree, having more radio units being coordinated by a single node, which would bring improved resource management and lower costs in the radio segment. On the other hand, larger distances between the O-DU and O-RU make the fronthaul traffic to share networking capabilities with other traffic flows, such as those belonging to the backhaul of 5G technology, with heterogeneous performance requirements. Hence, it becomes necessary to include QoS-aware solutions able to manage the effects caused by imposing higher centralization degrees.

In this paper we tackle this challenge, by looking at the scheduling policies of the intermediate nodes within the fronthaul network. We start by identifying a precise traffic model, which starts from a theoretical analysis, to be later validated by real measurements. We then use a realistic network setup to assess the benefits brought by QoS-aware scheduling policies at intermediate nodes. To carry out such assessment, this paper outlines the design, implementation and validation of an appropriate setup to evaluate scheduling policies at the transport level in the fronthaul segment of next generation Radio Access Networks. We exploit the *ns-3* framework to mimic typical Hierarchical Quality of Service (HQoS) mechanisms used by vendors.

The main contributions of this paper can be summarized as follows:

- We employ a dual approach to characterize the fronthaul traffic, combining both theoretical models, based on the 5G New Radio (NR) standard, with real experimentation, using the *srsRAN* software in Time Division Duplexing (TDD) configuration.
- We setup a realistic topology, within the *ns-3* framework, to validate the appropriate functionality of different scheduling policies using as reference a network setup from Juniper [1] based on [13]. The topology definition includes the traffic marking at edge nodes, according to the O-RAN standard for open-fronthaul traffic.
- Then, we analyze the performance of the schemes proposed by vendors following O-RAN specifications [12] and those with modified HQoS configuration, and compare their effect over backhaul traffic.

The rest of this paper is structured as follows. Section II reviews the state-of-the-art, and identifies the gap this work targets. In Section III we depict the proposed methodology, and how it is implemented within the framework of the *ns-3* simulator. Section IV discusses the fronthaul traffic model that we use to assess the performance of the proposed techniques, which combines a theoretical approach with real measurements. The benefits brought by the QoS-aware scheduling policies are discussed in Section V, while Section VI concludes the paper, providing an outlook of our future work.

II. Related Work

There are not many previous works looking at QoS-aware scheduling policies at the fronthaul transport network using realistic traffic models. However, there are some relevant references which shed light on various aspects of such networks, in the context of 5G deployment, paying special attention to latency challenges and network performance optimization.

Martins et al. [10] discuss the challenges of deploying Reinforcement Learning (RL) policies in real-time scenarios, specifically in fronthaul networks. The authors propose using policy distillation to extract simpler Decision Tree models to reduce inference time without hindering the performance. In the same line, the authors of [15] discuss a Traffic Pattern Adaptive Mechanism designed to bound packet delay and its variability in 5G Fronthaul networks. The proposed mechanism, Time-Window with Timeout (TWT), aims to balance packet delay and jitter while ensuring efficient traffic aggregation and management.

Zhang et al. [16] delve into the optimization of routing and packet scheduling in the context of the 5G Open Radio Access Network (O-RAN) Fronthaul architecture. Particularly, they focus on enhancing communication efficiency between O-DU pools and O-RUs to meet the diverse service requirements of 5G networks. The proposed solution, based on Dynamic Programming, is able to select optimal routes and packet sizes for transmitting traffic across multiple O-DU pools.

Pérez et al. investigate in [14] the performance of the Common Public Radio Interface (CPRI) protocol over ring-star topologies used by mobile operators. The study comes up with a number of key findings: benefits of bidirectional ring topologies in reducing average queuing delays, queuing theory models to analyze delay metrics, and packetization strategies in minimizing worst-case aggregated queuing delays. Besides, the work in [8] proposes alternatives for latency planning for both fronthaul and 5G slicing.

These works collectively underscore the significance of optimizing MAC and PHY layers, as well as the underlying network topology configurations, to minimize latency in fronthaul networks. As can be observed, there exists a gap in the evaluation of the configuration of intermediate nodes and HQoS in transport segment, which integrates both open-fronthaul and backhaul traffic. There are only a few works that touch this aspect.

Budhdev et al. design and implement in [4] a Fronthaul Slicing Architecture (FSA), which uses a wireless scheduler to identify the slice/user for each fronthaul packet. An evaluation over a real testbed evinces FSA's capability to handle up to 80 Gbps of fronthaul traffic with low-latency routing, while yielding a significant reduction in flow completion times.

Closer to our work, Balogh et al. compare in [3] the performance of Fair Queuing (FQ) Queue Scheduling Disciplines (QSD) with Weight Round Robin (WRR) in a bottleneck topology. The study suggests that FQ algorithms show worse computational performance compared to Round Robin (RR) algorithms. However, the paper hints that an accurate bandwidth allocation in WRR could lead to better performances.

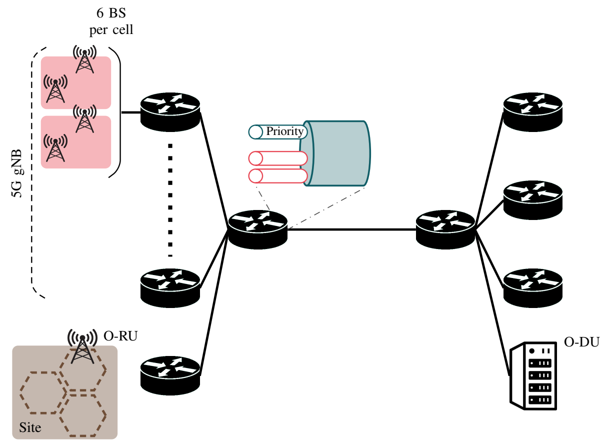

Fig. 1: Evaluation platform setup in *ns-3*, with a *DiffServ* architecture and fronthaul and backhaul service convergence.

This work diverges from existing research by focusing on the assessment of performance of the HQoS mechanisms proposed by different vendors, specifically for the fronthaul network. Unlike previous efforts, which mostly concentrated on enhancing fronthaul at the MAC/PHY layers, this study assesses the behavior of various Quality of Service (QoS) approaches in practical scenarios, where technologies converge. It also analyzes the impact of prioritizing certain technologies over backhaul services, thus pointing out the circumstances where HQoS would be required to ensure optimal performance.

III. Evaluation Platform

We have crafted and implemented a platform within the *ns-3* framework that allows us to study the performance of scheduling policies at intermediate nodes within the fronthaul network. In a nutshell, our setup exploits the traffic-control capabilities to mimic the behavior of a router. Additionally, we have also developed some utilities to ease the systematic deployment of different scenarios, enabling the configuration of various traffic patterns, belonging to either backhaul services or the fronthaul segment (such as 5G gNB, O-RU), the scheduling policies, as well as the Differentiated Services Code Point (DSCP) mapping.

Figure 1 depicts the evaluation scenario implemented in *ns-3* as described in [1]. As can be seen, it aggregates fronthaul and backhaul traffic, which is managed according to the policies of the routers in shared links. In addition, access routers provide configurable DSCP marking to traffic flows.

In order to guarantee an accurate model for both traffic volume and packet size distribution, we developed an *ofh-application* derived from the original *OnOff* application. This application mimics traffic generation of both the Control- and User-Plane, allowing the transmission of two flows (one per plane) each of them with their own characteristics, as detailed in Section IV. We use IP nodes in our experiments to implement the fronthaul routers, and we can thus use the *traffic control* tool within *ns-3*, which provides a functionality akin to the Linux Traffic Control. Positioned between the *NetDevices* and any network protocol atop, it takes responsibility for packet processing, and it is able to execute various actions, such as scheduling and dropping packets. To emulate edge router behavior, we have extended the default functionalities by enabling per port DSCP packet marking for classification purposes when necessary.

In order to mimic the schemes proposed in [6], we have implemented a two-level HQoS scheduler, which considers priority and non-priority queues with different QSD. In addition, for the non-priority queues, a WRR scheduling policy has been included. All the QSD disciplines can be configured to perform mapping between underlying queues and DSCP marking in the packets. Altogether, the operation of the scheduler is as follows. When the prioritization queue has any packets pending to be served, the scheduler will always prioritize them for transmission. On the other hand, when the prioritization queue is empty, the WRR is employed, and the other queues are served based on the maximum quantum of packets allowed for transmission. The quantum of each queue is reset to its initial value when it reaches zero.

It is worth noting that the *ns-3* development used in this work has been made available to the research community at a public repository[1]. In addition, to guarantee the reproducibility of the results discussed in the paper, all scenario building and configuration utilities are also included.

IV. Open Fronthaul Traffic Modeling

Understanding traffic patterns of Open Fronthaul networks, as well as the load distribution between User-Plane and Control-Plane, is pivotal for assessing the performance of HQoS strategies and to evaluate whether they can improve the overall QoS. Figure 2 depicts the O-RAN protocol architecture from which only the Control and User Planes are integrated in the fronthaul. Over such stack, the precise characterization of each traffic flow serves as a fundamental input to tackle the system setup. In this work we assume Time Division Duplex (TDD) mode where each slot comprises 14 symbols with the normal Cyclic Prefix (CP), where two symbols are assigned to the Control-Plane, and the rest of them to the User-Plane [9].

To establish the packet size of the User Plane, we adopt eCPRI over Ethernet (without IP/UDP), using fronthaul block compression techniques, like Block Floating Point (BFP) [7]. The Physical Resource Block (PRB_{size}) is defined as the payload data width (9 bits by applying BFP9 compression) multiplied by the number of subcarriers, accounting for both the I and Q components, with an additional 8 bits for overhead:

$$PRB_{size} = Bytes(data_{width} \cdot N_{subcarriers} \cdot 2 + 8) \tag{1}$$

Then, the number of available Resource Blocks (RBs) within the given bandwidth (BW) and Subcarrier Spacing

[1]https://github.com/tlmat-unican/ns3-hqos-open-fronthaul

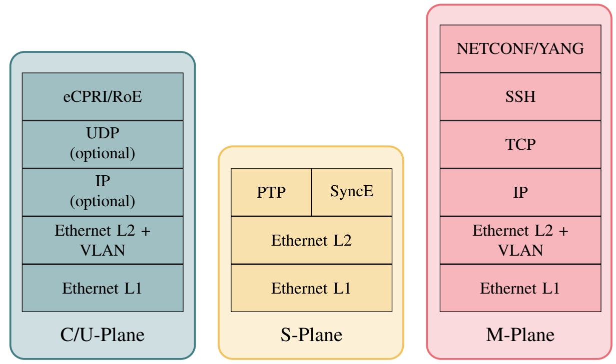

Fig. 2: O-RAN protocol architecture.

(SCS) considering guard bands defined in [2] can be calculated as:

$$N_{RB} = \left\lfloor \frac{BW - 2 \cdot BW_{guard}(SCS, BW)}{SCS \cdot N_{subcarriers}} \right\rfloor \tag{2}$$

Based on it, we can calculate the total number of bits allocated to user data, iq_{data}, considering the Physical Resource Block (PRB) size, as follows:

$$iq_{data} = N_{RB} \cdot PRB_{size} \tag{3}$$

Besides, the packet size ($Packet_{size}$) encompasses both user data and packet header overhead:

$$Packet_{size} = iq_{data} + Packet_{header} \tag{4}$$

Finally, the data rate in each configuration is estimated by considering the packet size, the number of symbols per slot (i.e., 12), the modulation order (2^{μ}), and converting it to bits per second:

$$Rate = Packet_{size} \cdot 12 \cdot 2^{\mu} \cdot 8 \tag{5}$$

On the other hand, Control-Plane packets are considered to be of 64 bytes and are sent twice per slot. The above expressions provide a quantitative model for evaluating traffic characteristics and allows us to estimate data rates, which is essential to assess the behavior of QoS-aware strategies in fronthaul networks. The corresponding theoretical values are summarized in Table I.

To validate the values obtained, we use the *srsRAN*[2] software which is a modular software fully compliant with the O-RAN architecture. In particular, we instantiate an *srsRAN* O-DU sending Open Fronthaul Traffic to a the loopback interface (i.e. the O-RU is not present). The traffic generated is then captured and analyzed identifying the plane corresponding to each packet. For the experiment, we use two FR1 bands, 20 MHz and 100 MHz, with a SCS of 30 KHz, and two TDD configurations: *6d3u* and *7d2u* in a Single-output Single-input (SISO) mode. Figure 3 depicts the average rate yielded by the experimental setup and the theoretical approach. As can be seen, it evinces the validity of the proposed model, with an error less than 10% between theoretical and experimental values. It is worth mentioning that the theoretical values have

[2]https://www.srsran.com/

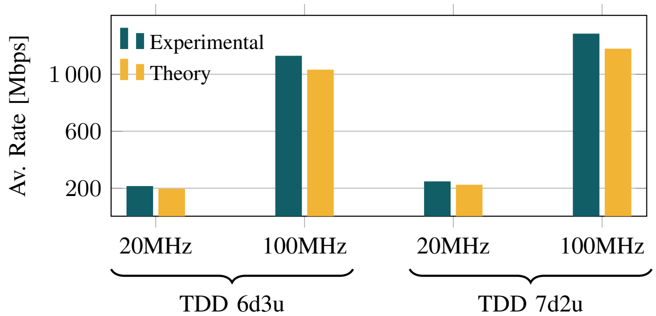

Fig. 3: Traffic analysis and experimentation with srsRAN Project Software (DL Mode) using SCS of 30 kHz across various Bandwidths and TDD configurations in a 1 SISO carrier mode.

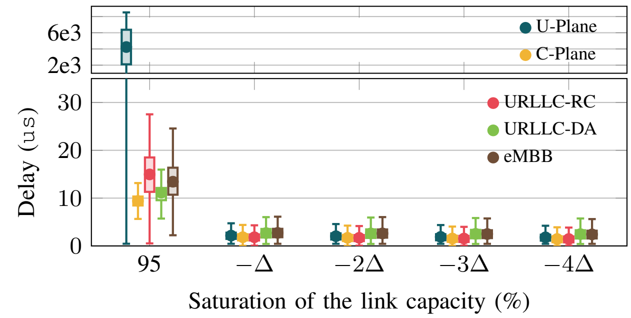

Fig. 4: E2E delay without HQoS scheduling polices ($\Delta = 1\%$)

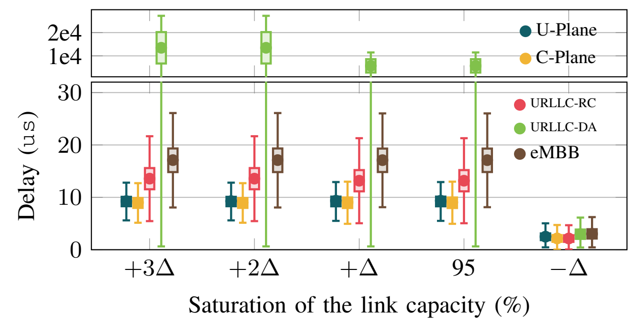

Fig. 5: E2E delay with CU-Plane traffic prioritization and WRR with initial weights ($\Delta = 0.05\%$).

been scaled, according to the particular use of the downlink and special slots (DL + 1).

V. Validation and Use Case Application

According to the setup depicted in Figure 1, we analyze the impact of introducing HQoS scheduling policies at intermediate nodes, as well as how this might affect other services. Our scenario includes a site with 3 Open-Radio Unit (O-RU), with bandwidths of 20 MHz, and SCS of 30 kHz. Table II summarizes the parameters used in our network setup. Backhaul services are integrated as 5G gNB nodes, which require 12 BSs to cover an industrial area [5]. Three ingress nodes are considered, with two of them offering two different Ultra-Reliable and Low Latency Communications (URLLC) services (Remote Control and Discrete Automation), while the third ones aggregates traffic of enhanced Mobile Broadband (eMBB) services. All traffic flows are Constant Bit Rate (CBR) sources and it is assumed that buffers are not limited.

Initially, we study the performance of combining multiple flows from both backhaul and fronthaul, and we measure the End-to-end (E2E) delay between O-DU and O-RU (fronthaul), as well as the delay affecting backhaul traffic, while sweeping channel saturation (Δ). This analysis aims to understand the impact of the fronthaul traffic over background flows, since its latency requirements are very strict. As can be seen in Figure 4, when the aggregated traffic remains below 95% of the channel capacity, the impact of using QoS-aware scheduling policies is almost negligible, and the delays observed for the various flows are almost alike, exhibiting low values. However, once this threshold is exceeded (as it might happen in punctual congestion events), U-Plane traffic flow suffers from a relevant delay increase, which goes clearly beyond the requirements imposed on the fronthaul, being 75 μs for the one-way delay [1]. We can therefore establish this 95% as the threshold above which the use of QoS-aware schedulers is needed to keep delays at a reasonable level.

Applying the QoS model with one priority queue, we aim to meet latency requirements in the fronthaul. Consequently, enhanced Common Public Radio Interface (eCPRI) control and user plane traffic are mapped to the aforementioned priority queue. Non-priority queues are served with a WRR policy. In this context, considering the diverse traffic characteristics and packet sizes among the flows, weights (which correspond to the number of packets allocated) are assigned based on the arrival rates λ_i of the flows, taking into account both the aggregated rate and the packet size, $w_i = \frac{\lambda_i}{\sum_{i=0}^{N} \lambda_i}$. Notably, URLLC packets having smaller sizes (for instance, remote control services) have more stringent delay requirements, and the proposed weighting scheme takes this aspect into consideration by increasing its relative weight.

Figure 5 illustrates the E2E delay of each flow when weights are established according to the aforementioned strategy. Channel capacity has been swept around the saturation point that was previously identified to examine the impact of the scheduling policies. As can be seen in the figure, in all cases the fronthaul traffic meets the delay requirements. Then, in Figure 6, to emphasize that the second group of traffic from URLLC exhibits higher jitter and average delay compared to that from eMBB, we adjust the weights of the WRR scheduler, shifting some of the allocated resources (in this case 1% of the overall bandwidth) from eMBB traffic to the URLLC flows, by slightly adjusting the corresponding weights. As can be seen, this adjustment affects the eMBB traffic, while the delay strongly decreases for the URLLC traffic.

We now analyze the impact of moving the U-Plane traffic to the WRR queues to better balance the effect on the backhaul traffic. Figure 7 illustrates the behavior when assigning weights just considering all arrival rates. The outcome is similar to the one observed in the previous configuration,

TABLE I: Packet Size including VLAN Tag (36-Byte) and U-Plane data rate per antenna port.

		BW (MHz)												
	SCS (KHz)	5	10	15	20	25	30	40	50	60	70	80	90	100
Packet size (bytes)	15	736	1492	2248	3004	3760	4516	6084	7596	-	-	-	-	-
	30	344	708	1100	1464	1856	2220	3004	3760	4572	5328	6112	6896	7680
	60	-	344	540	708	904	1100	1464	1856	2248	2640	3032	3424	3816
U-Plane data rate (Mbps)	15	70.66	143.23	215.81	288.38	360.96	433.53	584.06	729.21	-	-	-	-	-
	30	66.0	135.9	211.2	281.1	356.3	426.2	576.8	721.9	877.8	1023	1173.5	1324	1474.6
	60	-	99.1	155.5	203.9	260.3	316.8	421.6	534.5	647.4	760.3	873.2	986.1	1099

TABLE II: Parameters used in the scenario setup.

Parameters SITE O-RU	
BW	20 MHz
SCS	30 kHz (n = 1)
3 cells	16 · 16 MIMO
U-Plane (per antenna port)	281.088 Mbps
C-Plane (per antenna port)	2.048 Mbps
Packet Size	1492 B (+UDP/IP)
5G gNB - Industrial Area from Table 2,3,4 [5]	
Base Station (BS)	12
URLLC - Remote Control	750 Mbps (160B)
URLLC - Discrete Automatation	833.33 Mbps (1358B)
eMBB - Urban Area	1666.67 Mbps (1450B)
Topology	
Access Link	100 Gbps
Bottleneck	≈ 55 Gbps

TABLE III: Standard Deviation of the E2E delay (μs) for 95.15% of channel capacity.

	CU-P PRIO	CU-P PRIO∗	C-P PRIO	C-P PRIO∗
U-Plane	1.48	1.58	3.31	3.15
C-Plane	1.43	1.53	1.55	1.76
URLLC-MC	2.79	3.08	3.08	3.21
URLLC-DA	7800.23	4.99	7800.50	4.75
eMBB	3.25	3893.04	3.38	3892.99

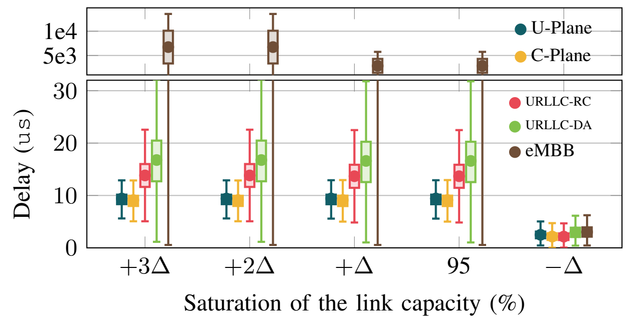

Fig. 6: E2E delay with CU-Plane traffic prioritization and WRR with modified weights ($\Delta = 0.05\%$).

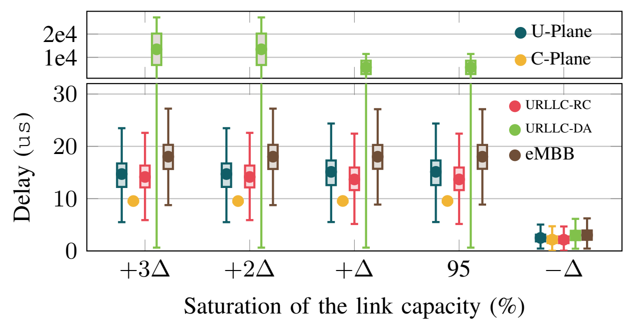

Fig. 7: E2E delay with C-Plane traffic prioritization and WRR with initial weights ($\Delta = 0.05\%$).

with the User-Plane traffic having a slight increase in its delay (average value and variability), albeit still conforming to the corresponding requirements of 75 μs. Then, we used the modified policy, shifting some resources from eMBB to URLLC. The results, as shown in Figure 8, evidence that even with small adjustments of the scheduling policies, we can strongly influence the corresponding delays and their jitter.

Finally, Table III provides a summary of the standard deviation of delay observed for all traffic flows and scheduling policies (indicated with ∗ those with adjusted weights), when the bottleneck link is at 95.15% of its capacity. As can be seen, the fronthaul traffic exhibits a very predictable behavior for all considered schedulers, while tuning the weights of the WRR scheme strongly affects the performance shown by the other flows.

VI. Conclusions

In this work, we have designed and implemented a methodology, exploiting the *ns-3* framework, to evaluate QoS-aware scheduling policies to be used by intermediate nodes in fronthaul networks. Using such platform, we have evaluated the performance of various schemes from the literature, studying not only the average end-to-end delay, but also its variability. We have seen that using scheduling techniques becomes of utter relevance to ensure that the stringent requirements of fronthaul traffic are met, even for punctual events of redundant link failures and temporary congestion. We have seen that this can be achieved even without using a priorization queue. The results also evince that the proposed solution is sensitive to its configuration, since small weight adjustments may have a strong impact on the results for certain traffic flows.

In our future work, we plan to extend the characterization by expanding the topology to explore the physical distance limitations between O-DU and O-RU. Moreover, we will also increase the number of O-RU sites, to assess the impact of traffic aggregation, precisely modeling the compression

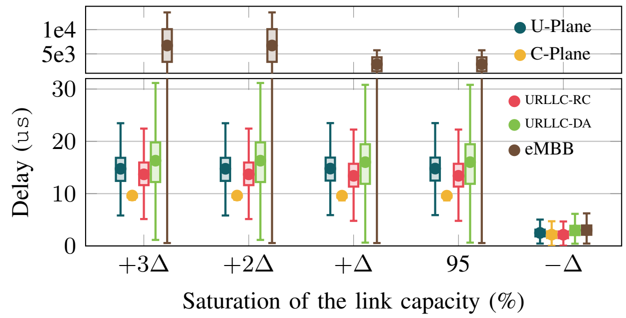

Fig. 8: E2E delay with C-Plane traffic prioritization and WRR with modified weights ($\Delta = 0.05\%$).

techniques that are used in real networks. Furthermore, we plan to introduce interleaving at the scheduler, to reduce the corresponding jitter, to avoid the bursty effect that characterizes the proposed schedulers. We will exploit the methodology presented in this paper to carry out this work, but we will also validate some of these configurations by means of experimentation over a real testbed, at Telefonica Technology and Automation lab premises.

Acknowledgment

This work has been supported by the Spanish Ministry of Economic Affairs and Digital Transformation and the European Union – NextGenerationEU with the projects: 6GBLUR/JOINT (TSI-063000-2021-57) and SITED: Semantically-enabled Interoperable Trustworthy Enriched Data-spaces (PID2021-125725OB-I00). We have also received funds from the regional Cantabria Government, Fondo Europeo de Desarrollo Regional (FEDER): "Grants for research projects with high industrial potential of technological agents of excellence for industrial competitiveness (TCNIC)" program (2023/TCN/002).

References

[1] "Juniper Validated Design (JVD) 5G Fronthaul Class of Service ACX7100, ACX7509".

[2] 3GPP. *TS 138 101-1 - V16.5.0 - 5G; NR; User Equipment (UE) radio transmission and reception; Part 1: Range 1 Standalone (3GPP TS 38.101-1 version 16.5.0 Release 16)*. TS. 2020.

[3] Tomáš Balogh and Martin Medvecký. "Performance evaluation of WFQ, WF2Q+ and WRR queue scheduling algorithms". In: *2011 34th International Conference on Telecommunications and Signal Processing (TSP)*. 2011, pp. 136–140. DOI: 10.1109/TSP.2011.6043756.

[4] Nishant Budhdev et al. "FSA: fronthaul slicing architecture for 5G using dataplane programmable switches". In: *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*. MobiCom '21. New Orleans, Louisiana: Association for Computing Machinery, 2021, pp. 723–735. ISBN: 9781450383424. DOI: 10.1145/3447993.3483247.

[5] Luca Cominardi et al. "Understanding QoS Applicability in 5G Transport Networks". In: *2018 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*. 2018, pp. 1–5. DOI: 10.1109/BMSB.2018.8436847.

[6] Luis M. Contreras, Ivan Bykov, and Krzysztof Grzegorz Szarkowicz. *5QI to DiffServ DSCP Mapping Example for Enforcement of 5G End-to-End Network Slice QoS*. Internet-Draft draft-cbs-teas-5qi-to-dscp-mapping-00. Work in Progress. Internet Engineering Task Force, Mar. 2024. 15 pp.

[7] Sandra Lagén et al. "Modulation Compression in Next Generation RAN: Air Interface and Fronthaul Trade-offs". In: *IEEE Communications Magazine* 59.1 (2021), pp. 89–95. DOI: 10.1109/MCOM.001.2000453.

[8] David Larrabeiti et al. "Toward end-to-end latency management of 5G network slicing and fronthaul traffic (Invited paper)". In: *Optical Fiber Technology* 76 (2023), p. 103220. ISSN: 1068-5200. DOI: https://doi.org/10.1016/j.yofte.2022.103220.

[9] Ana Larrañaga et al. "An open-source implementation and validation of 5G NR configured grant for URLLC in ns-3 5G LENA: A scheduling case study in industry 4.0 scenarios". In: *Journal of Network and Computer Applications* 215 (2023), p. 103638. ISSN: 1084-8045. DOI: https://doi.org/10.1016/j.jnca.2023.103638.

[10] Jean P. Martins et al. "Policy Distillation for Real-Time Inference in Fronthaul Congestion Control". In: *IEEE Access* 9 (2021), pp. 154471–154483. ISSN: 21693536. DOI: 10.1109/ACCESS.2021.3129132.

[11] ORAN. *O-RAN Working Group 1 (Use Cases and Overall Architecture WG). O-RAN Architecture Description (V11.00)*. TS. 2024.

[12] ORAN. *O-RAN Working Group 4 (Open Fronthaul Interfaces WG). Control, User and Synchronization Plane Specification (V14.00)*. TS. 2024.

[13] ORAN. *O-RAN Working Group 9 (Open X-haul Transport WG). O-RAN Xhaul Packet Switched Architectures and Solutions (V7.00)*. TS. 2024.

[14] Gabriel Otero Pérez, José Alberto Hernández, and David Larrabeiti López. "Delay analysis of fronthaul traffic in 5G transport networks". In: *2017 IEEE 17th International Conference on Ubiquitous Wireless Broadband (ICUWB)*. 2017, pp. 1–5. DOI: 10.1109/ICUWB.2017.8250956.

[15] Federico Tonini et al. "A Traffic Pattern Adaptive Mechanism to Bound Packet Delay and Delay Variation in 5G Fronthaul". In: *2019 European Conference on Networks and Communications (EuCNC)*. 2019, pp. 416–420. DOI: 10.1109/EuCNC.2019.8802050.

[16] Cheng Cheng Zhang, Kim Khoa Nguyen, and Mohamed Cheriet. "Routing and Packet Scheduling For Virtualized Disaggregate Functions in 5G O-RAN Fronthaul". In: *2021 IEEE Global Communications Conference (GLOBECOM)*. 2021, pp. 1–6. DOI: 10.1109/GLOBECOM46510.2021.9685802.