Packet-Optical Transport for Vertical Services in the Path to 6G

Paola Iovanna*
Ericsson Research
Ericsson
Pisa, Italy

Giulio Bottari
Ericsson Research
Ericsson
Pisa, Italy

Andrea Sgambelluri
Scuola Superiore Sant'Anna
Pisa, Italy

Piero Castoldi
Scuola Superiore Sant'Anna
Pisa, Italy

***Abstract* —5G evolution and 6G networks will have a crucial role in supporting the acceleration of industrial digitalization and new applications such as cloud robotics and remote-assisted surgery. The goal is to support differentiated quality of services including the critical ones, without increasing the cost/revenue ratio for network operators. The transport network plays a crucial role in supporting the physical connections of the interfaces of a mobile 3GPP network, including the interfaces between the Radio Access Network (RAN) and the Core Network (CN) and within the RAN itself. Here, optical technologies can play an essential role by providing challenging E2E QoS for vertical industries. The main challenge is to guarantee E2E QoS while avoiding over-provisioning on network resources. The use of a combination of packet and optical technologies, supported by AI-based orchestration mechanisms, can provide flexibility, and dynamically optimize the use of network resources. A specific case study was investigated and assessed through a pilot in the context of the H2020 5Growth project.**

Keywords—5G, 6G, mobile networks, transport, Industry 4.0, mobile fronthaul, mobile backhaul, smart manufacturing, robotics, integrated photonics.

I. Introduction

The acceleration of industrial digitalization and the rise of new applications such as cloud robotics and remote-assisted surgery are leading to a high demand for the new capabilities of 5G networks. At the same time, future use cases like massive twinning, immersive telepresence, and collaborative robotics are helping to shape the journey toward 6G by asking ever more challenging requirements. There is also the need to achieve ubiquitous radio access, with peculiar performance in specific deployment areas, supporting a high density of users and access points with the mobile the network. The goal is to satisfy all these needs without increasing the cost/revenue ratio for network operators.

The transport network has the role to support the physical connections of the interfaces of a mobile 3GPP network, following all the evolutions of radio generations and 3GPP releases. This includes the interfaces between the Radio Access Network (RAN) parts and the Core Network (CN) parts, referred to as mobile backhaul and the interfaces within the RAN itself, e.g., the digital interfaces between Remote Radio Units (RRU) and the related Digital Units (DU)/Baseband Units (BBU), referred to as mobile fronthaul. In essence, transport networks provide the connectivity of the RAN/CN functions among the multitude of radio sites. This network is constituted by wired and wireless transmission media which interconnect switching and/or routing systems in various topological arrangements.

In this context, optical technologies can play an essential role which will increase in the future generation of mobile networks. In fact, the evolution towards 6G impacts the transport network in several aspects. First, RAN and CN will be characterized by high modularity and cloudification. Second, an increasing number of services for vertical industries will require challenging E2E QoS (throughput, latency, availability, reliability) to be ensured from a local to a geographical dimension.

The main challenge in the deployment of a transport infrastructure is to guarantee E2E QoS while avoiding over-provisioning on network resources. Sizing the network at the peak of the expected traffic, as it was done in the past, is becoming more and more impractical in term of costs. Rethinking transport implies to handle the current and future services with mostly automatic operations, supported by AI-based orchestration mechanisms, to dynamically optimize the use of the network resources while fulfilling the agreed service levels. The right combination of packet and optical technologies allows to combine the flexibility of the packet layer with the high bandwidth and low latency of the optical layer. Anyway, simple downsizing and reusing conventional transport systems, tailored for regional and long-haul segments, does not match the constraints of cost, size and power consumption that are typical of the access area. Hence, a new paradigm is needed, based on novel photonic components and on smart "transport-aware" orchestration mechanisms that which favor a real integration of optical connectivity with the packet flows it carries.

To demonstrate key components of this new transport paradigm, including the mentioned orchestration engine, a specific case study has been investigated and assessed through a pilot hosted in the COMAU and TIM premises in Turin, Italy, in the context of the H2020 5Growth project [1]. COMAU is an Italian multinational company in the automation field, part of the automaker group Stellantis. Starting from 2020, an experimental area has been covered with an Ericsson's 5G radio network (RAN/CN) supported by three main innovative components: a transport network, including optical system deployed in Ericsson research laboratories, a transport aware orchestrator to manage different QoS and a monitoring module that interworks with the orchestration to enable remote analysis of network performances. The system allows to convey radio traffic with the desired E2E QoS for three different use cases which presented profiles of high interest for COMAU. In fact, in a smart factory, many use cases shall be typically supported in a parallel manner over the same radio/transport infrastructure.

* Contact author email: paola.iovanna@ericsson.com

This work was partially supported by the EC H2020 5GPPP 5Growth project (Grant 856709) and by the European Union under the Italian National Recovery and Resilience Plan of NextGenerationEU, partnership on "Telecommunications of the Future" (PE00000001 - program "RESTART").

More specifically, the pilot has demonstrated the support of low latency communication for synchronized digital twin of robots, massive machine type communication for real-time asset monitoring through pervasive sensors and enhanced mobile broadband to enable immersive telepresence for technical support from a remote location.

In addition to this introduction, the paper is constituted by five sections. *Section II* illustrates considerations on the end-to-end requirements of vertical services and the related slicing implications. *Section III* provides an overview of the transport enabling technologies. *Section IV* describes the pilot deployed on the vertical premises illustrating its main components. *Section V* provides insights on the network infrastructure which supports the pilot. *Section VI* reports tests and measurements made during three validation campaigns carried out by a monitoring platform integrated in the orchestration platform that allows to assess the performance of the overall integrated solution from remote.

II. Vertical services requirements in the path to 6G

Based on Service Level Agreements (SLAs), new services for vertical industries, including those that demand extreme performance, will require E2E QoS at scale over three main network deployment areas: local areas, confided wide areas, and general wide areas. Examples of local areas include factory shopfloors and warehouses, while examples of confined wide areas include ports and railways. General wide areas refer to larger geographical locations such as cities.

E2E QoS is the combination of the QoS in the radio layer RAN/CN and the QoS in the underlying transport layer. An optimal combination can be achieved by using automatic and dynamic techniques to create smart mapping between the RAN/CN QoS and the related transport QoS by considering the specific technology of the infrastructure. For example, the 5G QoS Identifiers (5QI) of a 5G RAN could be mapped to the corresponding Differentiated Services Code Points (DSCP) of the underlying transport.

Since the transport layer is logically separated from the radio layer and the expected radio needs are known, the conventional approach is to dimension and allocate transport resources on the peak of expected radio traffic. The downside of this approach is that it frequently results in the over-provisioning of the transport layer, which may not always be feasible or economically justifiable. As an alternative to this conventional approach, in which the QoS of the RAN/CN and the QoS of transport are orthogonally associated and independently configured, we have developed an approach that avoids over-provisioning by making the radio layer orchestration aware of the transport resources. The traffic flows, each aggregating homogenous services, are associated on a specific slice and mapped to the most appropriate transport connections in a shared transport network. It is important to note that the transport-aware approach requires some changes to current 5G slicing practices. Firstly, it requires that the service and its deployment area are associated to a particular slice. Secondly, it requires that the slice tracking areas (TA) are chosen to cover the deployment areas of the service. Thirdly, the E2E QoS parameters of the service need to be mapped to the corresponding network resources (RAN, CN and transport) associated to the slice by using a RAN/CN and transport abstraction to expose a suitable view of the network resources to the orchestrator [2]. When all of these conditions are met, a consistent association of the QoS of slices in the RAN/CN with the QOS of transport can be performed automatically, and both layers can be automatically configured.

III. Transport Enabling Technologies

Optical transmission and switching are fundamental tools for the transport infrastructure, but the challenge is to implement links and nodes whose cost is an order of magnitude lower than corresponding elements in conventional optical transport networks. For instance, hundreds of Gb/s transmission and wavelength selective switch (WSS)-based optical switches have already been commercialized for metro networks for many years. However, those systems are too costly for the radio access segment. One order of magnitude of cost reduction for optical transceivers and one or two for optical switches represent typical research challenges. Additionally, vertical industries, such as manufacturing plants and logistics systems, will require specific telecom architectures based on new forms of optical transport [3].New integrated silicon photonic technologies can enable low-cost transmission and switching systems in different segments of the network, simplifying the architecture of some of those systems while performing better than electronics for some key features. Silicon photonics allows for the integration of multiple optical and electronic components, such as waveguides, modulators, and detectors, on a single chip. This increases the density and functionality of devices, reducing costs and size. Silicon photonics allows for the transmission of data at high speeds, which is critical for 5G and 6G networks that require high bandwidth and low latency. Silicon photonics can be easily scaled up to meet the demands of future high-bandwidth networks [4][5][6].

Such new optical technologies are ideal to be combined with packet switching system at the edges of the transport network. In the context or radio access for 5G, its evolutions, and 6G, this integration is particularly important for the fronthaul transport of data between the antenna side and the core network. Specifically, packet-optical integration can be effectively used to transport Common Public Radio Interface (CPRI) and enhanced CPRI (eCPRI) signals over optical channels, which allows for the efficient and flexible transport of data, as well as the aggregation of multiple services on the same optical channel. In the pilot optical system has been combined with packet switches as first step to assess the packet-optical performances. It is planned to evolve the transport network architecture by including the integrated silicon photonics switches.

IV. Pilot on Vertical Premises

In a smart factory [7], many use cases are realized, indoor, in a parallel manner. They demand specific, and often challenging, performances to the telecommunication network which serve the industrial plant. Failing to provide such performances would immediately translate into bottlenecks in the manufacturing process. To investigate the use of cutting-edge cellular technologies in this scenario, supported by a packet-optical transport network, a specific pilot has been deployed in three years (including design, implementation, and validation) on the shop floor of COMAU connected to the TIM premises (i.e., the remote central office for the pilot) [8].

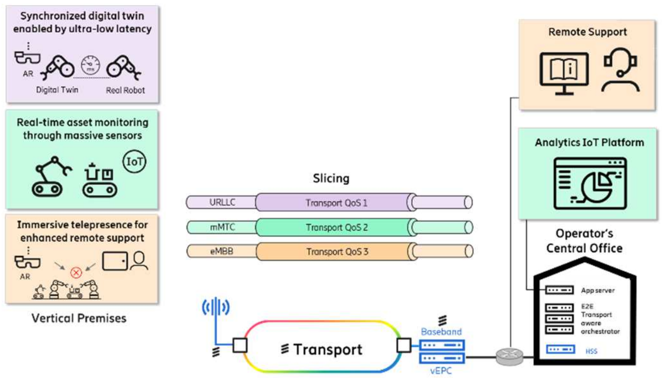

Fig. 1. Architecture of the pilot including transport network

Fig. 1 illustrates the pilot architecture at high level. The pilot architecture includes robotics systems and machineries, a 5G Non-Stand Alone (NSA) network supported by a shared transport network based on optical technologies, cloud platforms hosting a specific orchestrator [2] developed to manage the mobile and packet-optical networks. The clouds platform run also vertical applications related to the use cases in place. The pilot consists of three use cases, each with specific requirements for the radio network and, indirectly, for the transport and cloud infrastructures. The use cases are represented with three different light colors on the left side of Fig. 1.

A. *Synchronized Digital Twin*

A digital twin is a virtual model used in smart manufacturing to optimize a particular aspect of operations. It provides a digital representation of the entire operation, enabling manufacturers to determine how best to streamline the production environment without physically changing any processes. With the digital twin, manufacturers gain a rich visualization of the current and past state of the production environment, enabling them to plan for future states with "what-if" scenarios.

The first use case deployed in the pilot captures the motion of a real robot and, through an ultra-low latency radio link, produces a synchronized digital twin representation of the robot. The movement of the mechanical robot and of the respective virtual renderings are perfectly aligned in time. In a real plant, by gathering data in real time from all the robots and machineries, manufacturers can create a fully integrated digital view of the assets involved in the production line. Obviously, the effective use of the digital twin requires a massive amount of data to be acquired and elaborated in real-time. It is possible thanks to a highly reliable and stable connectivity providing a low deterministic latency. In this respect, this use case shall be served by an Ultra-Reliable Low-Latency Communication (URLLC) traffic profile.

B. *Real-time monitoring of industrial assets*

It is not efficient to send a technician to control the status of a machine that does not need attention, as it is not productive if a machine goes down because it has not received the attention it requires and it suddenly breaks. Automated monitoring of asset condition enables manufacturers to optimize maintenance, preventing unplanned stops on the production line and premature replacement of spare parts. Specifically, predictive maintenance, which uses data collected by a massive number of sensors, can determine exactly when maintenance needs to be performed by enforcing data analytics techniques. The status of key systems can be analyzed via a centralized dashboard, whether the assets are on a single factory or spread in multiple sites, and aggregated data can be used to pre-emptively schedule maintenance and to identify possible sources of downtime.

To investigate the use of cellular connectivity in this scenario, the second use case deployed in the pilot is dedicated to demonstrating such real-time, pervasive, monitoring of the industrial assets. Data captured from wireless sensors, placed on specific machineries, are sent to a software application running on a server which is in a data center of the network operator. This application elaborates the acquired data to determine actions for preventive maintenance, to enforce more accurate production planning forecast, for facilitate quality improvements and to capture other insights in real time. Monitored parameters are temperature, pressure, vibrations.

This use case requires a Massive Machine-type Communication (mMTC) profile to provide scalable and efficient connectivity for a massive number of sensors periodically sending few packets which contains the measured values.

C. *Immersive Telepresence for Enhanced Remote Support*

The third use case demonstrates immersive telepresence for enhanced remote support where the maintenance staff is assisted by a remote expert to investigate and solve a failure using augmented reality (AR) and step-by-step digital tutorials.

In this use case, AR is used to reduce training time for staff for complex tasks, minimize operator error, and allow remote assistance. Instructions can be rapidly visualized is specific AR devices or on a tablet, allowing experts to support on-site personnel remotely. To enable AR within the manufacturing environment, the facility must be able to provide reliable, , high throughput, real-time wireless data transfer to deal with the rich media and complex analysis required, especially as the AR glasses move with the employee throughout the floor. In this respect, this use case is supported by an enhanced Mobile BroadBand (eMBB) traffic profile.

V. NETWORK INFRASTRUCTURE

In the vertical premises of COMAU, the experimental area has been covered with 5G (RAN/CN) connected to the central office of TIM. The 5G network is deployed according to the NSA architecture based on Ericsson radio nodes and antennas. Here, the 5G RAN is used in conjunction with a 4G network and the related CN. In the initial commercial deployment of 5G radio, the NSA option has allowed operators to leverage their existing network investments in transport and CN rather than deploying an entirely E2E 5G network. Dedicated 5G terminals have been used to connect the various components of the pilot (robots, controllers, AR glasses, sensors, etc.).

A shared transport network connects the 4G's DUs and 5G's BBUs to the corresponding radio units and antennas, implementing a Centralized RAN (CRAN) architecture, where DUs and BBUs are centralized in a single hub location and serves several remote radio sites. The transport network, illustrated in Fig. 2, is a DWDM network with 25 Gbit/s optical channels operating on a ring topology. Two Optical Add and Drop Multiplexers (OADM) connects two packet switches located at a remote node and at a hub node.

Fig. 2. Transport Network arrangement

The following client traffics are transported by the network to show the concurrent support different traffic:

- One CPRI Option 3 flow (2.5 Gb/s) for connecting RRU to DU in 4G
- Two eCPRI (10 GBE) for connecting Advanced Antenna Systems (AAS) to the BBU is 5G

Cloud platforms, located on the vertical's site and on the operator's central office, enable the implementation of Network Function Virtualization (NFV) and the support of vertical applications running from remote. An E2E transport-aware slice orchestration, developed in the project, that automatically manages the interworking among radio, packet-optical transport, and cloud on the selected local deployment area.

On the left side, Fig. 3 reports a picture of part of the experimental area where is visible the rack which hosts most of the systems located on vertical's premises. Fig. 3 also reports a zoom on the rack with more details on such systems.

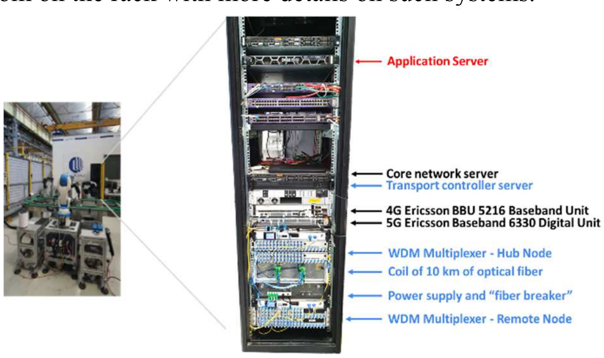

Fig. 3. Details of the equipment installed on the vertical premises

VI. Test and Measurements

Three validation campaigns [9] have assessed the target performances to validate the KPIs associated to the three use cases. Each validation campaign has been organized in sessions, each session lasting at least four days to obtain a significant set of measured values and to intercept possible "spikes" in expected performances. Most specifically:

- The *first validation campaign*, in September 2020, has tested the radio network and the transport network as separate entities. The radio network has been tested on the vertical premises in Turin while the transport network has been tested in the Ericsson laboratories in Pisa. This first campaign has shown that the optical network introduces negligible latency, compatible with critical requirements and enables the deployment of critical use cases also on a geographical distance.
- After the first validation, the transport network systems have been moved to Turin and integrated with the radio network to obtain the E2E infrastructure. On such integrated radio-transport network, a *second validation campaign* has been performed in June 2021. This campaign has provided important indications for the subsequent tuning of 5G network to achieve the best possible performances (especially in terms of round trip latency).
- In September 2021, a *third validation campaign* has verified the network performances in relation to the final KPIs targets validated on the fully integrated pilot after the mentioned network tuning. All the performance parameters were already complying with targets after the second validation campaign except latency for which a specific tuning has been operated before running this third validation campaign. In this final validation round, a specific monitoring platform has been integrated in the pilot to enable measurements from remote. The use of a centralized, automatic monitoring has demonstrated the ability to reduce costs by limiting on-site maintenance.

The set of measurements presented here are relative to the User Plane function of the transport network, to show that the performances of the network are adequate to support desired services without disruption. In that perspective, the reference methodology used is that described in IETF's RFC2544 [10].

Finally, the three use cases have been tested in practice by technicians who will be able to use these technologies in the real world to have an experiential verification of the usability and concrete functionality of each function implemented. For example, technicians have used AR glasses in the third use case to try the concrete improvements in their daily maintenance tasks. The following sub-sections detail the setup of tests and results of measurements.

A. Transport Network Test Setup

Preliminary tests on the transport network have been done in Ericsson labs in Pisa to assess that the performance of the transport network is adequate to support 5G and 4G fronthaul connections, based on the two 10GBE eCPRI links (5G), and on the CPRI link (4G) without any discontinuity.

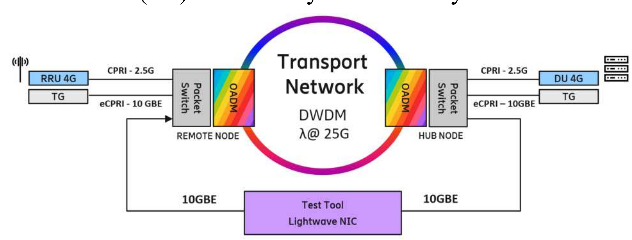

Fig. 4. Transport Network test setup

Fig. 4 reports one of the test setups used for the validation. A test instrument (NIC produced by Lightwave) has been used to perform RFC2544 compliant measurements on a 10 GBE links, while the other 10 GBE and the CPRI 2.5G links were fully loaded by traffic originated by a traffic generator. For what concerns the CPRI link, a functional test has been performed: the test showed that regardless the traffic load of the 10GBE clients, the 4G base station is up and running and no alarms are showed by RBS management system.

The following tests have been performed on the 10GBE links used for eCPRI:

- *Latency*: measures the time taken by a test frame to travel through a network device or across the network. Latency is the time interval that begins when the last bit of the input frame reaches the input port and ends when the first bit of the output frame is seen at the output port.
- *Throughput*: measures the maximum rate at which none of the offered frames are dropped by the device/system under test (DUT/SUT). This measurement translates into the available bandwidth of the Ethernet virtual connection.
- *Back-to-back burst*: measures the longest burst of frames at maximum throughput or minimum legal separation between frames that the device or network under test will handle without any loss of frames.
- *Frame loss*: defines the percentage of frames that should have been forwarded by a network device under steady state (constant) loads that were not forwarded due to lack of resources.

B. End-to-End Test Setup and Remote Monitoring

The measurements in the pilot have been conducted in a testbed located on the shopfloor of the vertical partner COMAU. In this lab setup, a 5G CPE communicates with the Ericsson 5G antenna AIR 6488 in the downlink direction (DL), from the antenna to the CPE, and in the uplink direction (from the CPE to the antenna). Two PC Engines APU2C4, labeled APU_108 and APU_109 in the figure, are connected to the two endpoints of the radio network assessing the transmission performances from the 5G CPE to the radio core network server. The measurements evaluate the performance of the entire radio-transport chain as the optical systems and the fiber coil (10 km) are traversed by the eCPRI flow supporting the radio traffic.

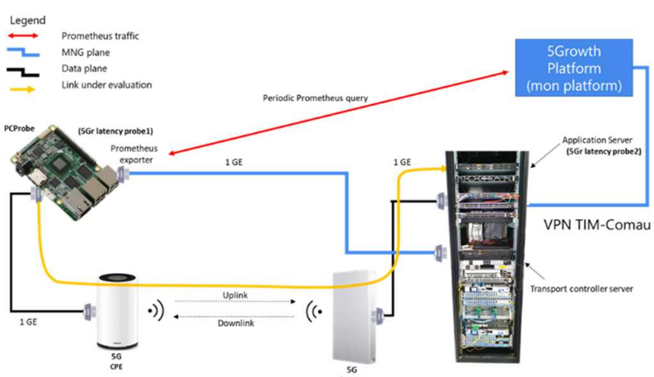

Fig. 5. End-to-End Test Setup

The monitoring platform developed within the 5Growth project has been used to perform automatic measurements on the deployed Network Service from remote, without the need to visit the vertical premises. In November 2021, as a complement of the third validation campaign, a specific campaign of additional measurements has been performed from the SSSA laboratories located in Pisa, Italy, connected to the monitoring platform located in Turin.

Fig. 5 shows the test setup including the monitoring platform represented with a blue box. For the campaign of additional measurements, two ad-hoc defined distributed 5Gr latency probes have been placed in the proper positions to monitor the uplink from the CPE and the application server: the first probe, running on the PC probe, behind the 5G CPE, and the second probe co-located with the application server. The probes have been adopted in the vertical premises to collect and report the data related to unidirectional latency, packet loss and jitter in a Prometheus collector format.

C. Throughput Measurements

The second validation campaign has measured a TCP average throughput of 789 Mbit/s in DL (reported in Fig. 6) and 43 Mbit/s in UL, on a time window of one week. The UDP average throughput has resulted in 855 Mbit/s in DL and 42 Mbit/s in UL. The packet loss has been measured as lower than 2*10-6.

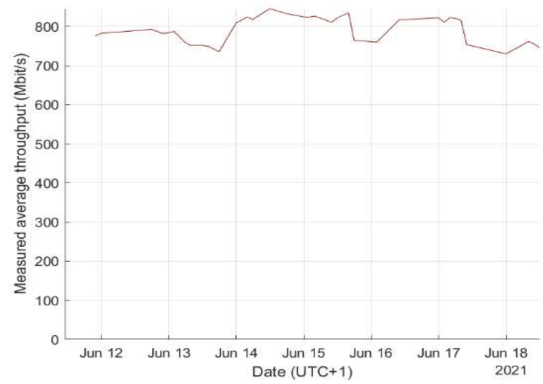

Fig. 6. TCP Throughput in downlink

After the second validation campaign, as said, the network has been tuned for latency optimization while the throughput settings have not been modified. However, new throughput measurements have been repeated with the following results: TCP 761 Mbit/s DL, 65 Mbit/s UL; UDP 952 Mbit/s DL, 66 Mbit/s UL.

D. Latency Measurements

The E2E latency experienced by the clients of the infrastructure, is constituted by the sum of the mobile contribution and the transport contribution. Results of the measurements demonstrate that the transport contribution is two order of magnitude less than the mobile one.

The first use case, i.e., synchronized digital twin, has a strict requirement on latency. This requirement comes from the need of having the digital twin (i.e., virtual representation) in perfect "visual" alignment with the real robot. In other words, the user looking at the digital twin, on a screen or in AR/VR glasses, shall have the perception of synchronization of the two robots: the real robot and its digital replica. The digital twin animation is a sequence of frames constructed by a specific rendering application running on a computer. Each frame, and the resulting animation, are based on the real-time coordinates of the robot axes positions coming from the (remote) robot controller. These coordinated are transferred across the 5G link. The refresh rate between two consecutive frames is 24 Hz which corresponds to 1/24 sec = 41.7 ms between two consecutive frames. When the digital twin is represented in VR or AR glasses, the refresh rate inside the glasses is higher than 24 Hz, to avoid the "cybersickness" effect for the user. Like the ones used in the pilot, commercial AR glasses have a refresh rate in the order 80

Hz or more (the higher the better) which corresponds to 1/80 sec = 12.5 ms. In summary, the AR application receives a new frame from the digital twin rendering application, every 41.7 ms but refreshes the image presented in the glasses to the human every 12.5 ms. The new image in the glasses can be a new one, just received from the rendering application, or the replica of the previous one (in case that a new frame is still not arrived at the end of the current 12.5 ms slot). The time budget of 41.7 ms is composed as follows: transmission over 5G + time waiting for the start of the next processing cycle in the AR program (max 12.5 ms) + processing time of the AR program for refreshing the new internal frame (12.5 ms). The time budget for transmission over 5G is then: 41.7 ms - 12.5 ms - 12.5 ms (max) = 16.7 ms. It is the maximum E2E latency that can be tolerated to have the correct user experience for the first use case.

Considering the automated measurement procedure, performed with the monitoring platform, the link between the CPE and the server where the application is running has been monitored, considering three parameters: the unidirectional latency in uplink (mp1), the jitter (mp2) and the packet loss (mp3). The link has been monitored for three days. Fig. 7 shows the results collected during the automatic measurements. The screenshot has been collected from the GUI of the monitoring platform, considering the Grafana dashboard instantiated to plot, in real time, the data stored in the Prometheus server. In particular, the top panel shows the unidirectional latency in uplink (mp1) evolution during the three days' time interval. The measured average value is 8.73ms (min 5.2, max 12.3), as shown in the legend of the plot. This result is in line with the data reported above, where the third on-site measurements have reported an average latency value of 16.1ms in round trip (DL+UL).

In the measurements done with the monitoring platform, where only the UL direction is considered and the virtualization environment of the application VM is in place, the obtained average result is about half (i.e., 8.73 ms), showing a good match between the two considered measurement procedures (i.e., on site vs remote). The panel in the middle of the dashboard shown in Fig. 7, reports the jitter (mp2) trend during the measurement tests, showing that the jitter presents an average value of 0ms, with values registered in the range -6.8 ms and 6.4 ms. The panel at the bottom of the figure shows the packet loss (mp3) trend. During the tests, no packets have been lost (0 packet lost).

VII. Conslusion

In conclusion, the rapid advancement of industrial digitalization is driving the need for new capabilities in cellular networks. Future use cases are shaping the journey towards 6G with increasingly challenging requirements. The transport network plays a crucial role in supporting the physical connections of the interfaces of a mobile 3GPP network, and optical technologies have the potential to play an essential role in meeting these demands. The main challenge is to guarantee E2E QoS while avoiding over-provisioning on network resources. The use of a combination of packet and optical technologies, supported by AI-based orchestration, can provide flexibility, and dynamically optimize the use of network resources. The optical network enables the support of critical use cases (e.g., in latency) also on geographical areas. This allows to reduce the cost of the services increasing the automation of network management and with the possibility to control the network from remote. Additionally, it opens new opportunities for use cases that require centralized control and support in geographically areas at a lower cost than dedicated networks can provide.

Fig. 7. Automated measurements

The results of a specific case study, investigated and assessed through a pilot in the context of the H2020 5Growth project, demonstrate the potential of this approach in a real industrial shop floor. The pilot performance has been assessed using a monitoring platform that can perform automatic measurements on deployed services, remotely, without the need for in-person visits. Monitored data has verified throughput, latency, packet loss, and jitter. The future plan is to incorporate in the nodes also optical switches that utilize integrated silicon photonics technology.

Acknowledgment

The authors wish to acknowledge Stefano Stracca, Fabio Ubaldi, Marzio Puleri, Luca Giorgi, Simone Panicucci, Claudio Casetti, and Annachiara Pagano for relevant discussions, information about their research work, and support in pilot deployment.

References

[1] H2020 5Growth project, http://5growth.eu/, accessed January 2023.

[2] P. Iovanna et al., "End-to-end network slicing orchestration – A Key Enabler for Industry-Vertical use Cases," in Ericsson Technology Review, vol. 2022, no. 2, pp. 2-10, March 2022.

[3] R. Sabella, P. Iovanna, G. Bottari and F. Cavaliere, "Optical transport for Industry 4.0 [Invited]," in Journal of Optical Communications and Networking, vol. 12, no. 8, pp. 264-276, August 2020, doi: 10.1364/JOCN.390701.

[4] Sorianello V. et al. "Polarization insensitive silicon photonic ROADM with selectable communication direction for radio access networks". 2016, Optics Letters Vol. 41, Issue 24, pp. 5688-5691.

[5] Nadimi Goki P et al. "Lossless ROADM by Exploiting low gain SOAs in fronthaul network". 2019 24th OptoElectronics and Communications Conference (OECC) and 2019 International Conference on Photonics in Switching and Computing (PSC), Fukuoka, Japan, 7-11 July 2019.

[6] T. Komljenovic, D. Huang, P. Pintus, M. A. Tran, M. L. Davenport and J. E. Bowers, "Photonic Integrated Circuits Using Heterogeneous Integration on Silicon," in Proceedings of the IEEE, vol. 106, no. 12, pp. 2246-2257, Dec. 2018, doi: 10.1109/JPROC.2018.2864668.

[7] R. Sabella et al., " Industrial automation enabled by robotics, machine intelligence and 5G," in Ericsson Technology Review, vol. 2018-1, no. 96, pp. 21-26, January 2018.

[8] 5Growth "Execution report of in-house use cases for Pilots" Deliverable 3.6 (D3.6), November 2021

[9] 5Growth "Final validation and verification report", Deliverable 4.4 (D4.4), November 2021

[10] RFC2544, "Benchmarking Methodology for Network Interconnect Devices", available online at https://tools.ietf.org/html/rfc2544