The 5G-Crosshaul Packet Forwarding Element pipeline: measurements and analysis

Nuria Molner*[†], Sergio González*, Thomas Deiß[‡], Antonio de la Oliva*
*Universidad Carlos III de Madrid,[†]IMDEA Networks Institute,[‡]Nokia

Abstract—This paper is focused on the 5G-Crosshaul Packet Forwarding Element (XPFE), which is the packet forwarding element of the 5G-Crosshaul network architecture. The XPFE integrates multiple technologies which allow to transport traffic from multiple tenants and of different nature over the same infrastructure. Hence, the paper is focused on the performance of this essential element of our network, providing some end-to-end delay measurements and analyzing the behavior of this delay. Likewise, we have fitted the appropriate probability distribution for these measurements, which allows us to infer some confidence intervals for the prediction of the maximum number of hops supported by delay-sensitive fronthaul traffic before being processed in a 5G-Crosshaul Processing Unit.

Index Terms—5G, 5G-Crosshaul, fronthaul, backhaul, Multi-tenancy, SDN, NFV, delay, latency.

I. Introduction and Motivation

Mobile data traffic is predicted to increase seven-fold between 2016 and 2021 [1]. Hence, a new generation of air interfaces is required to satisfy the increasing user demand. This new generation is the Fifth Generation of network architectures and will be characterized by an increment of the available bandwidth and a higher capillarity, increasing the density of the network.

Nowadays, mobile networks use dedicated paths for the different types of traffic which comes from different tenants. It results in an inefficient network utilization because some of the resources are unused for some time, while others are overused and cannot assure reliable quality of service (QoS) to the users. Thus, to be able to satisfy the increasing traffic demand giving some guarantees to the end users, some new solutions which use the resources more efficiently the resources and save costs of deploying more networks have appeared.

A proposed solution is to base the future networks on the software-defined network (SDN) principles and on the architecture of Open Networking Foundation (ONF) [2]. These features allow to deal with multiple technologies, integrating them and applying Network Function Virtualization (NFV) concepts and the ETSI (European Telecommunications Standards Institute) NFV MANO (Management and Orchestration) architecture. The SDN principle allows to separate data and control planes, while NFV allows infrastructure and function virtualization. The Virtual Network Functions (VNFs) enable a dynamic and optimized instantiation and connection, using the resources in an efficient way. Thus, the networks will be virtualized in as many elements as possible, giving the opportunity of using the same physical resources by different types of traffic and tenants. Some policies will have to be established over the physical resources to accommodate this diverse nature of traffic, but the resources will be used in a more efficient way.

Hence, because the nature of networks will change, new network elements have to be developed to support these new requirements and transport and allocate the new multi-technology traffic. Here the European H2020 5G-Crosshaul project plays an important role providing an architecture based on SDN/NFV for 5G transport networks which integrates fronthaul and backhaul segments. The integration of both network segments is done in the 5G-Crosshaul Packet Forwarding Elements (XPFE), which are software based switches that integrates features to adapt the legacy technologies with the new ones in the same network.

In the case of 5G-Crosshaul project, the goal is the integration of fronthaul and backhaul traffic. This adaptation requires the packetization of the fronthaul traffic since it currently is using a continuous constant bitrate transmission. This traffic is critical due to its strict requirements in terms of delay and jitter. There are different types of fronthaul traffic with diverse requirements, where the most delay-sensitive is the Common Public Radio Interface (CPRI) [3]. This packetization can be based on standards such as IEEE 1914.1, which focuses on the packetization of the CPRI, and the Time Sensitive Networking (TSN) Task Group of IEEE 802.1, which is developing new extensions to support the forwarding of Ethernet traffic with delay and jitter guarantees, including mechanisms such as frame preemption, expedited traffic forwarding, and jitter reduction techniques.

The rest of the paper is structured as follows: Section II explains multi-tenancy in the context of 5G-Crosshaul in Section II-A and gives an overview of the 5G-Crosshaul architecture, focusing on its main elements such as 5G-Crosshaul Common Frame (XCF) in II-B, 5G-Crosshaul Forwarding Element (XFE) in II-C emphasizing the 5G-Crosshaul Packet Forwarding Element (XPFE). Section III focuses on the XPFE pipeline providing a description of it in III-A, explaining its integration with multi-tenancy in III-B, providing some performance measurements in IV-A and an analysis of these measurements in IV-B. Finally, Section V concludes the document.

II. Architecture

The premise of the 5G-Crosshaul project architecture is a unified fronthaul/backhaul transport network for 5G allowing

the application of the multi-tenancy concept. Further details of this architecture can be found in [4].

A. Multi-tenancy

The approach to reduce significantly the CAPEX and OPEX and optimize the energy consumption is to share the physical resources among the different tenants, i.e., to support multiple tenants over the same physical infrastructure. To be able to allocate traffic from different tenants over the same resources it is necessary to make a distinction between them and virtualize the physical elements in a way that from the same physical resource, several virtual resources are created to be used by the different tenants without mixing them.

Hence, network slicing is a clear evolution in networking making the 5G system more flexible, sharable and cost-effective. Thus, by creating slices we can assign each of them for different purposes.

A key enabler of this multi-layer architecture is the 5G-Crosshaul Controller Infrastructure (XCI), it follows the ETSI NFV principles and is formed by the NFV Orchestrator (NFVO), the VNF Managers (VNFMs) and the Virtualized Infrastructure Manager (VIM); it also offers a Northbound Interface (NBI) and a Southbound Interface (SBI) for further interaction with the different elements of the network, for more details see [4].

It is worth mentioning than in a 5G-Crosshaul domain a tenant refers to the entity that owns and drives the instantiation of one or more network services. We differentiate 3 types of tenants inside the 5G-Crosshaul architecture:

- Over-the-Top (OTT) tenant: a service provider that does not require control over the infrastructure to connect distributed points in the 5G-Crosshaul domain.
- Mobile Virtual Network Operator (MVNO): a service provider who owns a virtual slice of the infrastructure. It has limited control and operation over the slice and can manage and optimize the virtual resources and also share the virtual slice with other tenants.
- Mobile Network Operator (MNO): a service provider who owns the 5G-Crosshaul physical infrastructure and offers the services to OTTs and MVNOs.

Due to the multi-tenancy support a MNO can share in a hierarchical manner its physical resources instantiating multiple 5G-Crosshaul MANO recursively. Each tenant has to agree on the service level with the MNO, and the MNO has to manage the resources to meet these agreements.

This multi-layer sharing is done through the API offered by the Multi-Tenancy Application (MTA) located on top of the XCI, for further details see [5]. In addition, Section III-B explains the methods used for the tenants' traffic differentiation.

Fig. 1 shows an example of recursive and hierarchical instantiation of the XCI over a physical infrastructure owned by the MNO tenant#0, the role of MVNO is performed by tenant#1 and tenant#2 controlling slices of the infrastructure layer, finally, tenant#3 and tenant#4 are OTT service providers without control over the physical infrastructure.
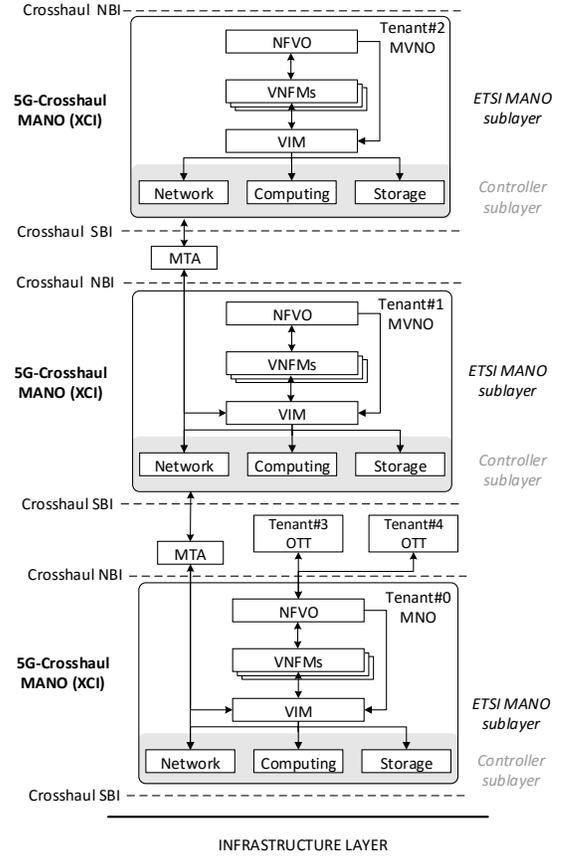

Fig. 1: Recursive MANO instantiation for multi-tenancy

In addition, to support multiple tenants in a network it is necessary to have network elements that support policies and agreements which are necessary to make the differentiation between the traffic coming from each tenant without mixing them. For this reason and the fact that the network also has to support different kinds of traffic, it is mandatory to create a common frame which supports all these requirements. This common frame is explained in II-B.

B. 5G-Crosshaul Common Frame (XCF)

To reach the integration of fronthaul and backhaul segments and support multiple tenants over the same infrastructure, a common frame format is necessary to transport these diverse types of traffic through the same infrastructure.

The common frame (XCF) used in the 5G-Crosshaul is based on Ethernet, including features of MAC-in-MAC [6] (or Provider Backbone Bridged Network (PBB)), such as tagging the packets with labels to differentiate the different tenants. These PBB frames allow the integration of multiple tenants to the XCF separated through the outer MAC header.

In addition, it allows the adaptation of the technologies of fronthaul (e.g., IEEE 1914.1 and 802.1 TSN) and backhaul traffic to be supported by this XCF. All these features are allowed thanks to the MAC-in-MAC mechanism showed in Fig. 2.
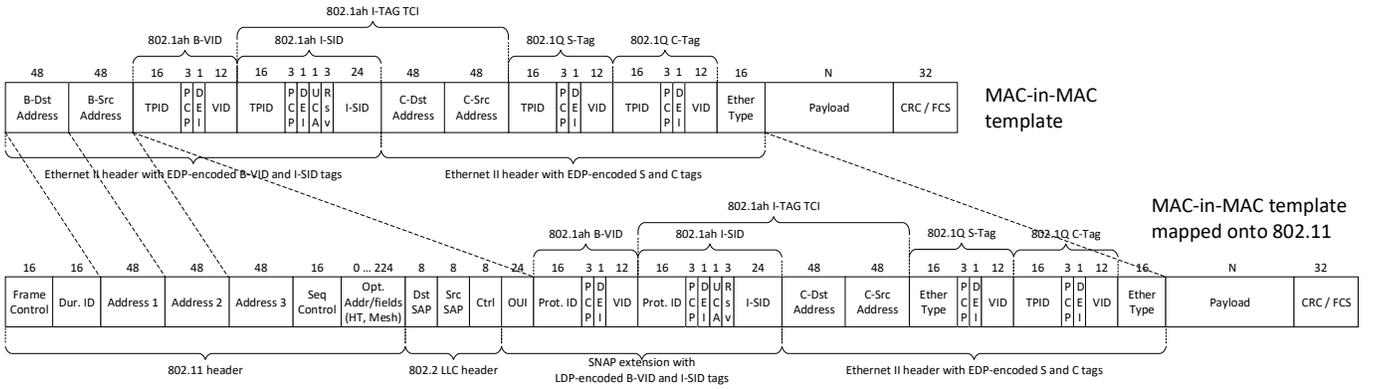
802.1ah I-TAG TCI spanning **802.1ah B-VID** and **802.1ah I-SID** | **802.1Q S-Tag** | **802.1Q C-Tag**

48	48	16	3 1 12	16	3 1 1 3 24	48	48	16	3 1 12	16	3 1 12	16	N	32	
B-Dst Address	B-Src Address	TPID	PCP/DEI/VID	TPID	PCP/DEI/UCA/Rsv/I-SID	C-Dst Address	C-Src Address	TPID	PCP/DEI/VID	TPID	PCP/DEI/VID	Ether Type	Payload	CRC/FCS	MAC-in-MAC template

Ethernet II header with EDP-encoded B-VID and I-SID tags   Ethernet II header with EDP-encoded S and C tags

MAC-in-MAC template mapped onto 802.11

802.1ah I-TAG TCI spanning **802.1ah B-VID** and **802.1ah I-SID** | **802.1Q S-Tag** | **802.1Q C-Tag**

16	16	48	48	48	16	0 … 224	8	8	8	24	16	3 1 12	16	3 1 1 3 24	48	48	16	3 1 12	16	3 1 12	16	N	32
Frame Control	Dur. ID	Address 1	Address 2	Address 3	Seq Control	Opt. Addr/fields (HT, Mesh)	Dst SAP	Src SAP	Ctrl	OUI	Prot. ID	PCP/DEI/VID	Prot. ID	PCP/DEI/UCA/Rsv/I-SID	C-Dst Address	C-Src Address	Ether Type	PCP/DEI/VID	TPID	PCP/DEI/VID	Ether Type	Payload	CRC/FCS

802.11 header   802.2 LLC header   SNAP extension with LDP-encoded B-VID and I-SID tags   Ethernet II header with EDP-encoded S and C tags

Fig. 2: 5G-Crosshaul Common Frame mapping

C. 5G-Crosshaul Forwarding Element (XFE)

The 5G-Crosshaul Forwarding Elements (XFEs) are the switching elements of the 5G-Crosshaul network. They integrate different link technologies, such as mmWave, Ethernet or Optical Fibre. XFEs are controlled by the 5G-Crosshaul Control Infrastructure (XCI), the control plane of our achitecture, that has a detailed view of the network, traffic and resources, and decides the route these XCFs have to follow in the network from its source to its destination.

XFEs are multi-layer switches, composed of a circuit-switching (XCSE) and a packet-switching layer (XPFE).

The Packet Forwarding Elements are the main elements of the XFEs and have dynamic flow tables which are modified by the controllers depending on the configuration and necessities of the network each moment. The flow-table modifications and the management of the switch are communicated via a dedicated network among the XPFEs and the SDN controllers. The 5G-Crosshaul common frame is the frame format used by the XPFE and is supported by all physical interfaces where packets are transported. Details of the XPFE flow-tables internally are given in III-A.

This component of the XFEs is designed to support different types of traffic over the same switching element and to support multi-tenancy within the network. It is enabled by the label mechanism based in MAC-in-MAC that includes the XCF, which is shown in Fig. 2.

III. XPFE PIPELINE

This section focuses on the configuration of the 5G-Crosshaul packet forwarding element and its performance, for which some measurements have been done and analyzed.

A. XPFE Pipeline Description

We defined an OpenFlow pipeline for the XPFEs to forward XCF frames, to encapsulate tenant frames into XCF frames and to decapsulate them vice versa. The pipeline is based on the L2-L3-ACL example in [7]. As the XCF is based on PBB we omitted the L3 part of [7] completely. We define three different paths through the pipeline as shown in Fig. 3. One path for forwarding XCF frames to another XPFE. One path

to decapsulate XCF frames and to deliver them to directly connected hosts. And one other path to handle frames of directly connected hosts, either by encapsulating them or by sending them to other directly connected hosts.

To determine the correct treatment of received frames (Table 0 in Fig. 3), we divide the ports logically into two groups of UNI (User Network Interface) and NNI (Network Network Interface) ports. At UNI ports the hosts of tenants can be connected, assuming non-XCF Ethernet frames are exchanged. At NNI ports other XPFEs can be connected, exchanging XCF frames. XCF frames have an outer VLAN around the PBB header, therefore the outermost Ethertype is the same for both XCF and non-XCF frames. As a consequence we use the UNI/NNI port distinction to distinguish traffic from tenant hosts and from other XPFEs and record this distinction in packet metadata. Frames received at a NNI port are distinguished further: When the destination address of a received frames matches a specific MAC address, acting as an XPFE identifier, then the frame is decapsulated, otherwise it is forwarded.

The *forwarding* path is the most simple, the egress port is determined based on the destination MAC address and the outer VLAN Id (Table 253 in Fig. 3), allowing different forwarding decisions for different services or tenants.

For frames to be *decapsulated*, a tunnel Id is recorded in metadata (Table 3 in Fig. 3) and the PBB header is removed. The egress port of the resulting frame is determined based on the destination address and tunnel Id (Table t,4 in Fig. 3). Multicast frames are forwarded to all UNI ports of the service (Table t,5 in Fig. 3). To prevent that a decapsulated multicast frame is encapsulated again, a split horizon is implemented for frames received at UNI or NNI ports. The distinction is based on the recorded metadata.

In the *encapsulation* path, the service of the tenant is determined based on ingress port and customer VLAN (Table t,1 in Fig. 3), thereafter the priority is determined (Table t,0 in Fig. 3). Both information is recorded in metadata. For the sake of simplicity, we assume that frames at a UNI port have a VLAN header, allowing to determine service and priority easily. The network operator may actually use a different
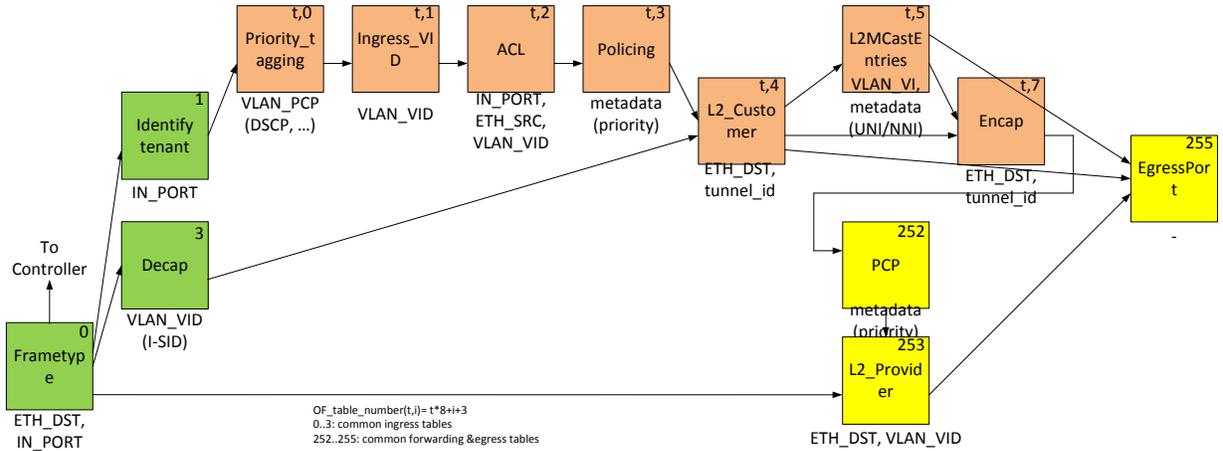
Fig. 3: OpenFlow Pipeline for XCF Forwarding and Encapsulation/Decapsulation

priority on the frame than a tenant. Priority remarking has to be negotiated among operator and tenant. Note, the original frame is not changed, the changed priority is carried in the outer VLAN.

We define access control and policing (Tables t,2 and t,3 in Fig. 3), allowing operators to enforce the SLAs with their customers. Frames passing these tables are delivered locally or forwarded via another XPFE. In the latter case, it is encapsulated (Tables t,7 and 252 in Fig. 3) and the same forwarding decisions (Table 253 in Fig. 3) as for received XCF frames are taken.

In the pipeline we mapped separate functionalities to separate tables, although some of the tables could be combined from a technical perspective. The complexity of the pipeline is comparable to the one for MPLS as described in [8]. This MPLS pipeline contains handling of OAM frames, which is still missing from our pipeline. OAM frames could be determined in the first table based on Ethertype, the UCA (Use Customer Address) bit in the PBB header, or their address, and forwarded to a local entity for further treatment or to another XPE for e2e monitoring.

B. Multi-tenancy Support

As explained in Section II-C, the XPFE supports multi-tenancy using the tags in MAC-in-MAC (or PBB). In the forwarding path it considers the VLAN Id or the PBB service ID, which results in support for a large number of tenants in a network. For encapsulation/decapsulation the corresponding parts of the pipeline are replicated per tenant (Tables t,0 to t,7 in Fig. 3).

OpenFlow supports up to 256 tables. In the shown pipeline, 7 tables are used per tenant and 7 for the common part. Reserving one table for future extensions about 30 separate tenants can be supported in this way. In case a physical server with many services of different tenants is connected to the XPFE, more than 30 tenants have to be handled for encapsulation/decapsulation. In such a case the flow entries for several tenants have to be added to the same set of tables.

The entries for different tenants are distinguished by a match for the tunnel id.

IV. PERFORMANCE ANALYSIS

This section focus on the performance of the XPFE in order to analyse the behavior of the end-to-end delay in a 5G-Crosshaul network where those are the forwarding elements.

A. Measurement Scenario

This section describes the performance measurements of the XPFE focusing on the procedures of encapsulation/decapsulation in XCFs. The goal is to measure and analyse the latency of the packets in a scenario of 4 switches connected as a daisy chain topology. Each switch consists on a PC with an Intel Core i5-3610ME@3.3Ghz processor and 8GB of RAM and is equipped with a 1Gbps Intel 82574L 4 port NIC.

The switch PCs use Ubuntu 16.04 and the software switch Lagopus [9], which is based on DPDK [10]. DPDK actively polls the NIC to avoid the overhead of interrupts. The forwarding and access to the NIC is done in two of the four logical cores, while the general Linux tasks and control of the switch is done in the other two cores. Also we configure Lagopus to write burst sizes of 1 and read burst sizes of 16 frames.

The testbed contains also a more powerful PC with the Moongen traffic generator [11] installed, generating traffic with a fixed rate and packet size. This PC uses an Intel i350 NIC with hardware timestamping.

Our target is to measure the end-to-end latency each packet experiences. For these measurements we use different packet sizes, i.e. 500B, 750B, 1000B, 1250B and 1500B. We also variate the rate at which the traffic is sent from the traffic generator, i.e. 250Mbps, 500Mbps 750Mbps.

B. Measurement Analysis

This section provides the statistical analysis of the measurements obtained from the scenario explained in Section IV-A.

The histograms of the measurements show two different behaviors in the delays depending on the parameters for rate
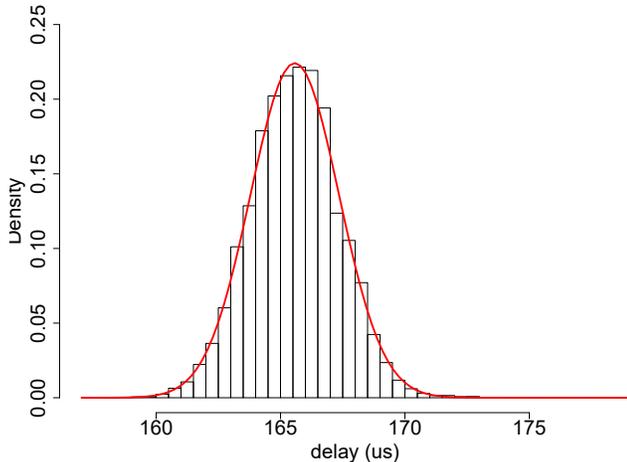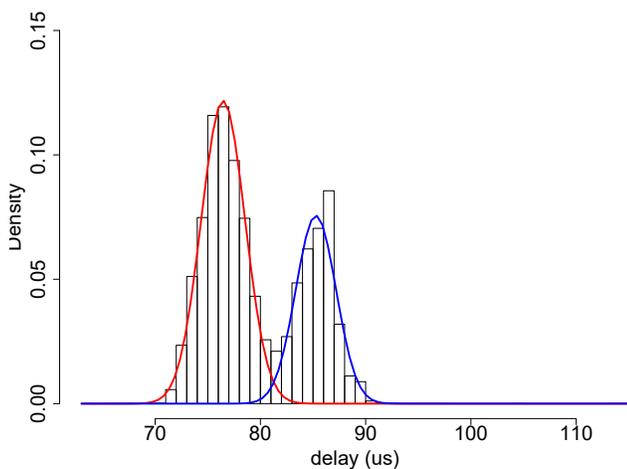
Fig. 4: Histogram pipeline 4 hops 750Mbps 1500B

Fig. 5: Histogram pipeline 4 hops 250Mbps 500B

parameters for the mixture of two Gaussians can be obtained using the Expectation-Maximization (EM) algorithm, which takes into account that one tail of each Gaussian can be inside the other Gaussian curve.

These procedures allows us to obtain the curves that better fits our data, as shown in colors red and blue in figures Fig. 4 and Fig. 5.

However, we showed graphically only one example of each of the two behaviors we observed in spite of having performed the measurements taking also 2 hops (2 switches in the pipeline), 3 hops and 4 hops. These results are gathered in Table I.

Furthermore, the observation of the behavior for each measurement individually allows us to obtain the values of the Gaussian or mixture of Gaussians distribution that better represent the data. From these analytic values we can infer confidence intervals for the mean, 1-percentile and 99-percentile values of the delay for each of the parameters considered. These confidence intervals enables to determine if the traffic we are considering in the 5G-Crosshaul can be transported in the network using some resources that are allocated some hops away. These confidence intervals are gathered in Table II.

With the measurements we performed we obtain that CPRI cannot be transported more than 4 hops, and only if it uses 500 Bytes as packet size and 500 Mbps as maximum rate. That happens because its delay requirement is 100 μs and the confidence interval shows that the 99-percentile value of delay for this parameter is between 90.57 and 91.61 μs with a probability of 0.95. These measurements are planned to be extended to 10 Gbps links in order to make more realistic assumptions and take a more adjusted number of hops from the sources available to support the CPRI flows. When we are dealing with other fronthaul types, they can be transported more than 4 hops independently of its packet size, because they delay requirement is at least 1 ms.

V. CONCLUSION

This document explains the 5G-Crosshaul Packet Forwarding Element, which is an essential element of the 5G-Crosshaul architecture. The XPFE enables the forwarding of packets of traffic for different nature traffic types. Thus, this element integrates multiple technologies to be able to deal with all the traffic types with their multiple requirements and their different nature. One important feature of this forwarding element is that it enables multi-tenancy due to the tags of MAC-in-MAC. The labels allow to differentiate the tenant but enabling at the same time to share the same physical resources. This last feature is important, since it allows to share the same physical infrastructure among tenants using resources in a more efficient way.

Thus, we measured the performance of the forwarding element in terms of the end-to-end delay suffered by the packets. A testbed based on the pipeline has been deployed to measure the performance varying several parameters such

and packet size. These behaviors are analyzed in this section, where we consider a representative example of each behavior.

On the one hand we have a first case, where all values are grouped around the mean, behaving as a Gaussian variable, as shown in Fig. 4. We tried other distributions such as Gamma, logNormal and exponential and among these different distributions the one with less error is the Gaussian.

On the other hand we have a second pattern, in which the values are grouped around two Gaussian shapes. We observe a bimodal distribution, as in the particular case of 4 hops, 250Mbps and 500B shown in Fig. 5.

Therefore, both behaviors can be explained with the Gaussian distribution. The first case fits with a single Gaussian distribution, while the second case is a mixture of two Gaussians. The single Gaussian distribution can be fitted using the mean and the variance of the data as estimators of the mean and the variance for the Gaussian curve. Meanwhile, the estimation of

Size	250Mbps			500Mbps			750Mbps		
	Average	1-percentile	99-percentile	Average	1-percentile	99-percentile	Average	1-percentile	99-percentile
2 hops									
500B	39.07	31.44	45.38	42.08	37.99	46.62	48.89	46.65	53.93
750B	45.66	34.01	52.98	52.16	47.51	56.67	56.43	54.66	59.78
1000B	51.90	38.70	62.81	59.66	54.06	64.07	66.51	64.27	68.67
1250B	56.32	41.01	71.66	68.25	60.82	74.29	73.44	70.82	76.95
1500B	61.23	46.42	84.05	80.76	72.49	84.97	83.56	80.34	87.36
3 hops									
500B	59.37	51.90	66.15	64.54	60.99	69.61	72.77	69.94	79.08
750B	72.28	57.97	79.98	77.89	72.86	81.10	82.56	80.16	87.97
1000B	89.11	70.78	93.69	92.17	85.95	98.62	98.01	91.38	103.31
1250B	87.73	78.07	109.15	103.22	97.14	109.35	108.60	105.42	113.08
1500B	107.66	89.17	123.70	121.21	111.40	126.10	123.01	119.25	127.30
4 hops									
500B	79.66	72.25	89.04	85.66	81.09	90.93	98.03	93.66	104.29
750B	102.49	87.29	108.48	105.15	99.70	109.51	110.64	107.74	115.84
1000B	121.29	104.35	128.98	122.80	117.80	126.73	130.75	128.09	133.19
1250B	125.98	113.50	145.54	141.60	134.17	147.33	145.24	141.75	150.22
1500B	152.49	129.22	164.06	157.15	149.46	166.74	165.57	161.44	169.77

TABLE I: Delay (μs) measurements pipeline

Size	250Mbps			500Mbps			750Mbps		
	Mean	1-percentile	99-percentile	Mean	1-percentile	99-percentile	Mean	1-percentile	99-percentile
2 hops									
500B	[38.94, 39.19]	[30.14, 31.35]	[45.01, 45.69]	[41.97, 42.21]	[36.99, 37.91]	[46.28, 47.22]	[48.82, 48.97]	[46.19, 46.81]	[53.41, 54.49]
750B	[45.40, 45.93]	[30.69, 33.08]	[52.96, 54.96]	[52.05, 52.27]	[47.25, 48.16]	[56.18, 57.11]	[56.37, 56.48]	[54.23, 54.89]	[59.41, 61.11]
1000B	[51.56, 52.20]	[34.26, 37.00]	[63.30, 64.51]	[59.51, 59.81]	[53.26, 54.42]	[64.88, 66.03]	[66.46, 66.56]	[63.46, 64.75]	[68.49, 69.84]
1250B	[56.06, 56.60]	[28.21, 33.13]	[79.54, 84.60]	[68.16, 68.35]	[60.06, 60.82]	[73.89, 74.57]	[73.36, 73.52]	[70.18, 70.77]	[76.37, 77.54]
1500B	[60.93, 61.55]	[40.15, 42.94]	[84.06, 84.82]	[80.62, 80.92]	[71.55, 73.42]	[84.49, 85.30]	[83.46, 83.66]	[79.51, 80.22]	[86.90, 87.62]
3 hops									
500B	[59.22, 59.50]	[50.58, 51.88]	[66.02, 66.72]	[64.44, 64.63]	[60.62, 61.32]	[67.77, 68.44]	[72.67, 72.87]	[68.99, 69.97]	[78.11, 79.24]
750B	[72.13, 72.42]	[55.19, 57.01]	[80.18, 81.09]	[77.79, 77.99]	[73.91, 74.59]	[81.19, 81.88]	[82.49, 82.64]	[79.76, 80.46]	[87.05, 88.27]
1000B	[88.91, 89.32]	[70.03, 73.49]	[93.91, 95.16]	[91.97, 92.36]	[84.61, 85.93]	[98.42, 99.77]	[97.90, 98.11]	[90.73, 91.54]	[103.35, 104.17]
1250B	[87.54, 87.91]	[77.21, 78.32]	[110.88, 113.86]	[103.05, 103.40]	[96.05, 97.34]	[109.12, 110.40]	[108.51, 108.69]	[105.03, 105.66]	[112.58, 114.32]
1500B	[107.36, 107.95]	[93.38, 95.61]	[123.24, 124.24]	[121.10, 121.33]	[109.86, 112.24]	[125.57, 126.40]	[122.90, 123.11]	[118.67, 119.44]	[126.55, 127.34]
4 hops									
500B	[79.38, 79.93]	[71.45, 72.48]	[88.43, 89.38]	[85.53, 85.82]	[79.66, 80.74]	[90.57, 91.61]	[97.92, 98.14]	[92.83, 94.20]	[103.66, 105.02]
750B	[102.32, 102.67]	[85.89, 88.61]	[108.31, 109.24]	[105.02, 105.30]	[99.02, 100.27]	[110.04, 111.26]	[110.54, 110.75]	[106.05, 106.83]	[114.44, 115.24]
1000B	[121.06, 121.53]	[102.90, 106.13]	[128.51, 129.83]	[122.67, 122.93]	[117.49, 118.47]	[127.19, 128.10]	[130.69, 130.82]	[126.20, 127.95]	[133.15, 135.02]
1250B	[125.77, 126.21]	[113.52, 115.16]	[145.58, 147.61]	[141.45, 141.77]	[134.68, 135.85]	[147.39, 148.61]	[145.14, 145.35]	[141.04, 141.82]	[149.41, 151.53]
1500B	[152.21, 152.78]	[127.38, 130.82]	[163.48, 165.43]	[157.02, 157.30]	[148.90, 149.88]	[165.43, 166.74]	[165.46, 165.67]	[161.06, 161.85]	[169.30 170.09]

TABLE II: 95% Confidence Interval for the mean, 1-percentile and 99-percentile of delay (μs)

as the packet rate, the size and the number of hops in the scenario.

Furthermore, these measurements have been analyzed and we found the data behaves in some cases as a single Gaussian distribution and in other cases as a mixture of two Gaussian shapes. Thus, from these distributions we have inferred confidence intervals for the mean, maximum and minimum values.

In addition, comparing the confidence intervals with the delay requirements of fronthaul traffic we conclude that the CPRI traffic, which requires 100 μs, cannot be transported more than 4 hops through the network, even less if the packet size is higher than 500 Bytes or the transmission rate is higher than 500 Mbps. That means that the 5G-Crosshaul processing units for this traffic type have to be allocated at most at 4 hops from the Remote Radio Heads that generate CPRI traffic.

Finally, this work is intended to be continued by better analyzing the delays and its behavior in a scenario with 10Gbps links including a mixture of different types of traffic which requires different QoS, which becomes a scenario that is closer to the real deployments.

Acknowledgment

References

[1] Cisco Visual Networking Index, "Global mobile data traffic forecast update, 2016-2021," *Cisco white paper*, 2017.

[2] Open Networking Foundation (ONF), "SDN Architecture, Issue 1.1," 2016, [accessed on 12 April 2017]. [Online]. Available: https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR-521_SDN_Architecture_issue_1.1.pdf

[3] A. De la Oliva, J. A. Hernández, D. Larrabeiti, and A. Azcorra, "An overview of the CPRI specification and its application to C-RAN-based LTE scenarios," *IEEE Communications Magazine*, vol. 54, no. 2, pp. 152–159, 2016.

[4] S. González, A. de la Oliva, X. Costa-Pérez, A. DiGiglio, F. Cavaliere, T. Deiß, X. Li, and A. Mourad, "5G-Crosshaul: An SDN/NFV control and data plane architecture for the 5G integrated Fronthaul/Backhaul," *Transactions on Emerging Telecommunications Technologies*, vol. 27, no. 9, pp. 1196–1205, 2016, ett.3066. [Online]. Available: http://dx.doi.org/10.1002/ett.3066

[5] X. Li, J. Núñez-Martínez, R. Casellas, S. González, C. F. Chiasserini, J. R. Sanchez, D. Siracusa, L. Goratti, D. Jimenez et al., "Innovations through 5g-crosshaul applications," in *Networks and Communications (EuCNC), 2016 European Conference on*. IEEE, 2016, pp. 382–387.

[6] IEEE 802.1 Task Group, "IEEE 802.1ah-2008 - IEEE Standard for Local and Metropolitan Area Networks âĂŞ Virtual Bridged Local Area Networks Amendment 7: Provider Backbone Bridges."

[7] Open Networking Foundation (ONF), *OpenFlow Table Type Patterns, Version 1.0*, August 2014.

[8] ——, *MPLS-TP OpenFlow Protocol Extensions for SPTN, Version 0.9, ONF TS-029*, January 2017.

[9] "Lagopus," [accessed on 12 April 2017]. [Online]. Available: {http://www.lagopus.org}

[10] "DPDK," [accessed on 12 April 2017]. [Online]. Available: {http://dpdk.org/}

[11] "Moongen," [accessed on 12 April 2017]. [Online]. Available: {http://scholzd.github.io/MoonGen/index.html}