5G key technology enableRs for Emerging media COntent pRoDuction services

Deliverable D5.2
Trials initial deployment

Version v1.0 Date: 2022/04/01

Document properties:

Grant Number:	957102
Document Number:	D5.2
Document Title:	Trials initial deployment
Editor(s):	Federico Maria Pandolfi (RAI)
Authors:	Jan Dürre, Norbert Werner (SEN); Mohamed Nabil Ibrahim (EDD); Baruch Altman (LU); Pablo Perez (NOK); Federico Maria Pandolfi, Giulio Stante (RAI); Emilio Garrido (TID); Panagiotis Matzakos (EUR); Paola Sunna (EBU; Julián Cabrera Quesada (UPM).
Reviewers	Adrián Rodrigo, Cristina Avellan, Nerea Cilleruelo (UPV); Assunta De Vita (RAI)
Contractual Date of Delivery:	2022/04/01
Dissemination level:	PU[1]
Status:	Final
Version:	1.0
File Name:	5G-RECORDS_D5.2_v1.0

Revision History

Revision	**Date**	**Issued by**	**Description**
1.0	2022/04/01	RAI	Final version

[1] CO = Confidential, only members of the consortium (including the Commission Services)

PU = Public

Abstract

Document D5.2 describes the outcomes of the initial trials planned in D5.1. These will be structured around the use cases, each one led by the champion in brackets: Live audio production (Sennheiser), Multiple cameras wireless studio (Ericsson) and live immersive content production (Nokia). The trials will focus on the demonstration and application of 5G technologies involving those applicable to RAN and Core and the newly developed integration components to support the different requirements and functionalities of each use case. The 5G and media production infrastructure capabilities in terms of relevant KPIs will be tested against different configuration and deployment options in order to stress the network conditions and evaluate the limits. In addition, performance and capabilities of edge computing and network slicing with specific SLA will also be assessed. This activity also targets the deployment of VNFs to realize specific functionalities in the targeted scenario. Along the project, the trials will be further developed in order to increase complexity in terms of the use case and its components.

Keywords

Trials, Deployment, Measurements, Testbeds, Results

Disclaimer

Executive Summary

This document describes the outcome of the first stage of trials for the project's use cases along with the changes of plans dictated by mutating conditions during the course of the project. Some of the measurements initially planned, in fact, had to be adapted to the testbeds that were possible to implement.

The **UC1** testbed deployed for the first round of trials was a simplified version of the envisioned architecture and included a single set of UE with 5G modem and audio components to reduce the complexity of scheduling and slicing in the 5GS. Two parallel flavours of the same architecture have been tested using different RAN components and Core Network providers: *(i)* a monolithic version of the OAI RAN (gNB) and the OAI Core Network; *(ii)* a disaggregated RAN components version incorporating the Accelleran CU and the OAI DU. The first set of trials showed that the 5G testbed currently presents an unexpected amount and variety of unpredictable packet loss (up to 1%), especially in DL direction. Since in-depth analysis and optimization are needed to identify the causes of this, packet loss and reliability were not considered in this deliverable, as well as the synchronicity requirement. As for the network latency, the first measurements have shown that the overall E2E latency is still considered to be in the 10's of milliseconds.

In **UC2**, two trials phases were planned, one leveraging the available media components for the deployment and the second one integrating the newly developed features for the various components. Once a dedicated 5G network has been deployed for both phases, several tests have been performed for the two scenarios of this use case: the wireless studio scenario and the remote contribution scenario. Regarding the **wireless studio scenario**, the results of the tests were satisfactory overall and gave valuable information for the development of UC2. The KPIs analysis showed that the available UL throughput and packet loss needs to be improved but the regulatory framework does not offer much room for extra bandwidth. The current glass-to-glass latency value is very close to the proposed limit (40ms). Regarding the **remote contribution scenario**, the tests' results were satisfactory as well, as all the tests, with few exceptions, currently passed. The throughput and latency were well within the expected ranges for low traffic loads (under 50% load emulation) but started deteriorating severely when the network became more congested, especially in the multi-cam scenario. In all tests the UL packet loss rates at the application level were almost zero. Finally, both multi-cam, SMPTE 2110 video functionalities, remote audio communication and remote camera control were successfully validated, resulting in a good overall user experience, with just a couple of SMPTE tests not supported by the current setup.

In **UC3**, two trials phases were also defined. The aim of the first stage was to validate the E2E functionality. The tests focused on a live transmission from a reduced set of cameras from the trial location to a single user in a remote location. COVID restrictions made it impossible to have full access to the trial location and some measurements involving on-field camera deployment will be performed in a second phase. It is worth noting that the target KPIs have been measured successfully thanks to specifically developed video tools (e.g., motion-to-photon latency tool, offline view renderer). The general conclusion is that the system is working fine in the basic scenario (720p at 15-30 fps). Delay measures at the production segment (RTT and motion-to-photon) are slightly underperforming, especially under load, but functional. Delivery KPIs are not fully met either, but they are quite stable even under load. Regarding the functional validation of the automatic slice change, a more detailed analysis of the performance of this functionality will be done in the second phase of the project. During the testing some limitations have been found, mainly in mmWave components availability. However, in the second phase, some potential improvements related to the production console and the media delivery will be implemented.

Table of Contents

List of Figures

List of Tables

List of Acronyms and Abbreviations

The acronyms list has a special style defined as “acronyms”. Each acronym is separated by a tabulation with each definition. As is shown below:

3GPP	3rd Generation Partnership Project
4G	4th Generation of mobile communications systems
5G	5th Generation of mobile communications systems
5GC	5G Core
5GS	5G System
A/V	Audio / Video
AN	Access Network
BE	Best Effort (Slice)
BW	Bandwidth
COTS	Commercial Off-the-Shelf
CPE	Customer-premises Equipment
CPU	Central Processing Unit
CU	Central Unit
DHCP	Dynamic Host Configuration Protocol
DL	Downlink
DNS	Domain Name System
DoF	Depth of Field
DU	Distributed Unit
E2E	End-to-End
eNB	Evolved Node B
FTTH	Fiber-to-the-Home
FVV	Free Viewpoint Video
G2G	Glass-to-Glass
gNB	Next Generation Node B
HLS	HTTP Live Streaming
HTTP	Hypertext Transfer Protocol
HW	Hardware
IEM	In-Ear-Monitoring
IP	Internet Protocol
KPI	Key Performance Indicator
KQI	Key Quality Indicators
LTE	Long Term Evolution (3GPP 4G)
MEC	Mobile Edge Computing
MG	Media Gateway
MG	Media Gold (Slice)
MIMO	Multiple-Input Multiple-Output
MOCG	Media Orchestration Control Gateway
mmW	Millimeter Wave
NAT	Network Address Translation
NR	New Radio
NSA	Non-Standalone
OAI	Open Air Interface
OVS	Open vSwitch
PPS	Pulse-per-second
PTP	Precision Time Protocol
QoE	Quality of Experience
QoS	Quality of Service
RAN	Radio Access Network
RCP	Rate Control Protocol
RF	Radio Frequency
RTP	Real-time Transport Protocol
RTSP	Real-time Streaming Protocol
RTT	Round-Trip-Time
RU	Radio Unit

SA	Standalone
SCS	Subcarrier Spacing
SDI	Serial Digital Interface
SDN	Software-defined Network
SDR	Software-defined Radio
SMPTE	Society of Motion Picture and Television Engineers
SoC	System-on-Chip
SW	Software
TCP	Transmission Control Protocol
TDD	Time-Division Duplexing
TSN	Time Sensitive Network
UC	Use case
UE	User equipment
UL	Uplink
UDP	User Datagram Protocol
UPF	User Plane Function
USRP	Universal Software Radio Peripheral
VM	Virtual Machine
VNF	Virtual Network Functions

1 Introduction

1.1 Scope

This document (D5.2) focuses on the outcome of the first stage of trials for the project's use cases, as they were originally planned in D5.1 [4]. This deliverable also describes the changes to those plans (where needed by mutating conditions during the course of the project) by reporting, in detail, the final architectures that were deployed and measurements that were taken during the reference months.

This document uses as inputs the architectures and the list of measurements outlined in D5.1 [4], the KPIs defined in D2.1 [1] and the measurements tools listed in D4.1 [3].

This deliverable will be used as input for the final stage of trials and the overall technology validation, whose activities will be reported in the subsequent WP5 deliverable, D5.3.

1.2 Objectives

The objectives of this deliverable are:

- To provide a detailed description of the testbeds actually deployed during T5.2 activities.
- To report the measurements taken within each testbed.
- To compare and evaluate the aforementioned measurements against the KPIs previously defined in the project.
- To give insights and useful considerations for the deployment of the final stage of trials, which will lead to the validation of the involved 5G technologies.

1.3 Structure

This document is structured as follows:

- Section 2 describes the outcomes of the trials related to the live audio production use case (UC1).
- Section 3 describes the outcomes of the trials related to the multiple camera wireless studio use case (UC2).
- Section 4 describes the outcomes of the trials related to the live immersive media production use case (UC3).
- Section 5 presents the conclusions of this document.

Each section related to the three Use Cases (Sections 2 - 4) will also include the following subsections: *(i)* deployed testbeds architectures, *(ii)* measurements results, *(iii)* KPIs analysis and *(iv)* deployment considerations for the final stage of the trials.

2 Use Case 1: Live Audio Production

This chapter describes the outcomes of the first stage of trials in the context of UC1.Testbed architecture and measurements' results will be presented first and will be followed by a detailed KPI analysis (with regards to the KPIs defined in D2.1 [1]), as well as the deployment consideration useful for the final stage of the trials.

2.1 Deployed testbed architecture

Figure 1 depicts the currently deployed architecture at EURECOM site for UC1 (containing the 5G RAN + Core, the 5G UE and the Local Data Network blocks) which is based on the architecture presented analytically in D3.1 [2]. As described in D5.1 [4], for the first stage deployment, a more simplified version of this architecture was put in place using a single set of UE with 5G modem and audio components. For a more detailed description of the testbed architecture refer to D3.1 [2].

During the trials, two parallel setups based on the architecture in Figure 1 but with different RAN components and the Core Network provider were used:

- The first setup is based on a monolithic version of the OpenAirInterface (OAI) RAN (gNB) and the OAI Core Network to get measurement results for the first stage trials.
- The second setup is based on the disaggregated RAN components by incorporating the Accelleran CU and the OAI DU from EURECOM. This setup will be using the Cumucore core network.

The reason for using these two parallel setups was to be able to extract the end-to-end measurements for the first stage trials based on the interoperability of the monolithic OAI setup with the Sennheiser equipment, before the end-to-end interoperability between the disaggregated RAN components and the Cumucore Core Network became available.

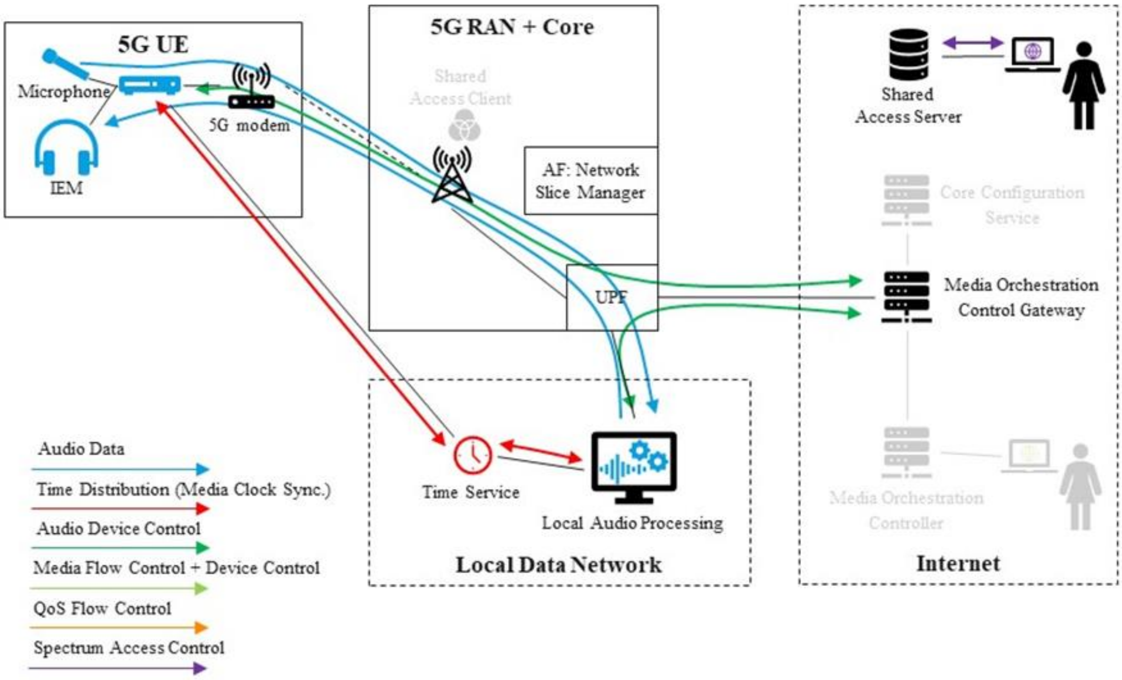

Figure 1 - UC1 deployed testbed architecture – First trial stage

In both setups the OAI RAN software running at the gNB server interfaces with the AW2S Jaguar (2x2, 20 MHz) RU [5]. For the UE portion, the SIMCOM 8200X [6] and the Quectel RM500Q-GL [7] COTS UE modules in embedded PCs were used.

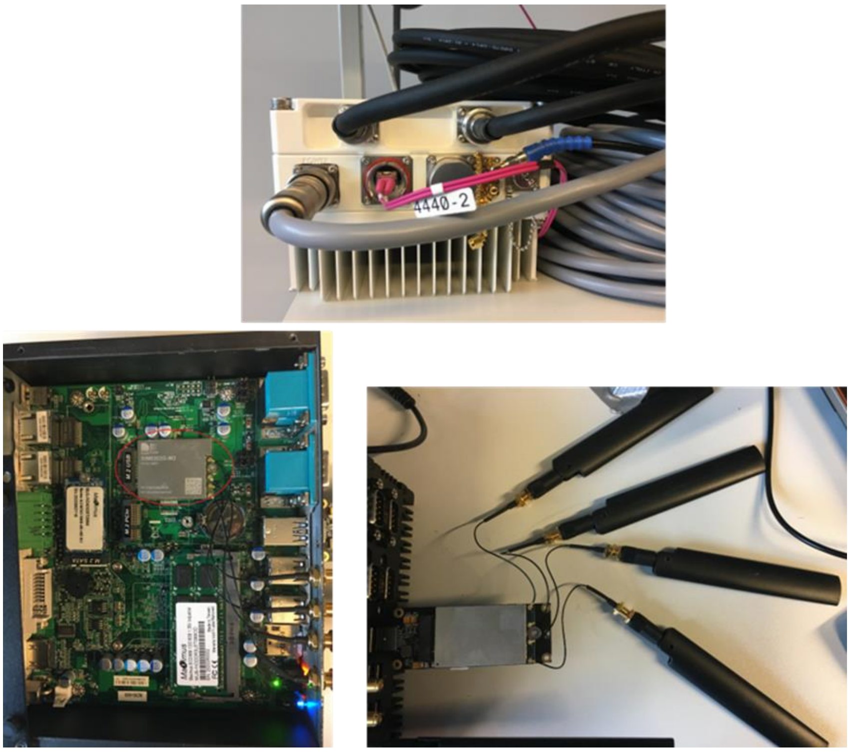

Figure 2 - AW2S Jaguar (2x2, 20 MHz) RU (top), SIMCOM (left) and Quectel (right) COTS UE modules used during first stage trials.

2.1.1 Updates on measurements planning

2.1.1.1 First stage trials: Monolithic RAN setup with single COTS UE

As stated in the previous section, for the first stage trials the OAI monolithic setup was used. Another update with respect to what is described in D3.1 [4], is the use of Rel. 15 COTS UE modules from SIMCOM and Quectel instead of a prototype UE based on OAI software and ETTUS USRP B210 RUs.

The first stage trials in UC1 had two main goals. First, to test and verify concepts, approaches, and separate components in a less complex architecture setup. Secondly, to identify needs for further optimization where application requirements are not yet met.

The reduced features of the first stage compared to the final setup are:

- The focus of the first stage was to reduce the latency and packet error ratio for wireless audio transmissions. Features to support time distribution over the 5GS will be considered further into the project. Therefore, time distribution and synchronization between audio devices is achieved in the first stage with a **dedicated wired connection** between time service and media devices (Figure 1).
- The first stage **deployed only a single 5G modem** to reduce the complexity of scheduling and slicing in the 5GS to test basic concepts of reducing the transmission latency. The audio platform attached to the modem can be configured as a microphone to test UL, as IEM to test DL, or both to test low latency UL + DL at the same time.

- During the first stage of the trials, the parameters and configuration methodologies for the 5GS were still in an optimization phase. Because of that, the **automatic procedures between network slicing manager and core configuration service were not yet available**. Therefore, the 5GS had to be configured and adjusted manually.
- Due to the separate and independent implementation of the Media Orchestration Control Gateway in this project and to Covid-related delays, the **MOCG was not available** for the first phase trial. The Media Orchestration Gateway had to be set up manually to control audio devices and setup media flows.
- Since in this stage there was **no shared access client in the RAN components**, the spectrum and power to use in the RAN were manually configured. However, there will be a demonstration during the first trial focusing on the interface to use between user and the shared access server for shared licensing purposes.

The following paragraphs will provide a summary of the different test cases as well as their updates with regards to D5.1 [4] and the corresponding configuration parameters.

Trial UC1.A.1

Initial measurement of pre-existing monolithic testbed implementation based on Open Air Interface as a reference point.

Table 1 - Parameters Trial UC1.A.1.

Parameter	Value
5GS (RAN + Core)	Monolithic OAI
Time sync between application and 5GS	No
UE / modem	1x COTS {1}
Audio streams DL	Local audio processing → 1x IEM (connected to {1})
Audio streams UL	1x Mic (connected to {1}) → Local audio processing
Slot duration/SCS	0.5 ms/30 kHz
Periodicity of the DL-UL pattern	5 ms
Audio network packet periodicity	1 ms
Target KPIs	5G network latency, packet error ratio

Trial UC1.A.2

Comparison of a synchronized and a not-synchronized interface between application and 5GS to evaluate potential implications.

Table 2 - Parameters Trial UC1.A.2.

Parameter	Value
5GS (RAN + Core)	Monolithic OAI
Time sync between application and 5GS	No / Yes
UE / modem	1x COTS {1}
Audio streams DL	Local audio processing → 1x IEM (connected to {1})
Audio streams UL	1x Mic (connected to {1}) → Local audio processing
Slot duration/SCS	0.5 ms/30 kHz
Periodicity of the DL-UL pattern	5 ms
Audio network packet periodicity	0.5 ms / 1 ms / 2.5 ms / 5 ms
Target KPIs	5G network latency, packet error ratio

Trial UC1.A.3

Optimization of DL/UL pattern periodicity to reduce latency, identification of additional bottlenecks. This testcase also implies time sync between application and 5GS for both 5 ms and 2.5 ms of DL-UL pattern periodicity.

Table 3 - Parameters Trial UC1.A.3.

Parameter	Value
5GS (RAN + Core)	Monolithic OAI
Time sync between application and 5GS	Yes
UE / modem	1x COTS {1}
Audio streams DL	Local audio processing → 1x IEM (connected to {1})
Audio streams UL	1x Mic (connected to {1}) → Local audio processing
Slot duration/SCS	0.5 ms/30 kHz
Periodicity of the DL-UL pattern	5 ms and 2.5 ms
Audio network packet periodicity	0.5 ms / 1 ms / 2.5 ms / 5 ms
Target KPIs	5G network latency, packet error ratio

Trial UC1.A.4

Extensive latency optimizations with further reduced DL/UL pattern periodicity and increased sub-carrier spacing. **Comparison of the monolithic OAI setup and the disaggregated 5GS including all partner components.**

Table 4 - Parameters Trial UC1.A.4.

Parameter	Value
5GS (RAN + Core)	Monolithic OAI / Disaggregated
Time sync between application and 5GS	Yes
UE / modem	1x COTS {1}
Audio streams DL	Local audio processing → 1x IEM (connected to {1})
Audio streams UL	1x Mic (connected to {1}) → Local audio processing
Slot duration/SCS	0.25 ms/60kHz
Periodicity of the DL-UL pattern	5 ms, 2.5 ms, 1.25 ms, 1 ms
Audio network packet periodicity	0.5 ms / 1 ms / 2.5 ms / 5 ms
Target KPIs	5G network latency, packet error ratio

2.1.1.2 Final stage trials: Disaggregated RAN setup with single COTS UE

The final trial stage expands, based on the first stage trials results, towards the fully featured use case setup. The focus of these trials will be the enhancement of latency, scalability with multiple audio devices, time synchronization over 5G, full remote controllability and dynamic spectrum assignment.

The architecture of the end-to-end setup supporting the final stage trials is shown in Figure 3. For a more detailed description of the testbed architecture please refer to D3.1 [2]).

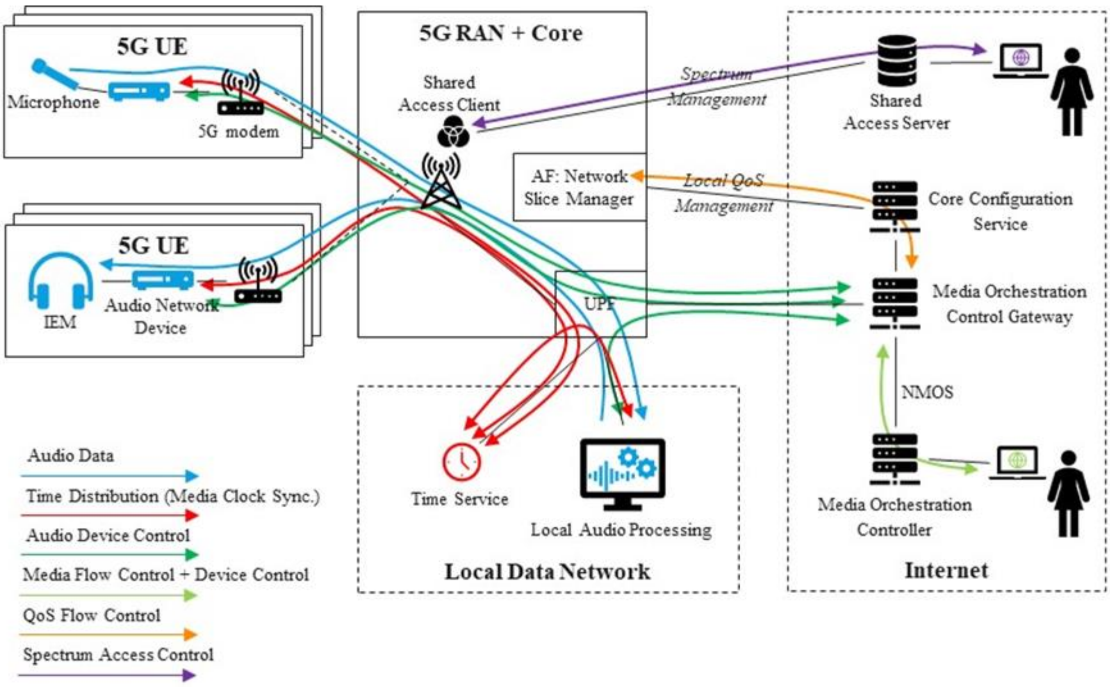

Figure 3 - UC1 Testbed architecture - Final trial stage

The following paragraphs will provide a summary of the different test cases envisaged for the final trial stage, as well as their updates with regards to D5.1 [4] and the corresponding configuration parameters.

Trial UC1.B.1

Introduction of a second UE to validate the concept for multi-device low-latency audio transmissions.

Table 5 - Parameters Trial UC1.B.1.

Parameter	Value
5GS (RAN + Core)	Disaggregated
Time sync between application and 5GS	Yes
UE / modem	2x COTS {1,2}
Audio streams DL	Local audio processing → 2x IEM (connected to {1,2})
Audio streams UL	2x Mic (connected to {1,2}) → Local audio processing
Slot duration/SCS	0.25 ms/60kHz
Periodicity of the DL-UL pattern	1 ms
Audio network packet periodicity	1 ms
Target KPIs	5G network latency, packet error ratio

Trial UC1.B.2

Further increase of number of UEs to evaluate scaling of multi-device low-latency audio setups.

Table 6 - Parameters Trial UC1.B.2.

Parameter	Value
5GS (RAN + Core)	Disaggregated
Time sync between application and 5GS	Yes
UE / modem	8x COTS {1,2,3,4,5,6,7,8}
Audio streams DL	Local audio processing → 4x IEM (connected to {1,2,3,4})
Audio streams UL	4x Mic (connected to {5,6,7,8}) → Local audio processing
Slot duration/SCS	0.25 ms/60kHz
Periodicity of the DL-UL pattern	1 ms
Audio network packet periodicity	1 ms
Target KPIs	5G network latency, packet error ratio

Trial UC1.B.3

Comparison of delivery of synchronized time information via dedicated Ethernet cable, over-the-top and with TSN-based mechanisms.

Table 7 - Parameters Trial UC1.B.3.

Parameter	Value
5GS (RAN + Core)	Disaggregated
Time sync between application and 5GS	Yes
UE / modem	2x COTS{1,2}
Audio streams DL	Local audio processing → 2x IEM (connected to {1,2})
Audio streams UL	2x Mic (connected to {1,2}) → Local audio processing
Slot duration/SCS	0.25 ms/60kHz
Periodicity of the DL-UL pattern	1 ms
Audio network packet periodicity	1 ms
Media Clock Sync	Ethernet / Rel 15 COTS / Rel 16 COTS
Target KPIs	Synchronicity

2.1.2 Uncertainties and risk assessment

As described in section 2.1, the validation of the end-to-end interoperability using the disaggregated RAN setup between the Accelleran CU and the OAI DU from EURECOM required more effort than initially planned and, as a result, did not allow the end-to-end first stage trials to be performed using the disaggregated RAN setup. Basic control and user plane interoperability has now been validated between the two components. However, further validation with respect to user-plane traffic tests using the Sennheiser low latency application are required, before we can move to KPI measurements collection. Once this is completed, it is planned to perform the next stage of measurements with the disaggregated RAN setup and compare its performance to the monolithic one.

Despite the initial issues, the improvements that have been made in the monolithic OAI RAN, both in terms of stability and KPI performance towards the final performance targets, will be reflected at the OAI DU as well, since they address common extensions of the OAI lower RAN layers software and the interoperability with the AW2S RU.

Furthermore, the interoperability between the Accelleran CU and the Cumucore core network has been separately performed and validated between the two parties. Similarly, the validation between Accelleran shared spectrum access client and RED Technologies shared spectrum access server has also been separately validated. As a result, it is not expected that the integration of these components will cause any significant problems/delays for the integration and testing at EURECOM site.

2.2 Measurement results

At the beginning of 2022, a first series of measurements related to Trials UC1.A.1, UC1.A.2 and UC1.A.3 were conducted (see Table 1, Table 2, Table 3). This first set of tests is part of the first trial stage, with a focus on 5G network latency and packet error ratio (also see above trial stage description). The following sections will show exemplary measurements that are representative of the results obtained.

The first set of trials showed that the 5G testbed currently presents an unexpected amount and variety of unpredictable packet loss (up to 1%), especially in Downlink direction. It is unclear where the packet loss is occurring at this point in time and is part of ongoing investigation. It is assumed that the cause of the unreliability is to be found in the gNB and/or the Linux drivers of the COTS UE modules towards the IP/application layers. In the current state, no meaningful statements about the reliability in the context of use case 1 are possible. Thus, packet loss and reliability are not considered in the following first analysis of measurement results.

2.2.1 Trial UC1.A.1. (Initial measurements)

The goal of this first series of measurements was to show the trade-off between packet errors and latency when employing packet loss recovery methods such as Radio Link Layer (RLC) Acknowledgements and the MAC Layers Hybrid Automatic Repeat Request (HARQ). As mentioned before, the 5G testbed showed an unexpected packet loss behaviour, leading to non-reproducible measurement results for this analysis.

Figure 4 and Figure 5 show the 5G end-to-end latency of each audio packet sent through the 5GS over the course of about 30 minutes, measured from Ethernet transmission to reception in the Local Audio Processing and the Audio Network device. The measurements were taken in RLC unacknowledged mode (UM) and without HARQ.

Audio packets were sent with a periodicity of 1 ms, while the 5GS was able to provide a transmission periodicity pattern (TDD) of 5 ms. The pattern consisted of 10 repeating slots with 500 µs length each. 7 out of 10 slots were used by the scheduler for Downlink, one slot could be used for Downlink or Uplink, and two slots were exclusively for Uplink (DDDDDDDxUU). While all audio components were synchronized to the same time source with a dedicated Ethernet cable (see Figure 2), the time basis of the 5GS was a different one. This is also evident from the measurements in Figure 4 and Figure 5. Both graphs show typical slopes of the E2E latency caused by communication systems with unrelated and drifting clocks for their respective timing grid. Secondly, the graphs show multiple distinct slops. This is caused by the fact that the 5GS and the audio system used different periodicities for the transmission scheduling and packet sending during the measurement.

In Downlink the E2E latencies range up to 11.55 ms, while Uplink E2E latencies range up to 11.56 ms.

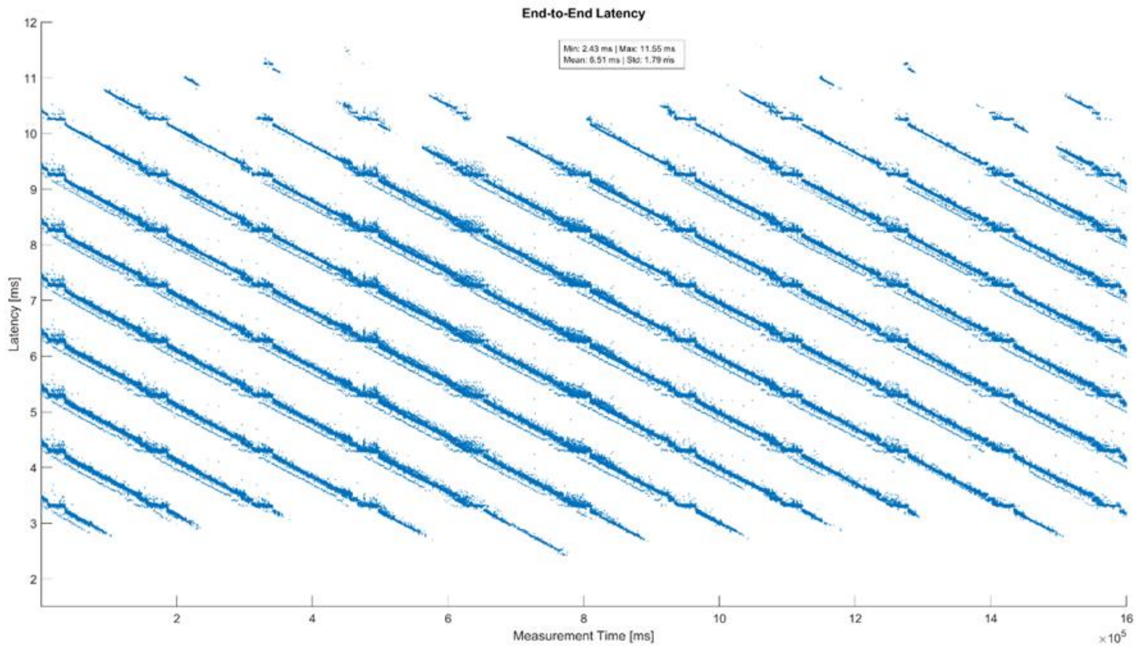

Figure 4 - E2E Latency, 5ms TDD, 1ms audio periodicity, no audio/5GS sync, Downlink

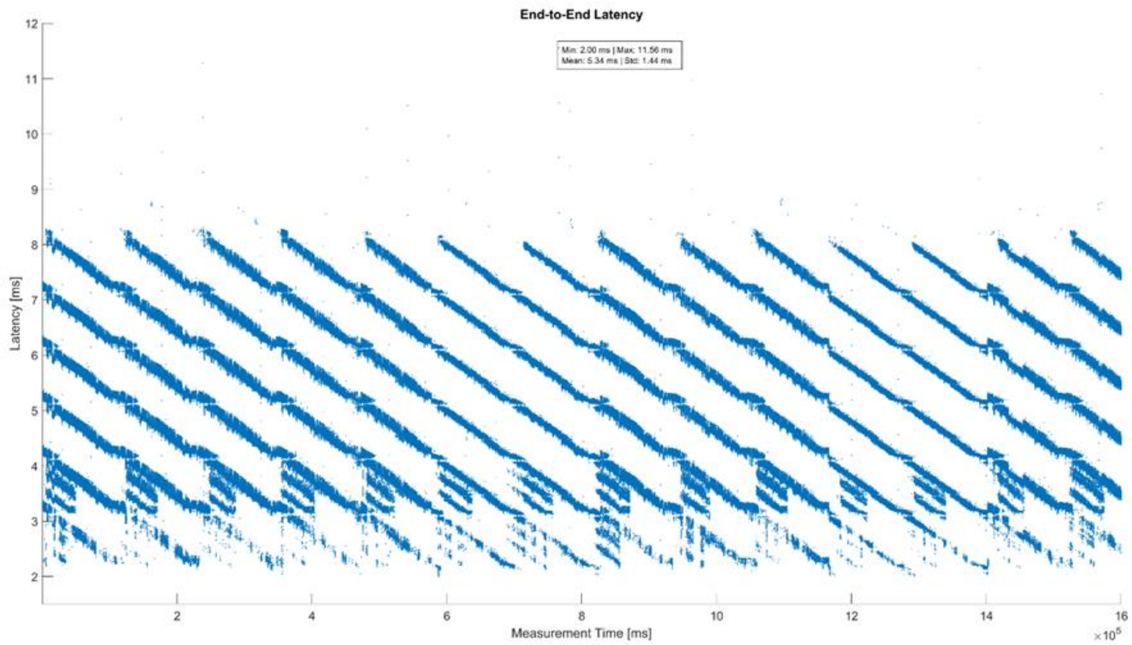

Figure 5 - E2E Latency, 5ms TDD, 1ms audio periodicity, no audio/5GS sync, Uplink

2.2.2 Trial UC1.A.2. (Sync between audio and 5G system)

This set of trials was designed to show the benefit of the availability of interfaces allowing clock synchronization of all parts of an end-to-end communication system. For a proof of concept, the synchronization between 5GS and audio components was achieved by connecting a GPS antenna to both separated subsystems and thus using GPS time as the same time source. Given that a future 5G systems would be able to distribute time information with sufficient precision over the air via PTP protocol, this could be used to synchronize the overall application.

Figure 7 and Figure 6 show the E2E latency of every audio packet that was sent through the 5GS over the course of about 30 minutes, with GPS synchronization.

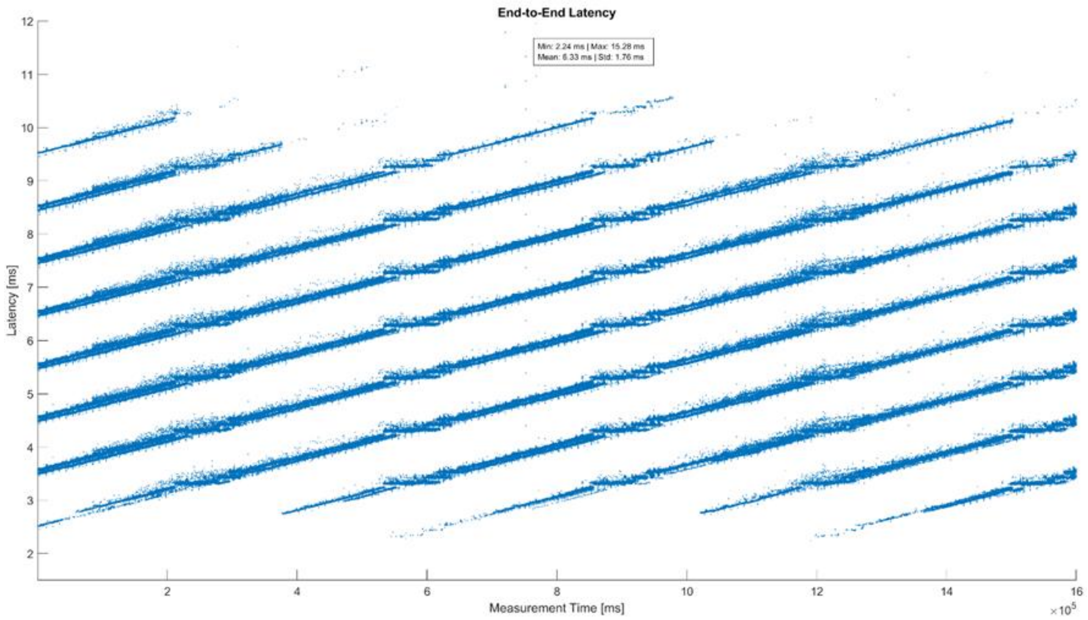

Figure 7 - E2E Latency, 5ms TDD, 1ms audio periodicity, GPS sync, Downlink

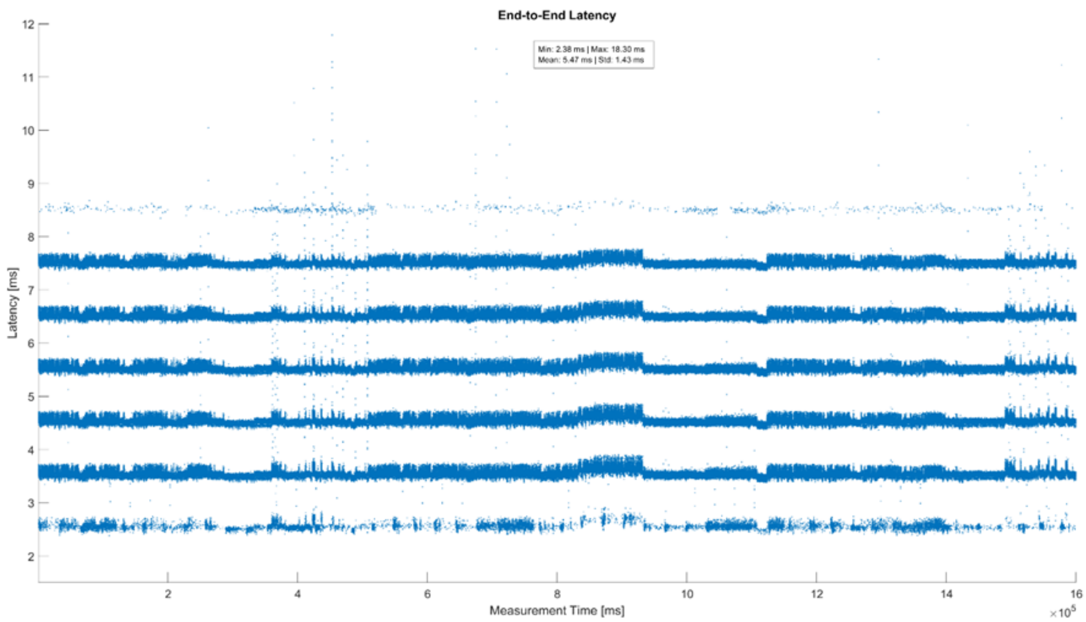

Figure 6 - E2E Latency, 5ms TDD, 1ms audio periodicity, GPS sync, Uplink

In Downlink direction (Figure 7) it can be seen that slopes are still present, while in Uplink (Figure 6) the slopes are no longer there. Here Uplink shows the expected behaviour when all components are in sync. Furthermore, this implies that in Downlink direction at least one component of the end-to-end signal path is not in sync to the GPS time source. At this point in time, it is unclear which component is behaving unexpected, one assumption is that a part of the 5GS UPF is not processing in sync.

In Downlink the E2E latencies range up to 15.28 ms, while Uplink latencies range up to 18.30 ms.

2.2.3 Trial UC1.A.3. (Reduction of transmission pattern periodicity)

One of the goals of UC1 is to optimize the E2E latency of a 5G testbed and to evaluate if 5G technology is able to meet the latency requirements of a URLLC scenarios such as professional live audio productions. For that reason, a set of major optimizations were implemented in the gNB to reduce the TDD pattern to a period of 2.5 ms with 5 slots of 500 µs each. Here, two slots are exclusively used by the scheduler for Downlink and Uplink, one slot remains flexible (DDxUU).

As mentioned before, Downlink direction of the current testbed is still being optimized to enable synchronized transmissions. Therefore, we only considered Uplink direction for the following measurements.

Figure 8 shows the E2E latencies of each audio packet sent with a periodicity of 1 ms through the TDD optimized 5G tested in Uplink direction. In a direct comparison with a TDD pattern of 5 ms (Figure 6) it is evident that the majority of packets benefit from the reduced pattern. The potential penalty of additional wait time when a packet missed a slot for transmission and having to wait for the next TDD period is directly related to the pattern periodicity and, thus, reduced for shorter periods.

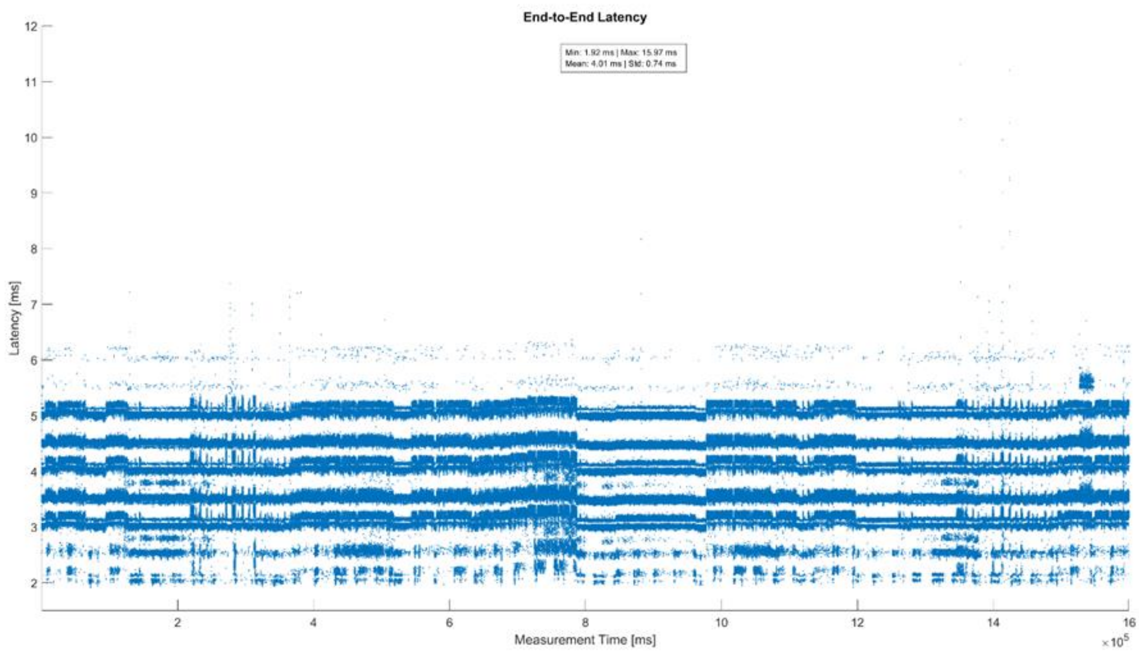

Figure 8 - E2E Latency, 2.5ms TDD, 1ms audio periodicity, GPS sync, Uplink

Observed E2E latencies range up to 15.97 ms in this measurement.

Figure 9 shows an E2E latency measurement with the same periodicity of audio transmissions as the TDD pattern. As can be seen in the graph, the multiple lines of latencies related to difference in periodicity patterns, are here reduced to one. The position of the remaining line depends on phase relation between 5GS slot timing grid and audio transmission periodicity grid, which is currently not optimized. This shows that, in general, it can be beneficial to align the periodicities in all involved communication components.

It should be noted here that increasing the periodicity in audio components leads to additional latency of the same size for collecting audio data. By increasing the audio periodicity in this trial, it was possible to reduce the E2E transmission latency of most packets, while also adding 2.5 ms of latency to every packet.

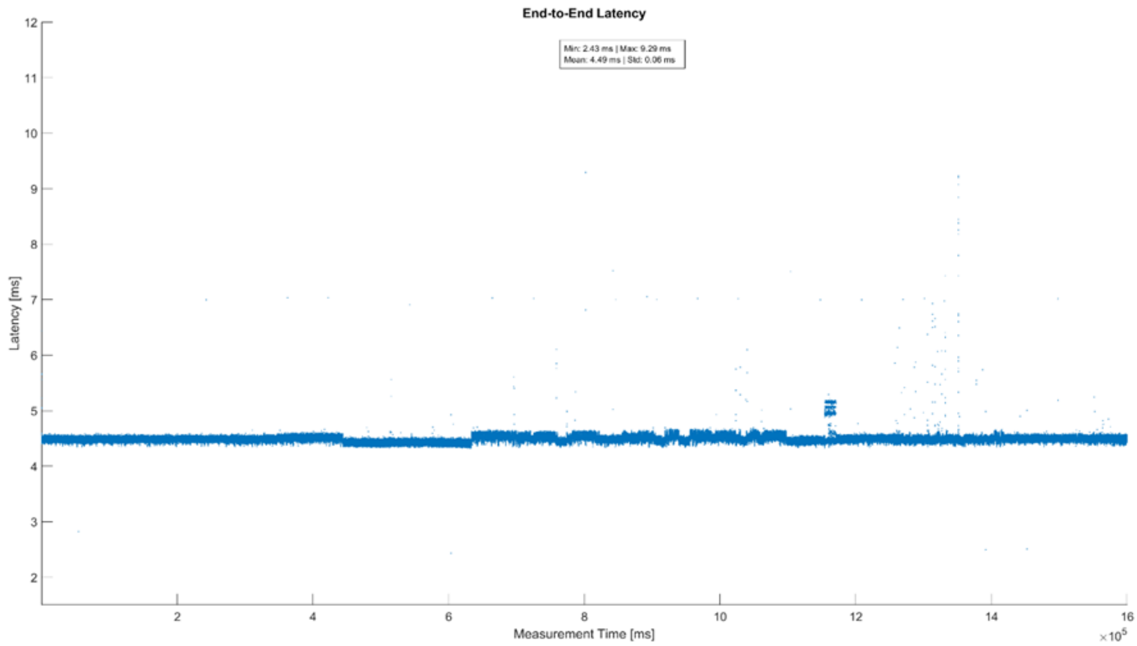

Figure 9 - E2E Latency, 2.5ms TDD, 2.5ms audio periodicity, GPS sync, Uplink

This trial shows that the TDD pattern in a transmission system is of major importance for the overall application latency.

E2E latencies in this measurement range up to 9.29 ms. Figure 10 shows the latencies in a CDF graph.

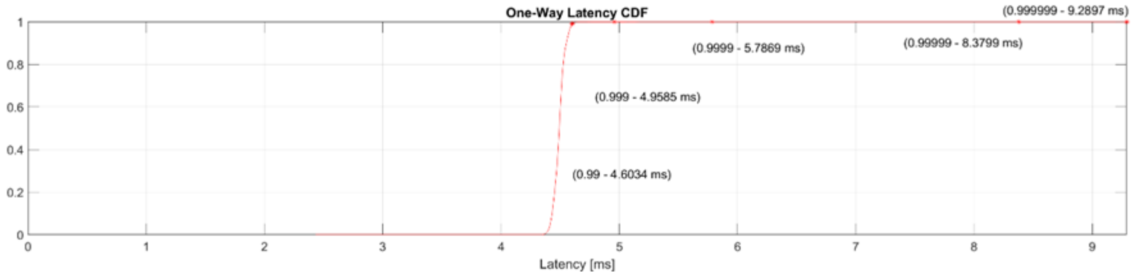

Figure 10 - E2E Latency CDF, 2.5ms TDD, 2.5ms audio periodicity, GPS sync, Uplink

2.3 KPI analysis

2.3.1 Network latency

Requirement: network latency < 1 ms.

The network latency consists of E2E latency and transfer interval. The E2E latency was measured from Ethernet transmission to reception in the Local Audio Processing and the Audio Network devices. To obtain the network latency it is necessary to add the respective audio periodicity / transfer interval.

Although it was possible to optimize some components of the 5GS, our first measurements have shown that the overall E2E latency is still considered to be in the 10's of milli seconds. While the majority of packets greatly benefit from the reduced TDD pattern, it is also evident that still many packets are outlining in the current testbed. Packets that are too slow for a potentially deployed jitter buffer have to be counted lost. When considering a required packet error ratio of better than 10^{-6}, the operating point so far had to be chosen between 10 ms and 12 ms E2E latency (e.g., Figure 10), sometimes higher.

Furthermore, we recognized that at least one 5G component in Downlink direction is not working in sync to the selected time basis. This currently hinders the optimization of the Downlink E2E latency. Next steps to continue the latency optimization are:

- To synchronize all E2E 5G Downlink components
- Further optimizations required to reduce overall number of E2E latency outliers
- Further reduction of TDD pattern or slot length to potentially meet millisecond E2E latency.

2.3.2 Synchronicity

Requirement: Synchronicity < 500 ns.

This KPI was not considered in the trials yet.

2.3.3 Packet error ratio

Requirement: Packet error ratio < 10^{-6}.

As mentioned above, we have not yet considered the packet error ratio in use case 1 since the current testbed shows some unexpected behaviour with respect to lost packets.

In-depth analysis and optimization are needed to identify the causes of packet loss in the 5GS.

2.4 Deployment considerations for final stage

The aim for the final stage trials is to put in place the configurations indicated in D3.1 [4] that can lead to the target radio latency reduction required for UC1. More specifically, either the combination of 1.0 ms DL-UL pattern periodicity and 30 KHz SCS/0.5 ms slot duration, or 1.25 ms DL-UL pattern periodicity and 60 KHz SCS/ 0.25 ms slot duration will be used to reach the target of ~1 ms radio latency at the OAI DU.

Initially, it was considered to test this configuration with Rel. 16 COTS UEs that are more likely to support such a low latency target comparing to Rel.15 devices. However, given the possible unavailability of Rel. 16 COTS UEs in the market until late 2022, we aim to start testing the aforementioned configurations with our Rel. 15 UE modules first.

3 Use Case 2: Multiple Camera Wireless Studio

This chapter describes the outcomes of the first stage of trials in the context of UC2. Testbed architecture and measurements' results will be presented first and then followed by a detailed KPI analysis (related to the KPIs defined in D2.1), as well as the deployment consideration, useful for the final stage of the trials.

3.1 Deployed testbed architecture

The first stage of trials of UC2 focused on the integration of both the available media components (phase 1) and the newly developed features for the UC components (phase 2). The following sub-sections discuss the architecture and components of the two user scenarios examined in the use case: the wireless studio and remote contribution.

Wireless studio scenario:

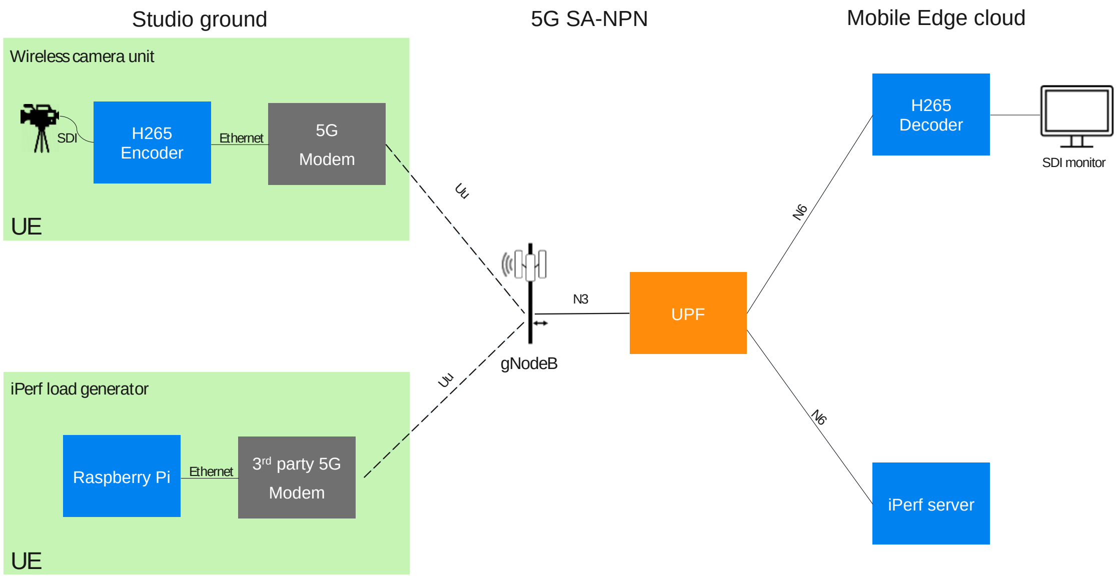

Figure 11 - Wireless studio deployed system architecture

In the wireless studio scenario, the plan for phase 2 was to integrate the 5G modem, the H.265 encoder and the Media Gateway into the 5G NPN test network and perform preliminary tests. However, due to difficulties in the integration between the H.265 encoder and the Media Gateway, it was decided to temporarily replace the Media Gateway with another H.265 decoder based on the Xilinx board used in the encoder.

Figure 11 depicts the architecture used during phase 2. This architecture will be the base configuration for the architecture targeted for the final stage and is divided into three different tiles:

1. The first (left) tile is called "studio ground" and consists of two UEs: one which resembles an on-air camera unit with an SDI camera, H265 encoder, and a 5G modem and another "synthetic" UE which generates extra traffic that competes with the on-air camera on the radio resources.
2. The second (center) tile represents the 5G NPN components, which is composed by the RAN represented by the gNodeB and the CN represented by the UPF.
3. The third (right) tile called "Mobile Edge Cloud" consists of the servers and the H265 decoder, those are used to interact with the devices in the studio ground. This section was used during phase 2 to validate the different components and test the performance under various network loads.

The architecture described in Figure 11 will be slightly modified to include the Media Gateway in later stages of the project, as this component was not available at this stage.

Figure 12 shows the main media components used during the second phase of testing in Ericsson's lab.

Figure 12 - The setup in the 5G lab

Remote contribution scenario:

The basic setup for the second scenario, remote contribution, followed the same workflow of Phase 1: from Ericsson 5G lab in Aachen to the RAI Turin media lab. The Aachen network was 5G SA NPN, single node in lab, with DDSU TDD configuration.

To enrich this setup and the related tests, a LiveU LU800 5G Pro was shipped to Denmark and placed in TV2 studio where was set up to transmit video streams from there to RAI media lab using the Danish commercial 5G network. Therefore, the deployed architecture and the tests covered 3 sites, 2 different 5G networks and multiple scenarios.

Furthermore, in addition to the components integrated and tested in phase 1, phase 2 also tested the integration with:

a) LiveU 800Pro multi-cam capabilities (supporting encoding and transmitting feeds from up to 4 cameras simultaneously, splitting the available UL BW between them);
b) LiveU intercom and IFB audio servers, to test studio-field production communications over the 5G network and equipment and
c) Camera-control devices performing shading and iris adjustments from remote, using the LiveU bi-directional IP PIPE between the studio and field unit and the camera connected to it. The control equipment for this part was provided by Cyanview.

Some of the transmission tests were also conducted loading the Aachen 5G lab network with traffic emulation sourced by other UEs (i.e.: Raspberry Pi + 5G modem), connected to the network via the 5G RAN.

The following diagram depicts all the configuration variations put together.

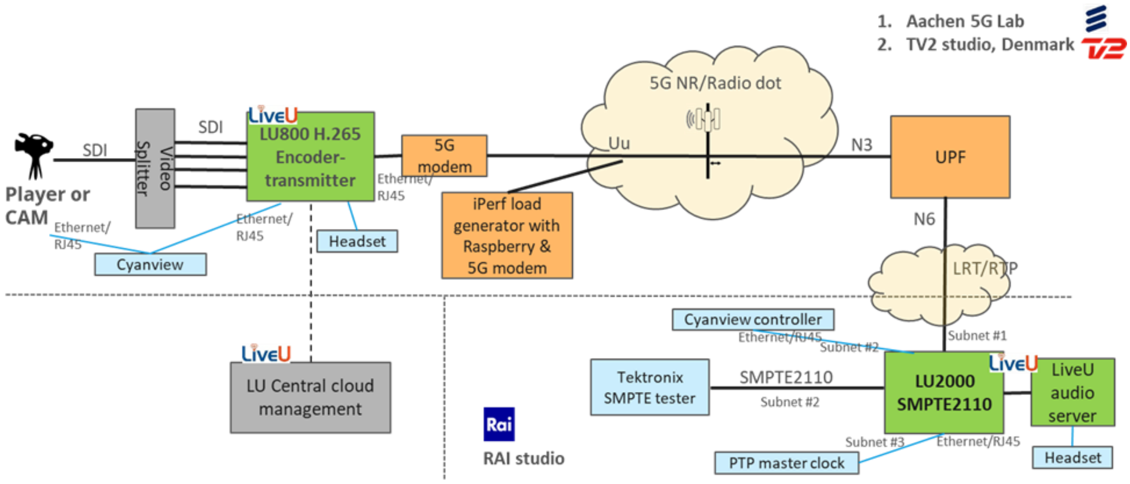

Figure 13 - Remote contribution deployed system architecture

Figure 14 and Figure 15, taken from the RAI media lab, show the LiveU LU2000SMPTE HEVC decoder-receiver, LiveU audio device, screens showing the multi-cam transmissions received from Aachen via the LiveU units, Tektronix Prism testing equipment, and other media equipment.

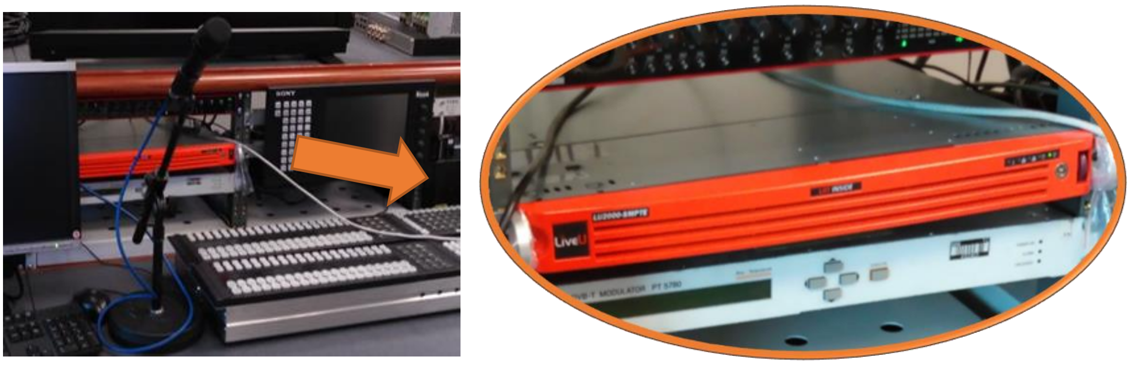

Figure 14 - LiveU LU2000SMPTE HEVC decoder-receiver

Figure 15 - Multi-cam screens, Tektronix Pirsm and other media equipment

3.1.1 Updates on measurements planning

Wireless studio scenario

In the wireless studio scenario, the first stage of trials was sub-divided in two sub-stages. An early-stage trial was focused on getting a hands-on experience with the newly installed 5G lab. The target was to run tests which could assess the network capabilities using well-known network performance metrics such as: throughput, latency, and packet loss.

- **Trial UC2.A.1:** Estimation of the network bitrate and packet loss
- **Trial UC2.A.2:** Estimation of the network uplink bandwidth

During the second sub-stage, the video encoder/decoder was available for testing and G2G latency over 5G network was successfully measured together with the remote control of the camera.

- **Trial UC2.A.3**: Glass-to-glass latency measurement
- **Trial UC2.A.4:** Camera remote control

For the final stage trials, no changes in the plan discussed during D5.1 [4] is expected.

Remote contribution scenario

Regarding the remote contribution scenario, the tests planned in D5.1 [4] were updated as follow:

- **Trial UC2.A.5:** Basic transmission of a single HD video feed using a single LU800 device per test source location (multi-encoder, splitting the UL BW available between them)
- **Trial UC2.A.6:** Multi-cam transmission of 4 HD video feeds using a single LU800 device per test source location (multi-encoder, splitting the UL BW available between them)

- **Trial UC2.A.7:** SMPTE ST 2110 compliancy of LU2000SMPTE HD video output, according to the JT-NM ST2110 compliancy tests [8]

In addition to the previously planned tests, other test cases were performed:

- **Trial UC2.A.8:** Transmission tests under traffic load emulation from Aachen 5G lab at 50%, 75% and 90% of its UL capacity
- **Trial UC2.A.9:** Audio intercom and IFB from RAI media lab in Turin to/from Aachen 5G lab
- **Trial UC2.A.10:** Remote-control of camera shading using Cyanview integrated with LiveU IP PIPE

It is worth mentioning that many of the transmission benchmark tests, audio transmission tests and remote-control tests were also performed using Danish 5G NSA commercial network of TDC.

3.1.2 Uncertainties and risk assessment

This section contains a description of the issues encountered during the initial deployments of the test-beds and their possible impacts on the project's deadlines.

In the initial planning for the multi-camera wireless studio scenario, the end-to-end chain was meant to include:

- Encoder(s)/Decoder(s) [ImageMatters]
- Modem(s) [Fivecomm]
- Media Gateway (RTP<->ST2110) [Bisect]
- Media Orchestration Control Gateway [Bisect/BBC]
- 5G SA network [Ericsson]

The ImageMatters encoders used had to be capable to provide:

- Ultralow latency encoding profiles
- Genlock capability
- RTP timestamping

During the first test session, none of the components listed above was available, except for the 5G SA network. Therefore, cameras capable to compress the video signals and a 3rd party modem was used, along with a simplified setup.

This situation slightly improved during the 2nd test session, when only the encoder (with limited features, only ultralow latency encoding), the modem and the 5G SA were available.

A 3rd round of tests will take place in Aachen in April 2022. At this point in time, the codec (with limited features), the modem and the media gateway (RTP<->ST2110) will be integrated and tested, return video included.

A 4th round of tests is planned in Q3 2022 with the final codec board, if available.

Meanwhile, the basic functionalities of the MOCG (input sources connected to the MG and exposing the sender/receiver) will be demonstrated at the end of April 2022 and in Q3 2022 with the final board, if available.

3.2 Measurement results

3.2.1 Trial UC2.A.1 (wireless studio)

Table 8. UC2.A.1 measurement parameters

Measurement description	Result	Comments
Estimation of the network bitrate and packet loss using SCREAM	Passed	An estimate of 100 Mbps is measured on the uplink with packet loss detected when SCREAM attempts to go above 100 Mbps

Figure 16 depicts a snapshot of the measurement of the uplink capacity using SCREAM. The tool can be configured to start at a specific bitrate and try to keep the bitrate as close as possible to such bitrate. In the graph given below the minimum rate is set to 100 Mbps. The tool starts at 100 Mbps and attempts to move to a higher bitrate until the packet loss is very high, the client estimates the network bandwidth to be below 100 Mbps.

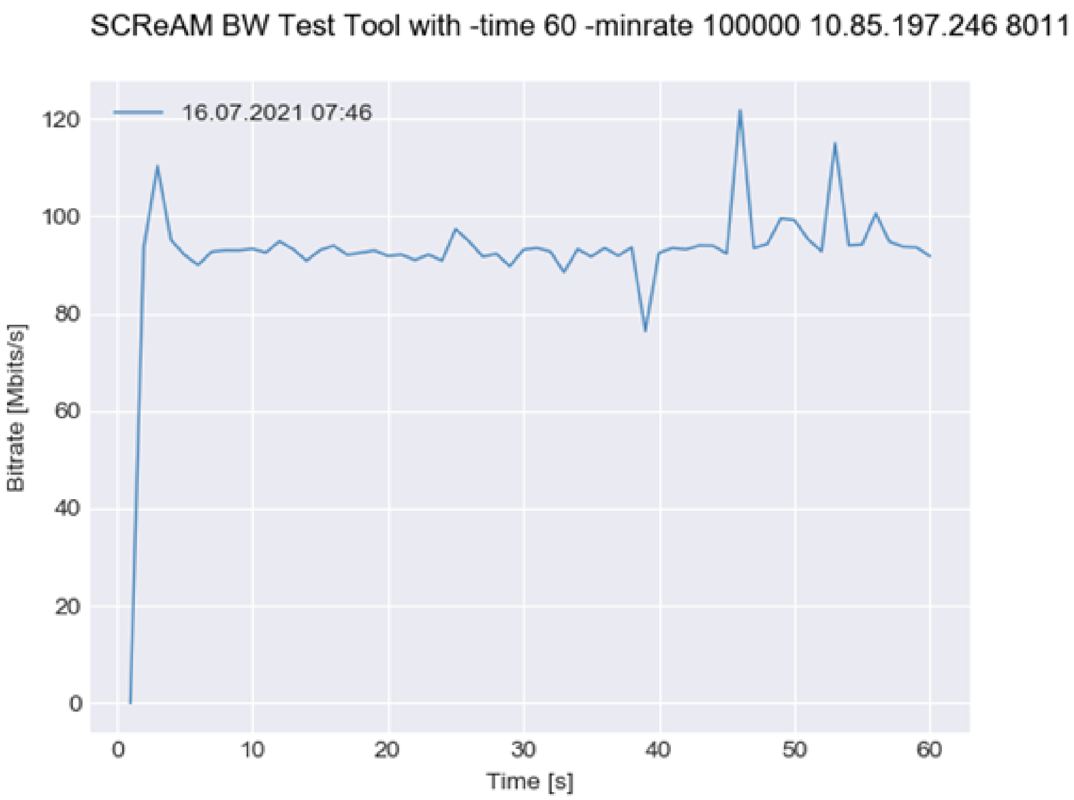

Figure 16 - Bandwidth measurement using SCREAM

The packet loss signaled from the receiver to the SCREAM client during the same measurement session is shown in Figure 17**.** The increase in the packet loss rate can be correlated to the attempts of the congestion algorithm used by SCREAM to increase the bitrate above the configured 100 Mbps. The spikes in bitrate are faced by a spike in packet loss.

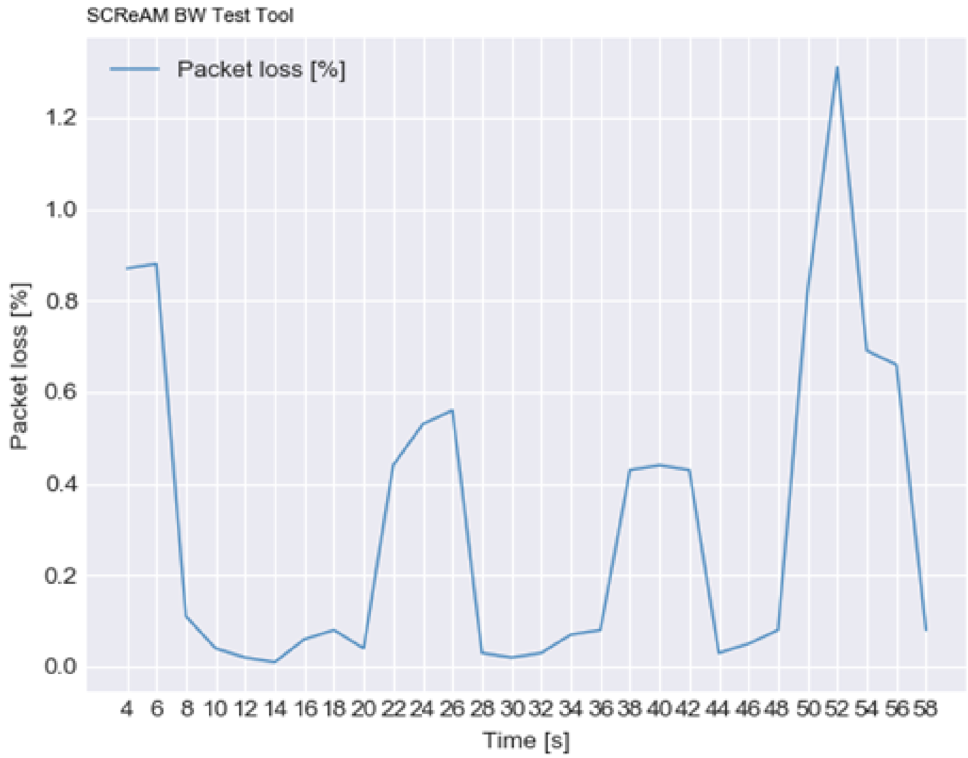

Figure 17 - Packet loss with minimum rate configured to 100 Mbps

3.2.2 Trial UC2.A.2 (wireless studio)

Table 9. UC2.A.2 measurement parameters

Measurement description	Result	Comments
Estimation of the network uplink bandwidth using iPerf	Passed	The UDP iPerf stream resulted in 100 Mbps uplink

The target for this measurement was to repeat the throughput measurement using a generic tool such as iPerf. The tool can be used in both UDP and TCP modes. The UDP mode enables a client to send a specific bitrate through the uplink without any indication if the packets are received by the server, while the TCP mode depends on a congestion algorithm to adjust its throughput.

In the test, the UDP client was configured to send 100 Mbps to the edge server. The bandwidth is calculated by observing the bitrate arriving at the edge server. This method indicates the maximum network throughput, but it is prone to packet loss.

Figure 18 depicts the maximum throughput observed by the edge server given that the client sends consistently 100 Mbps on the uplink.

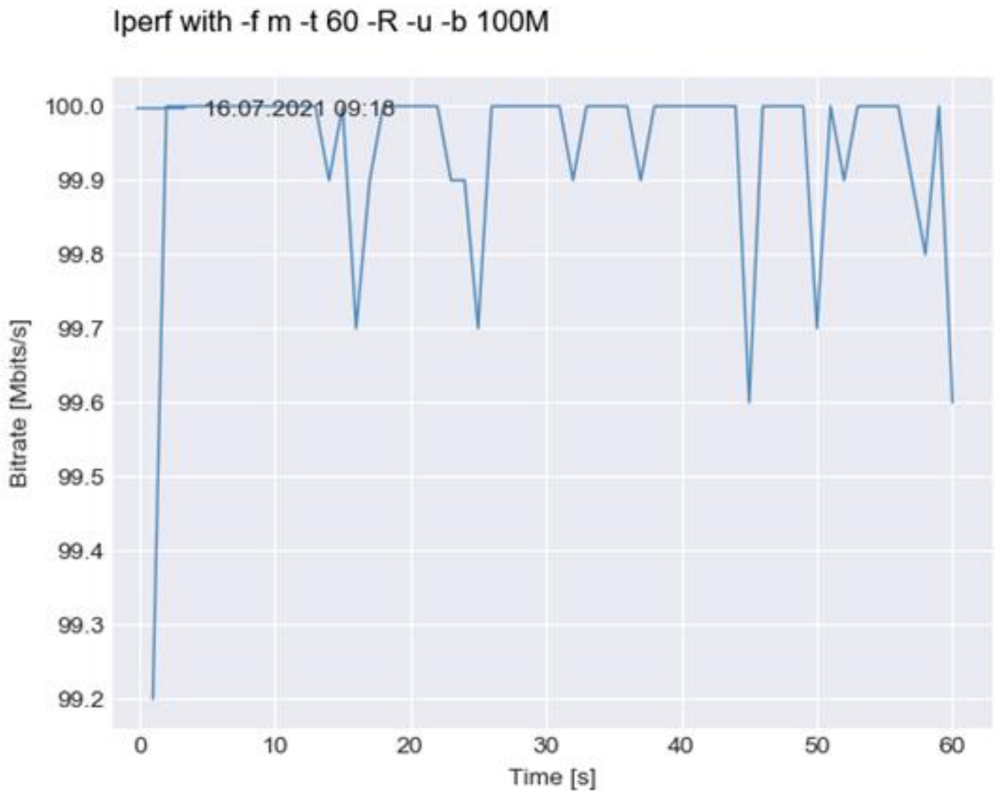

Figure 18 - Uplink throughput estimates using iPerf in UDP mode

3.2.3 Trial UC2.A.3 (wireless studio)

Table 10. UC2.A.3 measurement parameters

Measurement description	Result	Comments
G2G latency measurement using an oscilloscope	Partially met	G2G= 57 msbut the stream wasn't stable

During phase 1, accurate measurement for G2G latency was critical to evaluate whether the system was able to achieve the strict latency requirements described in D2.1.

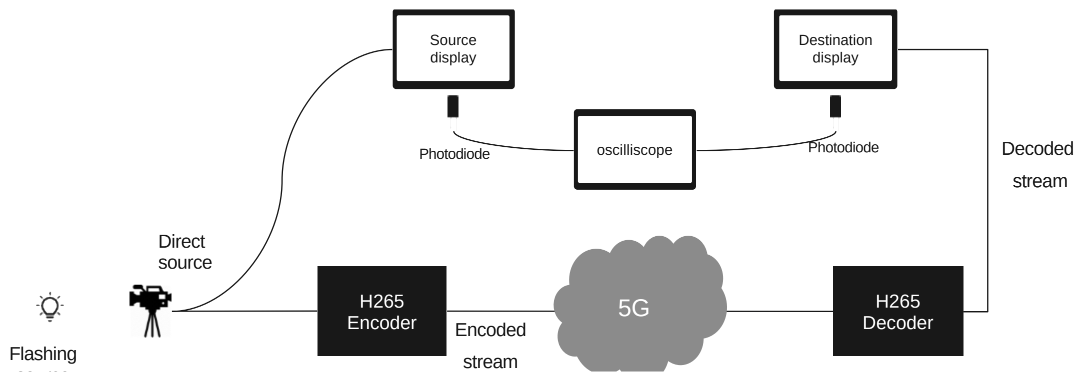

Figure 19 - Glass to glass latency measurement setup

The measurements setup, represented in Figure 19, was based on two photodiodes attached to the source and destination displays. The camera captures a periodically flashing light that is detected by both photodiodes. The diodes then send the electrical signals to an oscilloscope. The G2G latency was determined by measuring the time shift between the signal captured from the source and the signal captured from the destination photodiodes. The measurements were performed using the Image Matters encoder and decoder boards.

Figure 20 depicts a screenshot from the oscillator screen. It can be observed that the difference between the start of the signal pulse train, captured by source and destination photodiode signals, is 57.60ms. The low latency values were achieved by omitting the de-jitter buffer at the receiver side and using sliced-based encoding at the encoder side. However, the receiver crashed when packets didn't arrive on-time, therefore, for better stability of the stream, a de-jitter buffer can be added, which has the side effect of increasing the G2G latency. A more detailed analysis of the jitter versus the buffer size will performed during the next test sessions.

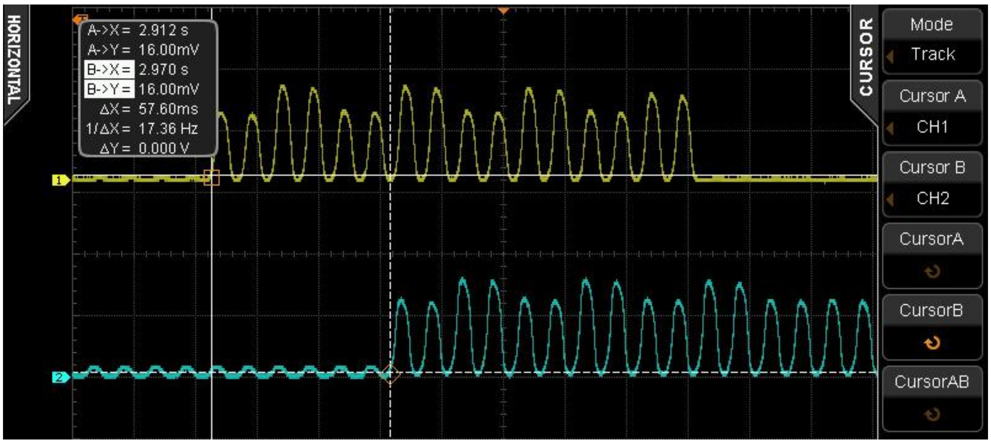

Figure 20 - The photodiodes source (yellow) and destination (blue) signals

3.2.4 Trial UC2.A.4 (wireless studio)

Table 11. UC2.A.4 measurement parameters

Measurement description	Result	Comments
Camera control validation over 5G using commercial RCP	Passed	The test was subjective. It depends on an expert controlling the camera both over wired network and 5G network and reports any change in experience.

The test case used an RCP currently available on the market to control the cameras. Typically, such functionality uses wired technology to ensure low latency communication between the control stick and the camera. The test, therefore, was initially executed using an ethernet connection and then it was replaced by 5G to evaluate the difference in performance on a real usage basis.

The person controlling the RCP didn't recognize any difference between Ethernet and 5G. Unfortunately, it was not possible to make any detailed analysis to compare the ethernet and 5G because of the lack of measurement tools within the RCP itself.

A video documenting the measurement was captured, which can show that there is no lag between the usage of the controller movement and the controlled cameras. A screenshot of the control session is shown in Figure 21.

Figure 21 - Remote control experience evaluation

3.2.5 Trial UC2.A.5 (remote contribution)

Table 12. UC2.A.5 measurement parameters

Measurement description	Result	Comments
Basic transmission tests, 0.6 & 1 sec end-to-end config, 1 HD 1080p feed, max capped at 30Mbps; 3 test cases	Passed	Achieved 30mbps UL for the single 1080p feed from Aachen and TV2 to RAI at 0.6 and 1.2 sec latency

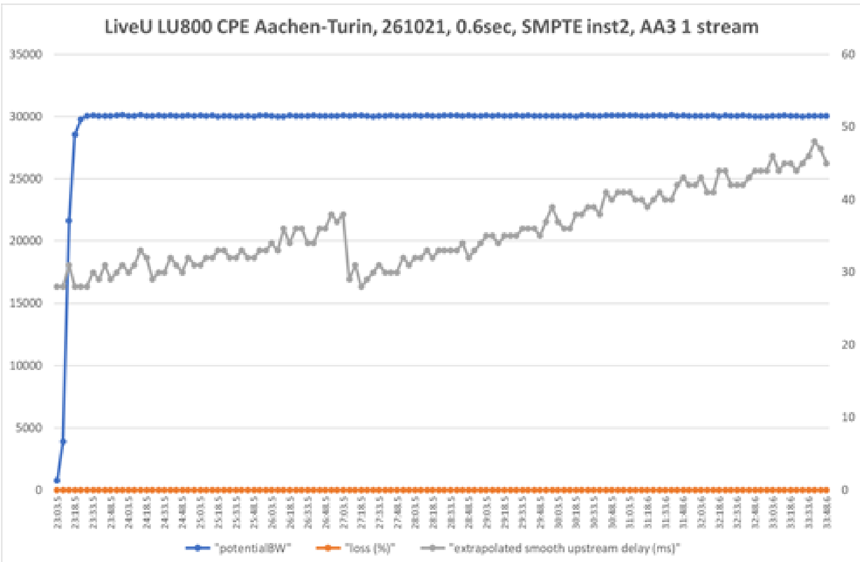

Figure 22 - Basic transmission test single stream from Aachen

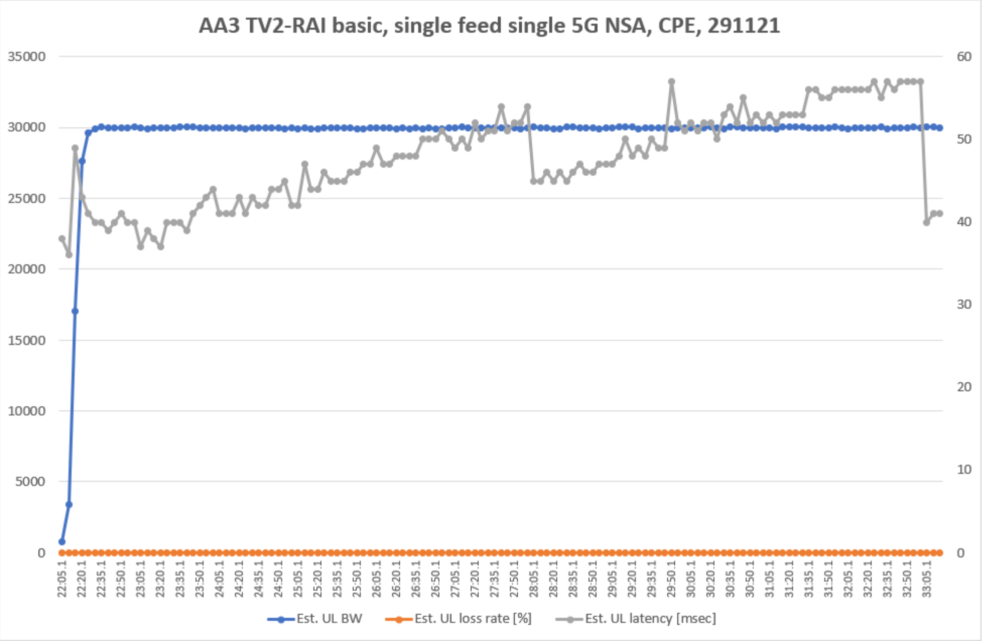

Figure 23 - Basic transmission test single stream from TV2

3.2.6 Trial UC2.A.6 (remote contribution)

Table 13. UC2.A.6 measurement parameters

Measurement description	Result	Comments
Multi-cam transmission, @ 1 sec and 1.4 sec end-to-end config; 4 x HD 1080p feeds, max capped at 60Mbps; 4 test cases	Passed	Achieved 60Mbps UL from Aachen and from TV2 to RAI at all latencies, stable; all videos came out ok and synced

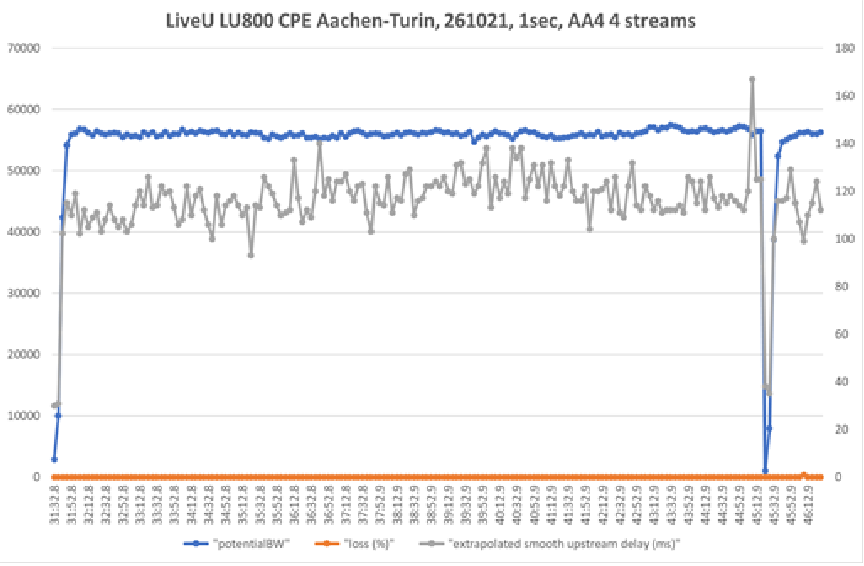

Figure 24 - Multi-cam transmission test, 4 streams

3.2.7 Trial UC2.A.7 (remote contribution)

Several tests have been performed at RAI's labs in Turin regarding the SMPTE ST 2110 compliancy of the IP output for the LiveU's decoder LU2000-SMPTE. Among JT-NM's SMPTE ST 2110 compliancy test list [8], 16 tests have been selected and evaluated. Results are reported in the following tables.

Table 14 - ST 2110-10 compliancy tests results for LU2000-SMPTE receiver

ST 2110-10 test description	Result	Comments
TX provide SDP	Passed	
SDP validated via SDPoker and/or manually	Not Supported	Two attributes are missing ('ts-refclk', and 'fmtp' for audio stream).

Table 15 - ST 2110-20 compliancy tests results for LU2000-SMPTE receiver

ST 2110-20 test description	Result	Comments
Stream present	Passed	
Multicast address correct	Passed	
Video format correct	Passed	
Decoded by reference RX	Passed	
No visible errors	Passed	
No errors reported by PRISM	Passed	
Sender N and/or NL and/or W	Passed	Narrow Gapped sender profile.
Cmax compliant	Passed	
VRXfull compliant	Passed	

Table 16 - ST 2110-30 compliancy tests results for LU2000-SMPTE receiver

ST 2110-30 test description	Result	Comments
Stream present	Passed	
Multicast address correct	Passed	
DSCP marking according to AES67	Not Supported	Media streams and PTP packets marked with DSCP value 0 (Default)
Stream audible	TBC	Audio stream is audible but it doesn't seem clear. Further tests are needed.

Table 17 - ST 2022-7 compliancy tests results for LU2000-SMPTE receiver

ST 2022-7 test description	Result	Comments
Video stream redundancy working	Passed	

Figure 25 - LiveU received SMPTE video in RAI's Lab

3.2.8 Trial UC2.A.8 (remote contribution)

Table 18. UC2.A.8 measurement parameters

Measurement description	Result	Comments
Transmission under load, at 50%, 75%, 90% load emulation; 1 HD 1080p stream max capped at 30 Mbps and 4 HD 1080p streams max capped at 60 Mbps; 9 test cases	Passed	As long as the needed transmission rate was below what was left after the load emulation, the LU800 used the bandwidth that the network allocated to it as remaining from the load emulation. Allocation was more or less relative to what remained from loading it and the LU800 adapted its transmission accordingly. So, transmission was sustained at 30Mbps single stream at 50% load (50Mbps were remaining) and sub 10Mbps when 90% load was done. Video was still consistent, not broken and adaptive. Also, stability of the video parameters experienced by the LU800 when the network was loaded were unstable and jittery, even in this synthetic conditions.

Examples of transmission performance under load (max capped in the device @30Mbps for single feed and @60Mbps for multi-cam) can be found below. A few insights:

- UL latency increases and available UL is decreased with more network loading (50%, 75%, 90%)
- LU800 video encoding and transmission adapts to available BW, both decreasing and increasing the bit-rate
- All 3 streams are still encoded and transmitted even at high network load, BW is split internally between them

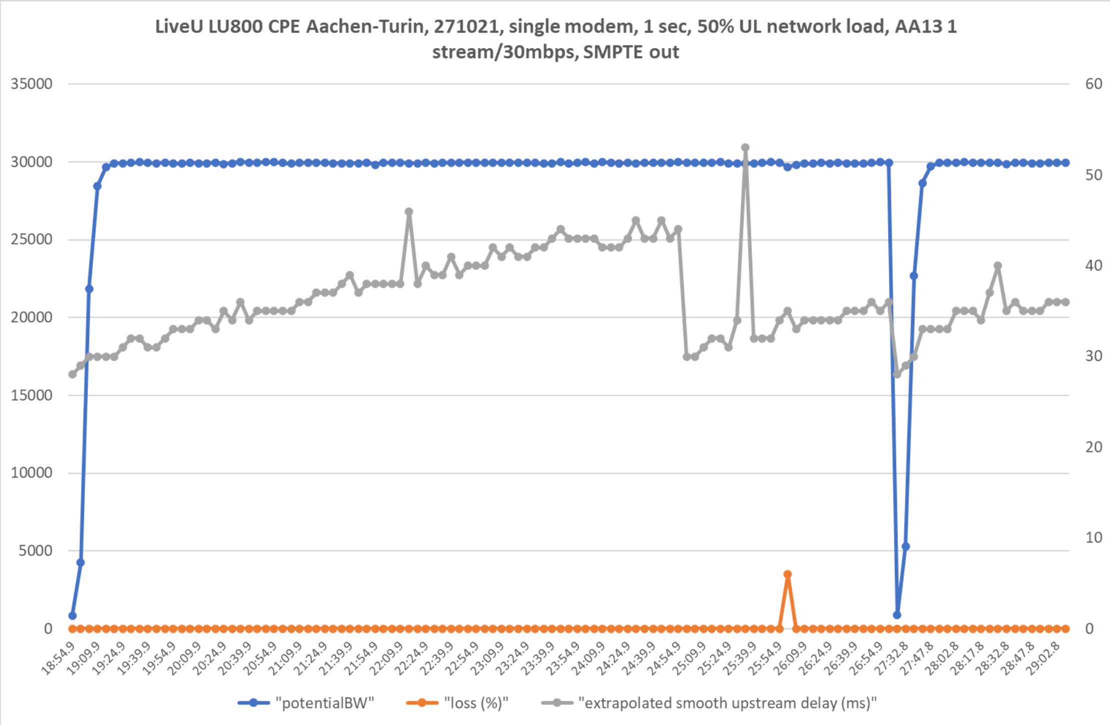

Figure 26 - Transmission under load, 50% UL, single stream

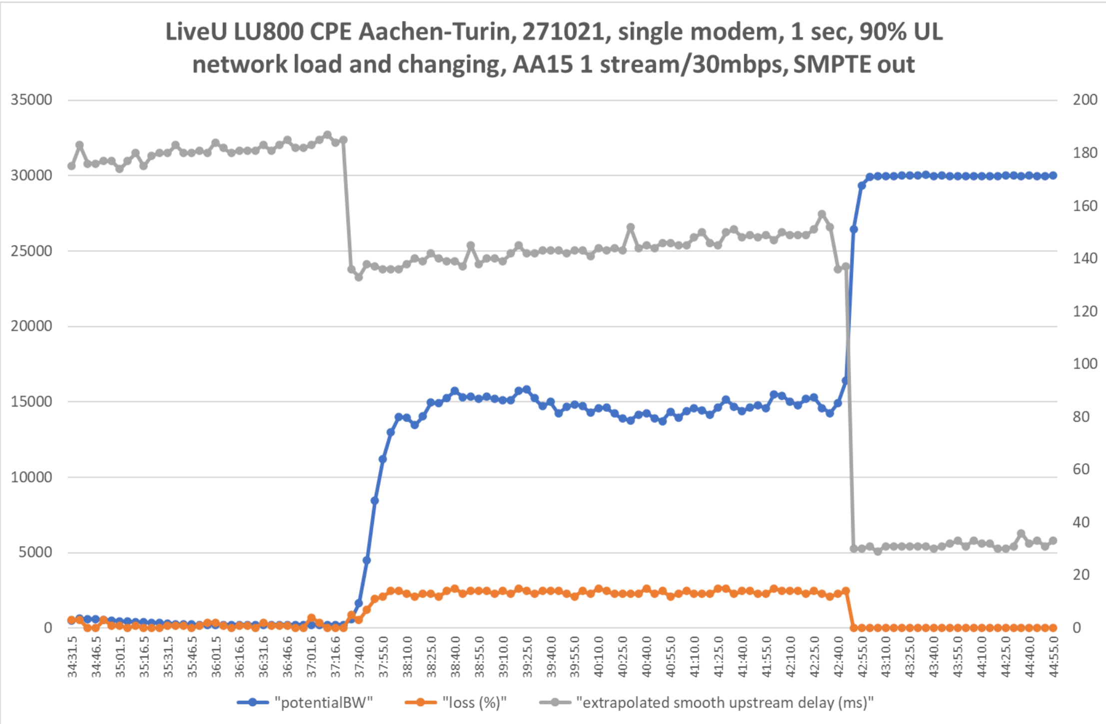

Figure 27 - Transmission under load, 90% UL, single stream

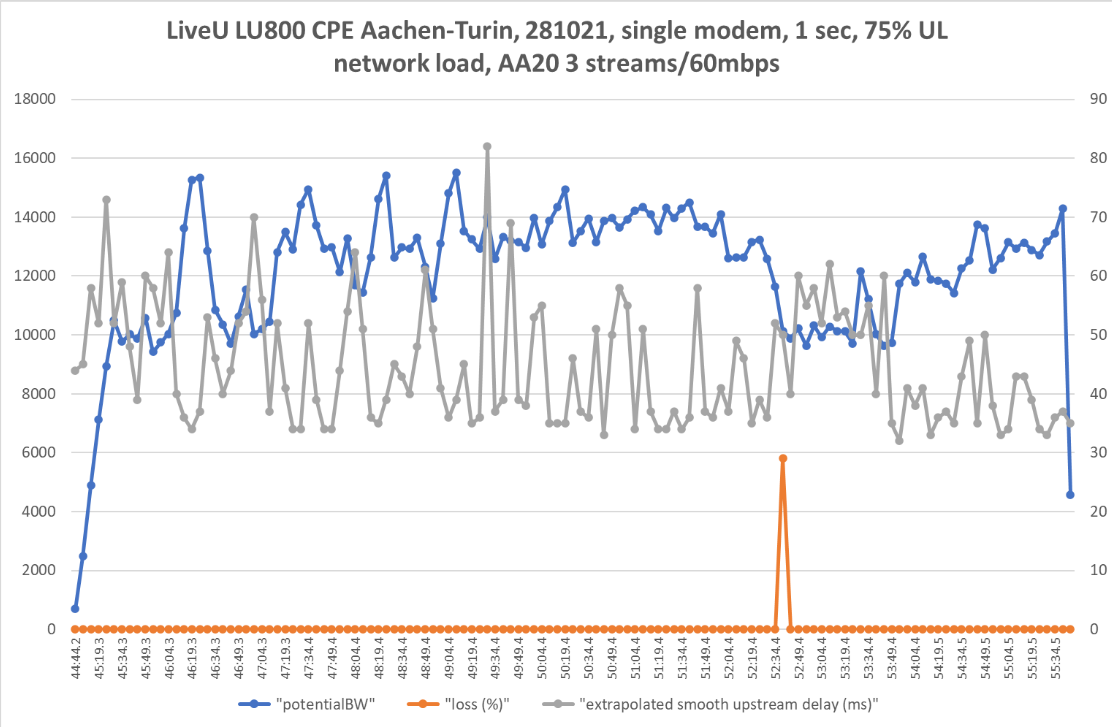

Figure 28 - Transmission under load, 75% UL, 3 streams

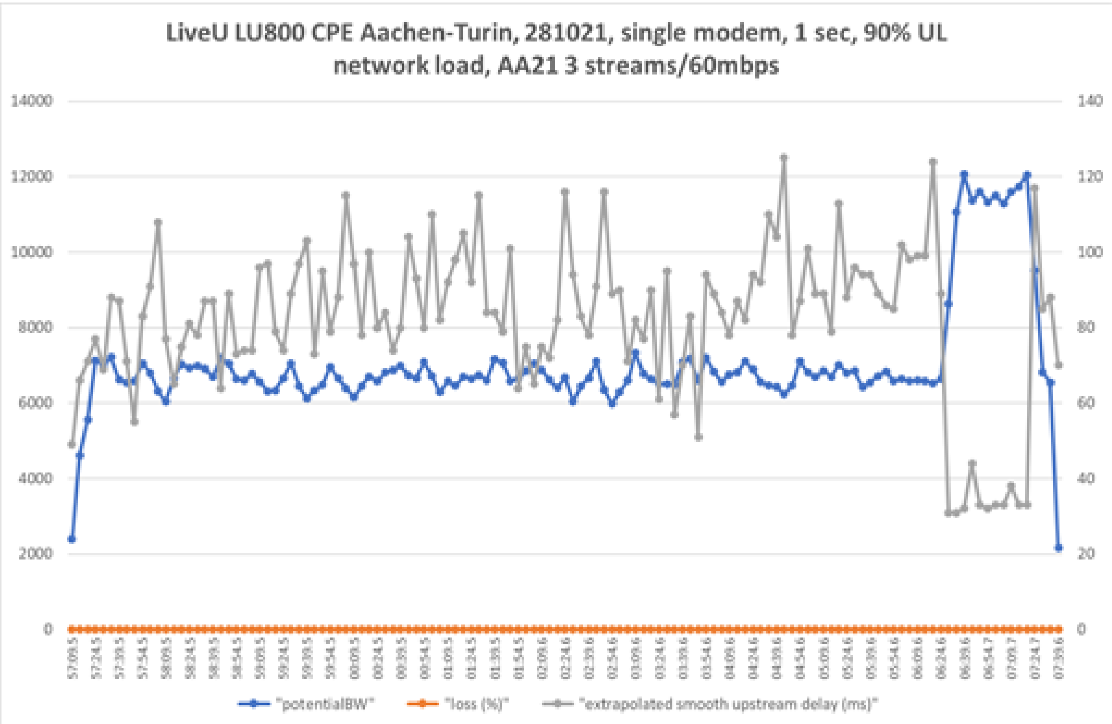

Figure 29 - Transmission under load, 90% UL, 3 streams

3.2.9 Trial UC2.A.9 (remote contribution)

Table 19. UC2.A.9 measurement parameters

Measurement description	Result	Comments
Remote contribution audio (IFB, intercom), several tests	Passed	-

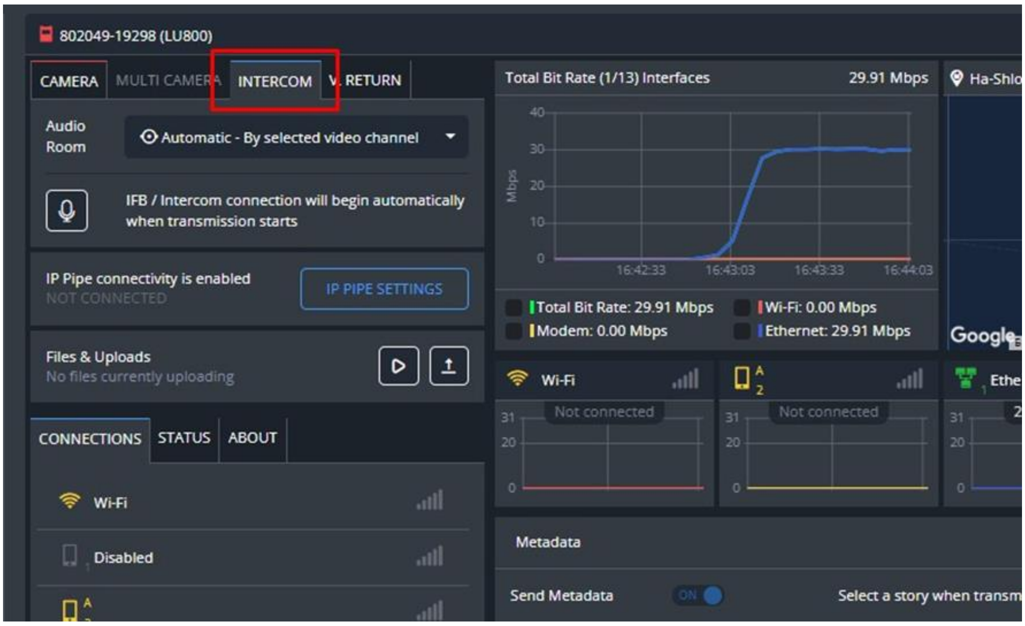

Figure 30 - LiveU intercom/IFB connectivity

3.2.10 Trial UC2.A.10 (remote contribution)

Table 20. UC2.A.10 measurement parameters

Measurement description	Result	Comments
IP PIPE (over public internet & 5G) end-to-end low traffic volume, using ping	Passed	Approx. round trip times: Minimum = 51ms, Maximum = 95ms, Average = 54ms. 0% packets lost
IP PIPE (over public internet & 5G) end-to-end for controlling cameras from remote, few test cases/changes	Passed	Approximately 0.8 sec latency from command to video change back in the remote studio (5G via LiveU DL, Cyanview processing, camera processing, video flow via LiveU back up)

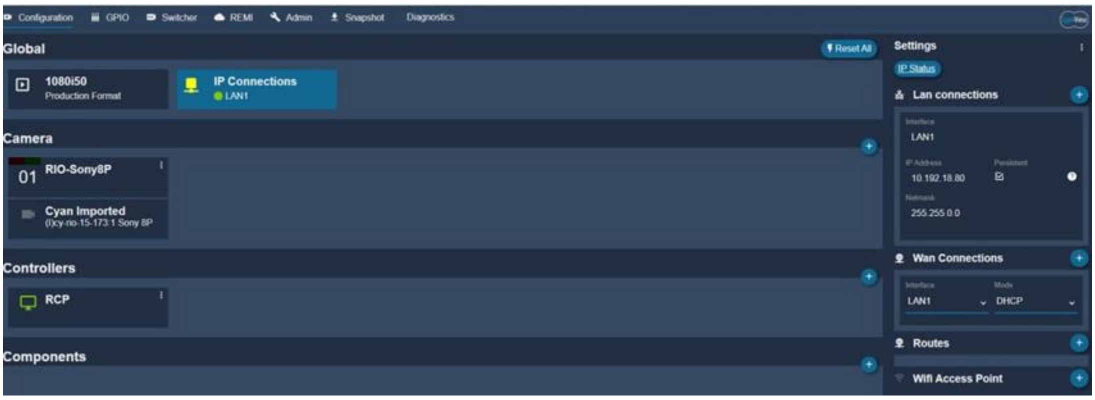

Figure 31 - LiveU IP PIPE connectivity

Figure 32 - Cyanview remote camera control via LiveU IP PIPE

3.3 KPI analysis

3.3.1 Uplink throughput (wireless studio)

The uplink throughput is relevant to various network configurations such as the available frequency bandwidth, the TDD pattern used, and the modulation scheme. The uplink throughput measured by SCREAM and iPerf has ranged between 90 Mbps to 100 Mbps. The measurements were collected using a single UE attached to the network, hence, all the uplink resource blocks are consumed by the UE. However, the available bandwidth is shared among all the UEs accessing the uplink resources at a given time.

Table 21 - 5G System configuration

Parameter	Value
3GPP Release	15
Spectrum	3700 MHz – 3800 MHz
5G Network type	Standalone
Carrier bandwidth	100 MHz

According to D2.1, the system should support five cameras with 50 Mbps each, which requires a total uplink throughput of 250 Mbps. The 5GS can achieve the desired throughput by allocating more carrier bandwidth. However, the regulator assigns up to 100 MHz for industrial usage. See more in D2.2 about regulatory framework.

Assuming to use the available carrier bandwidth, additional tests are needed to understand the minimum video bit-rate needed to ensure an acceptable quality. More sophisticated approaches (e.g., using different bit-rates for the on-air camera and preview cameras) could be also investigated. 5G QoS can also be used to ensure that the on-air camera gets a higher priority in terms of bandwidth and packet loss. 5G QoS analysis will be performed during the next test sessions.

According to D2.1, the ideal G2G latency value is 40 ms. The value measured in the first test is 57.6ms with an unstable signal. There are various factors affecting G2G such as encoder/decoder capability and network latency. The encoder and decoder are using very low latency configurations and operate with a small buffer to achieve such low latency. However, a larger buffer can be used at the decoder side to ensure stream stability against varying jitter. Further tests will be carried out and reported in D5.3

3.3.2 UL throughput, latency, packet loss rate (remote contribution)

The KPIs for UL throughput, UL latency and UL loss rate for the E2E chain over the Ericsson’s 5G lab environment and the public internet, were mostly met. In this setup, 30 Mbps for single UL transmission was achieved and maintained stable at 600 msec latency, as long as the network was empty or with load emulation below 50% (which means at least 50 Mbps were available for the transmission). When the load emulation was raised to 75% and 90%, the available UL bandwidth left for the transmission forced the LiveU devices to reduce the transmission rates by adapting their video encoder settings accordingly. However, the network latency was not stable for the whole period, with several cases of latency largely exceeding the 100 msec max seen in most stable transmissions. Moreover, the available bandwidth, as assessed by LiveU monitoring at the application level, was not stable at all times. These phenomena, on behalf of the 5G lab infrastructure, were expected because the infrastructure had to satisfy the other traffic as well as the video traffic. This was the reason why the adaptive video encoding and transmission algorithms were employed.

When multiple streams were transmitted (max. 15mbps per stream, 60-70 Mbps total for 4 streams), similar performance measurements were experienced, only more severely. When the available network bandwidth was sufficient (at 0% load emulation) 4 video streams at 60-70 Mbps total and bandwidth split evenly between them was achieved. When the available bandwidth dropped, the video encoding and transmission settings were adapted evenly between the streams, yet the experienced network stability was worse than previous tests, probably because the transmission attempts made by both the traffic emulator and the video transmission were exhausting the network even further and with more momentary changes in the demand patterns.

Regarding the IP PIPE bi-directional communication path, a latency of ~80-90 msec one direction (DL, from RAI media lab to the LU800 in Aachen) was measured using ping.

This means that the commands from remote can arrive quickly and control the field device with a satisfactory overall user experience.

In all tests the UL packet loss rates at the application level were almost zero.

3.3.3 Multi-cam via the LU800Pro (remote contribution)

This KPI was met as 1 - 4 camera streams were successfully transmitted simultaneously over the lab 5G (under several load emulation settings) and the public internet to the RAI media lab. Bandwidth was split by the LU800Pro between the transmitted streams evenly and in accordance to the total available UL. Latency was not affected by adding more streams under the same network conditions.

3.3.4 SMPTE 2110 10-20-30 video output compliance (remote contribution)

The output of the LU2000SMPTE SMPTE video was measured and passed the compliance test of the Tektronix Prism at the RAI lab. The LU2000SMPTE supported 3 different IP sub-networks: one for the video packets coming in from Aachen, one for the SMPTE video IP output and the 3rd for the PTP master clock. The validation of this KPI means that the video output can be integrated with a broadcaster's IP SMPTE network.

3.3.5 Remote audio communication (remote contribution)

This KPI was met using the LU800Pro and LU200SMPTE, enabling bi-directional audio communication between a remote producer/director and a field camera operator. In some cases, the audio quality was not great, yet intelligible, at no noticeable latency.

3.3.6 Cameras remote control (remote contribution)

This KPI was not initially considered for this UC scenario but, due to the progress made during the T5.2 timeframe, it was possible to add it to the achieved KPIs. Using the bi-directional LiveU LU800 - LU2000 IP PIPE in parallel with the video UL transmission, the remote control of camera shading/iris settings using Cyanview camera control devices was performed and the functional test was passed, giving a satisfactory user experience.

The overall E2E latency from sending the command (DL) to changing the camera shading/iris and displaying back the video result took about 700-800 msec. This latency value is acceptable in a remote production/contribution environment for less demanding scenes (e.g.: quasi-static lighting conditions, slow subjects), allowing to reduce on-site resources by having a remote expert team controlling several events, etc. It is worth noting that, whenever it would be possible to reduce the latency of the video UL path, this round-trip latency would also be reduced accordingly.

3.4 Deployment considerations for final stage

During the initial stage, only a limited number of components were available for testing. Therefore, the trials had to use products available in the market and perform individual validation of the components developed during that stage. It was also not possible to perform tests in mobility within the 5G lab. To overcome those limitations of the initial stage, the following points will be considered for the final stage:

- Execute an E2E validation test using all the developed components and re-evaluate the KPIs.
- Perform tests in mobility using the field trial network and simulating a portable camera, encoder and modem scenario even if the components will not be integrated all together with the right form factor
- Configure traffic prioritization in the 5G network based on QoS and network slicing to ensure a consistent performance for high priority traffic (e.g., on-air camera)

4 Use Case 3: Live Immersive Media Production

This chapter describes the outcomes of the first stage of trials in the context of UC3. Testbed architecture and measurements' results will be presented first and will be followed by a detailed KPI analysis (with regards to the KPIs defined in D2.1 [1]), as well as the deployment consideration useful for the final stage of the trials.

The focus of the first stage is validating the end-to-end functionality of the use case. The following scenarios from D2.1 will be validated:

- Content Production Simple Scenario:
  - Low scene complexity (few objects, short depth range)
  - Narrow FVV angle, with 1 DoF (within arc defined by reference cameras)
  - Single rendered view, live only
  - Up to 5 simultaneous uplink streams
- Content Delivery Scenario:
  - End-to-end multimedia gold slice for view renderer and Premium end user.
  - Partial multimedia gold slice for local producer and on-premise end user (from near-edge to edge cloud).
- Deployment Scenario 5G Theater:
  - 5G NSA
  - Over "existing" 5G infrastructure, integrated in the mobile network
  - Small deployment, limited local user services (a few downlink clients)

The next figure shows the simplification of the use case for the first stage.

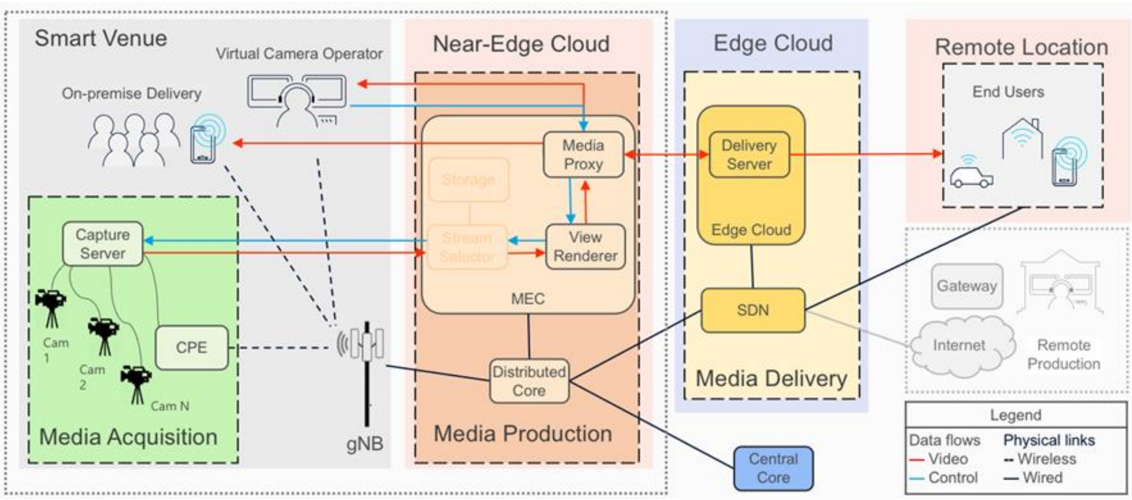

Figure 33 - UC3 functional architecture

4.1 Deployed testbed architecture

The technical architecture of the deployed testbed is depicted in the following figure.

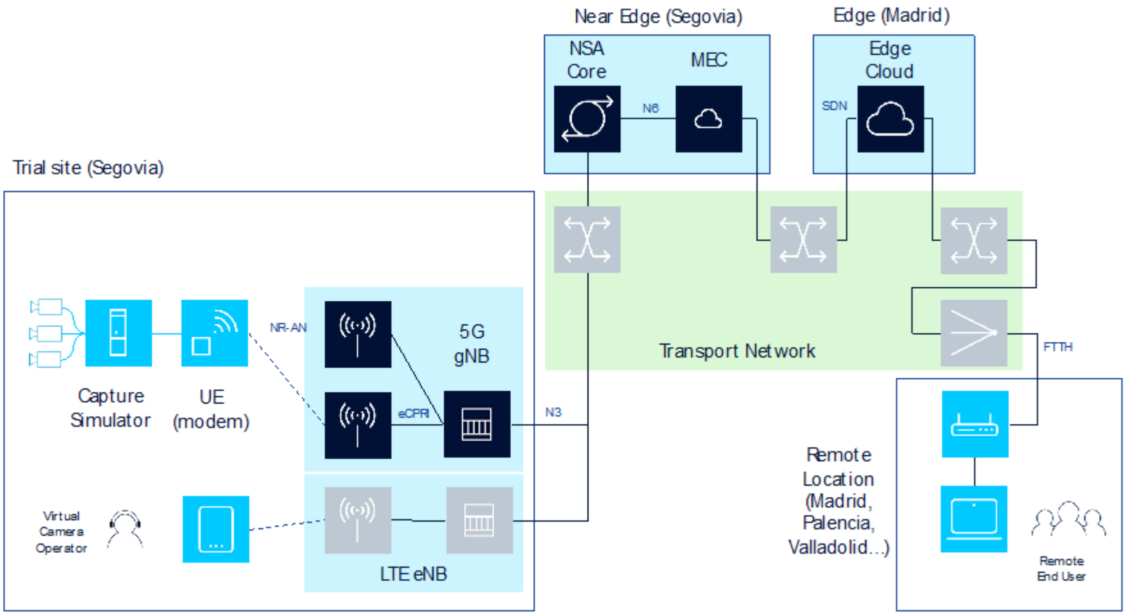

Figure 34 - UC3 trial network architecture

The aim of this first-phase testbed has been doing an end-to-end trial of the network functionality in all its areas, particularly:

- Deploying and validating a mmWave trial location in Segovia, Spain.
- Deploying and validating the edge cloud infrastructure required for the project.
- Providing end-to-end connectivity over Telefónica commercial transport network, including QoS management (slicing).

a b c

Figure 35 - UC3 RAN access: radio units (a, b) and BBU (c)

For the Radio Access Network, a full gNodeB was deployed in a location close to Segovia (Spain). This RAN used band n258 in 8CC DL / 2CC UL configuration, with 100 MHz in each band (for a total of 800 MHz). Two Nokia AirScale mmWave Radio Units were deployed, as well as one LTE unit for the NSA control plane. A Nokia Airscale Baseband Unit was installed in the same trial location and connected to Telefónica transport network. The site has been jointly operated by Nokia and Telefónica using Telefónica experimental PLMN.

The near-edge cloud was installed in Segovia at 5km from the trial site. It includes the user plane functionality of a distributed network core (Nokia cloud core), as well as the MEC running the production functionality. Physically it is based on Nokia OpenEdge servers with Intel Xeon CPUs and NVIDIA Tesla T4 GPUs. Logically, it implements ETSI MEC functionalities using MicroStack to administrate and orchestrate the VNFs.

The delivery edge cloud, located in Madrid, is made up of three layers: network, compute and storage. Physically, it is built with six whitebox switches in leaf-spine Clos topology, seven whitebox open computing platform (OCP) servers, and one OCP storage server with just a bunch of disks (JBOD).

In addition, several remote user locations have been established in different Spanish cities: Madrid, Valladolid and Palencia. These remote locations are connected via residential FTTH access to Telefonica network so that the end users will access the content in real time.

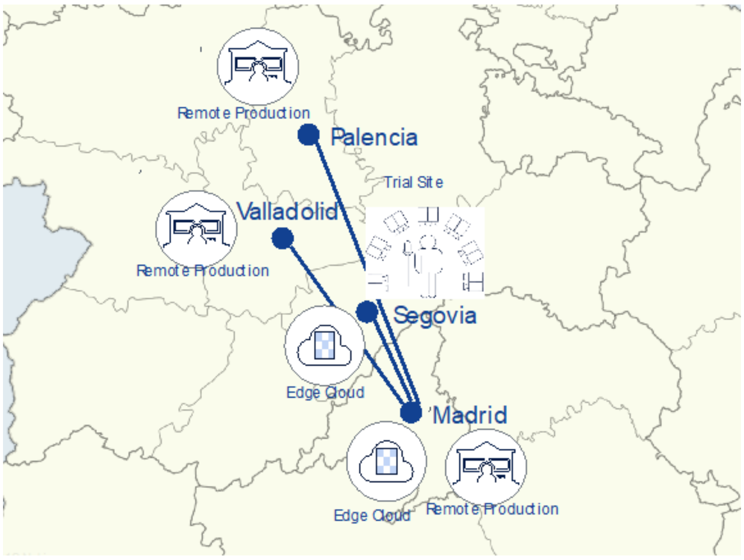

Figure 36 - UC3 trial locations

End-to-end connectivity and QoS management were guaranteed through Telefónica transport network. Two transport slices have been defined, covering the network from the near-edge MEC to the delivery edge, and from the delivery edge to end users: best effort (CS0) and multimedia gold (AF41).

More details on the specific components involved and how they operate can be found in deliverable D3.1 [2].

4.1.1 Updates on measurements planning

The initial focus foreseen for the phase 1 trial was an end-to-end funtional tests with a live transmission from a reduced set of cameras from the trial location (Segovia) to an individual user in a remote location (e.g. Valladolid). However, COVID restrictions during the omicron wave had made it impossible to have full access to the indoor part of the trial location, particularly the sections where the trial stage was to be deployed. Due to this, some of the measures involving camera deployment in the trial location have been deferred to a second phase, and the focus was shifted to consolidating the network part, particularly the field integration of the three main enablers of the use case: mmWave radio, edge cloud and end-to-end QoS management to several locations.

All the foreseen trial conditions have been executed, but some of them presented specificities or limitations which will be completed in the second phase of the project and are described in a case-by-case situation below.

4.1.2 Uncertainties and risk assessment

Two relevant issues have been found during the deployment of these testbeds and the execution of the trials:

- Due to COVID-related restrictions, the access to the public part of the trial location (where cameras were planned to be deployed and tested) has suffered from strong limitations.
- Related to the COVID pandemic and its effects, including the limitations on the supply chain and the unavailability of UE FR2 chipsets with 5G Standalone capabilities, it will not be possible to deploy a 5G-SA production-ready “portable” NPN deployment for the second phase of the project.

To prevent these two issues to impact on the project milestones, a mitigation plan has been established, which will be described later in section 4.4.

4.2 Measurement results

4.2.1 Trial UC3.A.1 (functional validation)

As described before, end-to-end functional trial with a real camera deployment has been deferred to the second phase of the project. Therefore, the camera flow has been replaced by executing the following path:

1. Offline capture of a video scene.
2. Send the video scene from the camera capture server through the uplink.
3. Render the virtual view in the MEC at the near edge cloud in Segovia.
4. Send the rendered view through the multimedia gold slice to the Media Delivery server in Madrid edge cloud.
5. Serve the content from the Media Delivery to a remote location.
6. Watch the video in a player.

All the segments have been validated. Figure 37 shows an image of the player with the video.

Figure 37 - Video player with the content for functional validation of UC3.A.1

4.2.2 Trials UC3.A.2/3 (uplink and render tests)

The critical path of the production pipeline is the link between the output of the individual cameras (or the capture server) and the rendered view. This covers three different but related KPIs:

- The cameras should produce a bitrate which is low enough to be supported by the 5G network.
- The radio uplink should support the bitrate produced by the cameras with negligible packet loss
- The MEC should support the rendering of the received frames in real time.

Several related measures have been taken to support this path, which cover use cases UC3.A.2 and UC3.A.3 from D5.1 [4], and even going beyond what was defined there.

Firstly, several performance tests have been done using iperf3 between the UE and the MEC, testing both uplink and downlink (not simultaneously) from several locations at several distances to the antenna (with and without line of sight). The results are shown in the following figure. Optimum performance is obtained at 20-60 meters distance from the antenna, which is also convenient for the installation of the modems close to the possible deployment locations for the cameras.

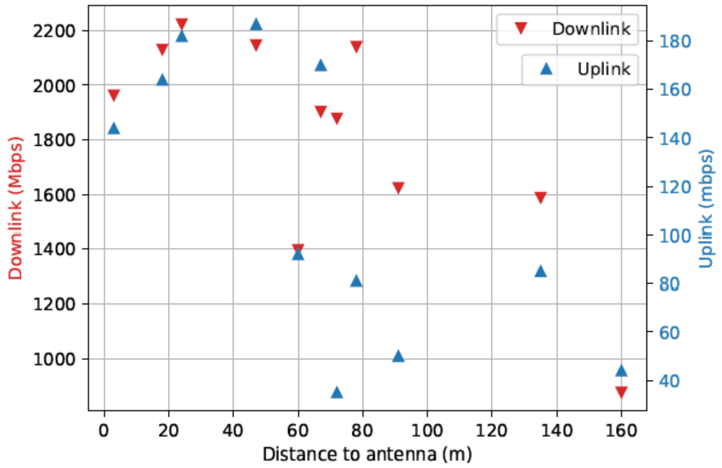

Figure 38 - Uplink and downlink throughput at different distances to the antenna

With the modem installed in one of those locations, several uplink tests have been done from the camera simulator to the view renderer.

The received average bitrate has been measured along with the percentage of packet losses for different number of camera video streams transmitted, resolutions, frame rates, and scene complexities. In particular, regarding the latter, three different scenarios have been considered in terms of the captured content. These scenarios involve different levels of complexity and therefore bitrate demands:

- Simple: small number of objects with little or no motion.
- Medium: more objects with limited motion and some occlusions.
- Complex: a lot of objects with significant motion and numerous occlusions.

Table 22 shows the achieved bitrate for these configurations.

Table 22 - Measurements Trial UC3.A.2/3. Values show uplink throughput in Mbps

FPS	30	30	15	15
#RGB-D	3	5	3	5
720p simple	33.8	54.2	-	24.7
720p medium	46.9	79.4	-	45.2
720p complex	135.7	212.6	-	120.0
1080p simple	167.5	-	55.6	90.1
1080p complex	>220	-	180.2	>220

All results in Table 22 with a white background indicate a packet loss rate which is negligible (< 10^{-5}). Results with a grey background have significant packet loss rates (> 2%), which makes the resulting video unusable from a video quality perspective.

At the view renderer, the same streams have been rendered and the performance in frames per second has been obtained. The results are shown in Table 23. As seen from the combined results, the system is functional at 15 fps with 720p and 3 to 5 cameras. It can work at 720p@30 or 1080p@15 but there will be limitations in the GPU rendering capacity.

Table 23 - Measurements Trial UC3.A.2/3. Values show rendering FPS

#RGB-D	720p simple	720p medium	720p complex	1080p simple	1080p complex
3	27.20	25.69	21.67	16.23	11.83
5	26.72	25.28	20.93	18.27	10.29
9	25.97	24.53	20.78	17.85	8.83

4.2.3 Trial UC3.A.4 (motion-to-photon latencies)

Motion-to-photon latencies have been measured using the tool depicted in Figure 39. The virtual camera control runs in a mobile device, which is connected through LTE to the eNB in the NSA deployment. Besides, due to COVID limitations, it was not possible to perform the test on Segovia trial deployment, but only in Madrid's integration laboratory.

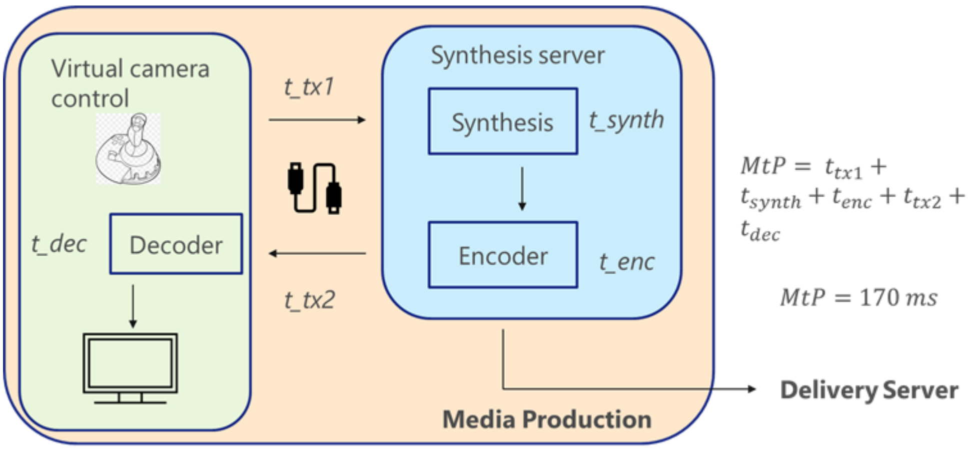

Figure 39 - Motion-to-photon measurement procedure

Measurement results are shown in Table 24 for different scenarios.

Table 24 - Measurements Trial UC3.A.4

Scene	Mean (ms)	Std (ms)
720p30 simple	300	89
720p30 medium	280	78
1080p30 simple	290	45

To validate latency-related KPIs, this trial was complemented with RTT ping tests from the UE to the MEC, which were already reported in D4.1 [3], but are shown here for completeness.

Table 25 - RTT measurements for UC3

Test	Duration	RTT min	RTT mean	RTT max
Ping	20 min	7.6 ms	11.6 ms	19.0 ms
Ping + 50 Mbps	20 min	6.8 ms	22.0 ms	70.1 ms
Ping + 150 Mbps	20 min	47.3 ms	150.2 ms	567.0 ms

4.2.4 Trials UC3.A.5/6 (performance difference in slices)

To test QoS differentiation between transport slices, a trial has been performed where the video is played by end users in residential locations and some background traffic is added as artificial noise: 0 (no noise), 250 Mbps, 500 Mbps, 800 Mbps, and 1 Gbps. During each test, the video player collected application and traffic metrics and sent them to a central monitoring system, where they were stored.

Figure 40 shows a summary of relevant results for a single remote location and three video qualities: automatic, 520p and 720p (fixed), all at 30 fps. They map the qualities which would be available from a 720p stream, as the result of the previous tests suggest.

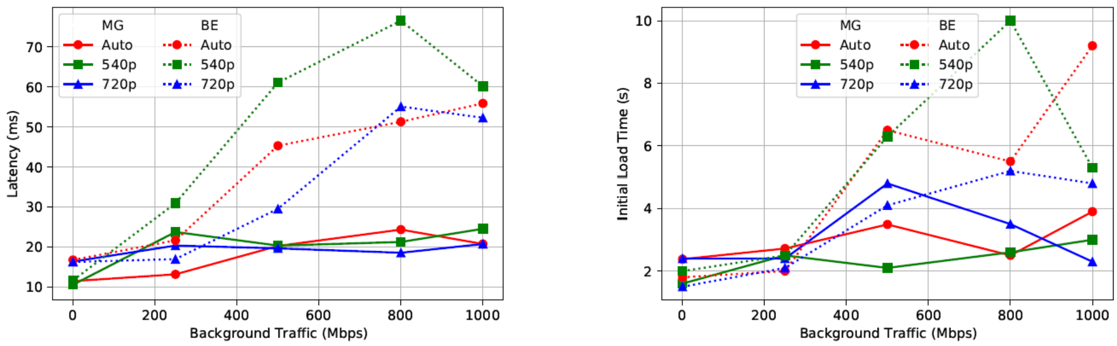

A more detailed analysis has been performed with several tests from several locations and several resolutions and bit rates. Figure 41 and Figure 42 show prioritized and no prioritized traffic KPIs (average round-trip time (RTT), average jitter and the initial load time) for different locations (Madrid, Palencia and Valladolid), five video qualities (automatic, 540p, 720p, 1440p and 2160p) and several scenarios with artificial noise added.

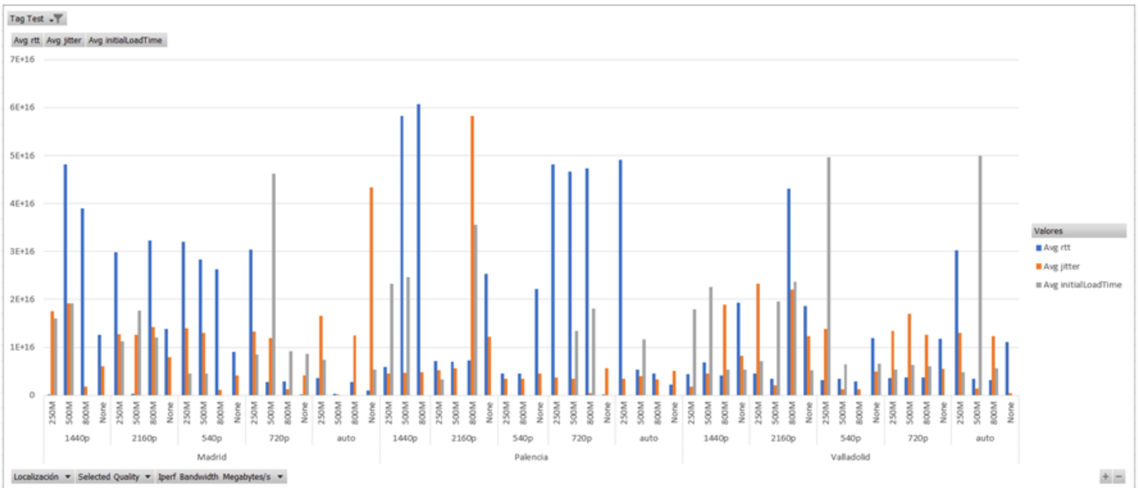

Figure 41 - Prioritized traffic from three Locations adding noise. (Jitter, RTT and initialLoadTime)

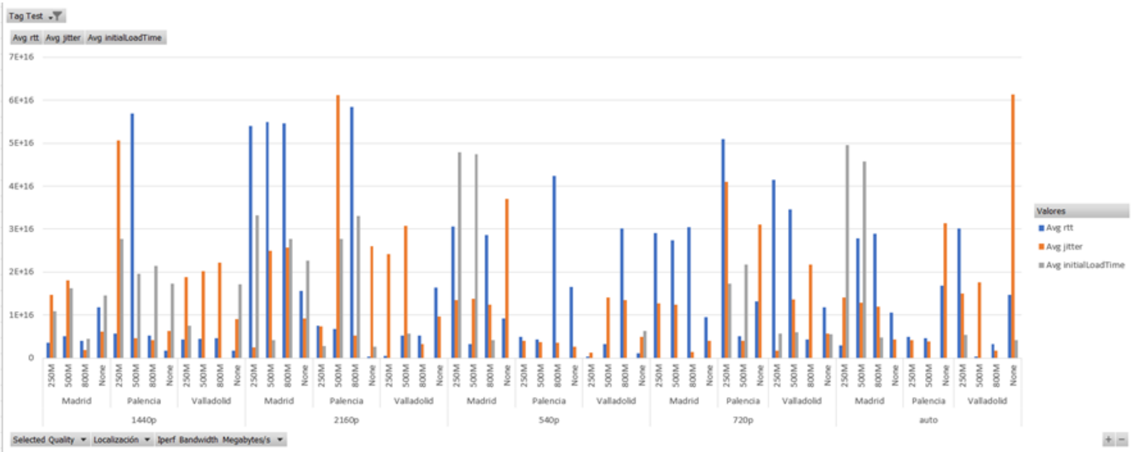

Figure 42 - NonPrioritized traffic from three Locations adding noise. (Jitter, RTT and initialLoadTime)

All these results cover the segment between the residential user and the media delivery (distribution), defined in trial case UC3.A.5. The contribution segment between the media delivery edge cloud and the production edge cloud (UC3.A.6) has been validated functionally: slices BE and MG are both working correctly. Detailed performance testing will be done in a next phase.

4.2.5 Trial UC3.A.7 (slice automation)

As seen from previous definitions, two transport slices are defined: one slice will be in multimedia priority and the other one in best effort priority. In this trial, test slice automation (i.e. move user between slices) will be tested for remote users. Figure 43 shows the diagram of the slice automation scenario.

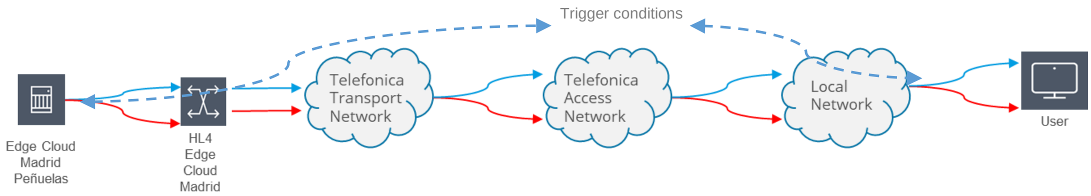

Figure 43 - Slice automation scenario

Unlike in the previous test cases, all the users in this scenario started their test session in the best-effort slice. Then some noise was added, until some of the monitored KPIs were impaired. When this happened, the slice automatically switched to multimedia gold via the slice selector functionality described in D3.1. The trigger conditions for slice change are shown in Table 26.

Table 26 - Conditions for slice change: trial UC3.A.7

	User on Slice	Latency	Jitter	Initial Load Time	Pause Count	User change to
Test 1	BE	> 200 ms	-	-	-	**MG**
Test 2	BE	-	> 3		-	**MG**
Test 3	BE	-	-	> 120seg	-	**MG**
Test 4	BE	-			>1	**MG**

Functional validation of the automatic slice change has been performed. A more detailed analysis of the performance of this functionality will be done in the second phase of the project.

4.3 KPI analysis

A summary of relevant KPIs presented in Table 27 - KPI analysis for UC3, assuming 720p30 content (see more in D4.1 [3]). We have assumed 720p video, at 30 fps. For the delivery KPIs (5, 6), the results consider what it is achievable in the Multimedia Gold slice under moderate background noise (250 Mbps).

Table 27 - KPI analysis for UC3, assuming 720p30 content

ID	KPI	Target	Current
1	Motion to photon	170 ms	280 ms
2	Uplink (camera)	100 / 50 Mbps	45 Mbps
2	Uplink (UE)	150 / 300 Mbps	180 Mbps
2	# Cameras	3 / 5	3
3	RTT UE-MEC	40 ms	11 – 150 ms
4	Render FPS	15 / 30 fps	22-27 fps
5	Remote user throughput	20 / 40 Mbps	20 Mbps
5	Delivery latency	150 ms	< 30 ms
6	Initial load time	1 s	3 s
6	Pause count	0	0

The basic conclusion is that the system is working in the basic scenario: 720p at 15 fps (end to end), 30 fps if we relax rendering real-time conditions (e.g. for replies). Delay measures at the production segment (RTT and motion-to-photon) are slightly underperforming, especially under load, but functional. Delivery KPIs are not fully met either, but they are quite stable even under load.

Potential improvements will be tested in the second phase of the project:

- The production console (virtual view controller) will be re-implemented over a PC platform, which allows us to connect it to the 5G RAN. This could improve motion-to-photon latency.
- Media delivery integration with client will be reviewed to check for potential improvements in start time and/or buffering periods.

However, beyond those two paths, there is little room for further improvement (e.g., in uplink capacity) due to the unavailability of the next generation of SA devices for FR2. Therefore, the strategy in Phase 2 will be to test the different parts of the chain as far as possible, even with its existing limitations, so that we can establish the requirements for the next generation of systems: radio links to support 1080p, 9 cameras; GPUs to support 1080p, 30 fps; and automatic slicing, including RAN.

4.4 Deployment considerations for final stage

During this phase of testing, some limitations have been found, mostly on mmWave availability (modems, radio, core): they are NSA only, strongly biased to DL (UL configuration limits reached), and there are high chances that this does not change in project lifetime. Therefore, for the final stage of the project, and to deploy the “5G Festival” scenario defined in D2.1, the following considerations have been done: *(i)* the deployment will be 5G-NSA. Therefore, it will be in a place with LTE availability (e.g., Nokia premises); *(ii)* The uplink of some of the cameras will be emulated, so that the total available throughput is increased and 1080p configurations can be validated; and *(iii)* QoS differentiation on the RAN will be tested, even if limited by the usage of NSA core. Additionally, the unavailability of the “stage” part of the deployment in Segovia has been an issue to be able to perform complete end-to-end testing. Therefore, in the final stage the availability of a location to deploy and use the cameras should be guaranteed.

5 Conclusions

This document described the outcome of the first stage of trials for the project's use cases along with the changes of plans dictated by mutating conditions during the project. Due to the pandemic crisis, in facts, many components and their functionalities suffered from delayed development and the integration phase was, in turn, deeply compromised. Some of the measurements initially planned had to be adapted to the testbeds that was possible to implement.

The **UC1** testbed deployed for the first round of trials was a simplified version of the envisioned architecture. This simplified testbed included a single set of UE with 5G modem and audio components to reduce the complexity of scheduling and slicing in the 5GS. Two parallel flavours of the same architecture have been tested using different RAN components and Core Network providers: *(i)* a monolithic version of the OAI RAN (gNB) and the OAI Core Network; *(ii)* a disaggregated RAN components version incorporating the Accelleran CU and the OAI DU. These two parallel setups allowed to extract the E2E measurements leveraging the interoperability of the monolithic OAI setup with the Sennheiser equipment, before the E2E interoperability between the disaggregated RAN components and the Cumucore Core Network became available. Basic control and user plane interoperability has been validated between the Accelleran CU and the OAI DU from EURECOM. The interoperability between the Accelleran CU and the Cumucore Core Network has been separately performed and validated between the two parties, as well as the validation between Accelleran shared spectrum access client and RED Technologies shared spectrum access server.

The first set of trials showed that the 5G testbed currently presents an unexpected amount and variety of unpredictable packet loss (up to 1%), especially in DL direction. Even if it is still a bit unclear where the packet loss is occurring, it is assumed that the cause of the unreliability could be found in the gNB and/or the Linux drivers of the COTS UE modules towards the IP/application layers. Since in-depth analysis and optimization are needed to identify the causes of packet loss in the 5GS, packet loss and reliability were not considered in this deliverable. The synchronicity requirement was not considered as well at this stage. Regarding the network latency analysis, the first measurements have shown that the overall E2E latency is still considered to be in the 10's of milliseconds. While the majority of packets greatly benefited from the reduced TDD pattern, it can be noted that still many packets are outlining in the current testbed. Packets that are too slow for a potential jitter buffer have to be counted lost. When targeting a required packet error ratio of, at least, 10^{-6}, the operating point, so far, had to be chosen between 10 ms and 12 ms of E2E latency, or higher. Furthermore, it was spotted that at least one 5G component in Downlink direction was not working in sync with the selected time basis. This currently hinders the optimization of the Downlink E2E latency. Final stage trials will target the radio latency reduction required for UC1, i.e., ~1 ms radio latency at the OAI DU, using either the combination of 1.0 ms DL-UL pattern periodicity and 30 KHz SCS/0.5 ms slot duration, or 1.25 ms DL-UL pattern periodicity and 60 KHz SCS/ 0.25 ms slot duration.

In **UC2**, two trials phases were planned, one leveraging the available media components for the deployment and the second one integrating the newly developed features for the various components. For the 2 phases, a dedicated 5G network was deployed at Ericsson's Aachen lab to support the use case needed connectivity. Once the network was up and running, several tests have been performed for the two scenarios of the use case: the wireless studio scenario and the remote contribution scenario. Even though many components were not available at the beginning of the test sessions, progressively more of them have been introduced into the test workflows to finally achieve a full integration of the components and functionalities into the testbeds.

Regarding the wireless studio scenario, a set of tests have been performed on the network and the E2E chain, including: *(i)* UL capacity, *(ii)* packet loss, *(iii)* glass-to-glass latency, *(iv)* camera control, and *(v)* compatibility between equipment. The results were satisfactory overall and gave valuable information for the development of UC2. The KPIs analysis showed that the available UL throughput and packet loss needs to be improved but the regulatory framework does not offer much room for extra bandwidth. To overcome this limitation, a mechanism for selecting flow prioritization and throughput control should be implemented at MOCG level. The current glass-to-glass latency value (57.6ms) is very close to the proposed limit (40 ms) but the decoded signal wasn't stable. Additional tests need to be performed adjusting the receiver buffer size versus the jitter.

Regarding the remote contribution scenario, a set of tests have also been performed, involving mainly the evaluation of the network and the components performance and integration: *(i)* UL bandwidth, *(ii)* latency, *(iii)* packet loss, *(iv)* multi-cam support, *(v)* compliancy with video standards, *(vi)* remote audio communication, *(vii)* camera remote control. The results of these tests were satisfactory as well, as all the tests, with few exceptions, currently passed. The throughput and latency were well within the expected ranges for low traffic loads (under 50% load emulation) but started deteriorating severely when the network became more congested, especially in the multi-cam scenario. This behaviour, though, was already expected. In all tests the UL packet loss rates at the application level were almost zero. Both multi-cam and SMPTE 2110 video functionalities were successfully validated with just a couple of SMPTE tests not supported by the current setup. Other relevant functionalities for a remote production such as remote audio communication and remote camera control gave good results as well, resulting in a good overall user experience with negligible latencies in both cases.

In **UC3**, two trials phases were also defined. The aim of the first stage was to validate the E2E functionality of the use case. The trial tests focused on a live transmission from a reduced set of cameras from the trial location (Segovia) to an individual user in a remote location (e.g., Valladolid). However, COVID restrictions made it impossible to have full access to the trial location. Due to this, some measurements involving camera deployment on the field will be performed in a second phase.

During the execution of several trials, *(i)* the functional validation, *(ii)* the link between the output of the individual cameras and the renderer view, *(iii)* the motion-to-photon latencies, *(iv)* the QoS differentiation between transport slices and *(v)* slice automation scenario have been evaluated. It is worth noting that predefined target KPIs have been measured successfully in this deliverable, using specifically developed video tools (motion-to-photon latency tool, offline view renderer, etc.) and standard network measurement tools (i.e. iPerf3). The general conclusion is that the system is working fine in the basic scenario (720p at 15 fps, 30 fps if we relax rendering real-time conditions). Also, delay measures at the production segment (RTT and motion-to-photon) are slightly underperforming, especially under load, but functional. Delivery KPIs are not fully met either, but they are quite stable even under load. Regarding the functional validation of the automatic slice change, a more detailed analysis of the performance of this functionality will be done in the second phase of the project.

It should be noted that some limitations have been found during the tests, mainly in mmWave components availability. However, even though there is not much scope to further include more improvements (due to the unavailability of SA devices for FR2), in the second phase some potential improvements related to the production console and the media delivery will be implemented. Therefore, the strategy in Phase 2 will be to test the different parts of the chain as far as possible, even with its existing limitations, so that we can establish the requirements for the next generation of systems: radio links to support 1080p, 9 cameras; GPUs to support 1080p, 30 fps; and automatic slicing, including RAN.

References

5G-RECORDS Deliverables

[1] 5G-RECORDS, "Use Cases, requirements and KPIs", Deliverable D2.1, 5G-PPP 5G-RECORDS project, 01/2022.
[2] 5G-RECORDS, "First Description of 5G components", Deliverable D3.1, 5G-PPP 5G-RECORDS project, 07/2021.
[3] 5G-RECORDS, "Integration of 5G components", Deliverable D4.1, 5G-PPP 5G-RECORDS project, 07/2021
[4] 5G-RECORDS, "Testbeds and Trials Roadmap", Deliverable D5.1, 5G-PPP 5G-RECORDS project, 08/2021.

External links

[5] "AW2S RU" [Online] https://www.aw2s.com/RRU.html, Retrieved 02/2022
[6] "Simcom module" [Online], https://www.simcom.com/product/SIM8200EA_M2.html, Retrieved 02/2022
[7] "Quectel module" [Online], https://www.quectel.com/product/5g-rm500q-gl, Retrieved 02/2022
[8] "JT-NM ST2110 Test Program March 2020" [Online] https://static.jt-nm.org/documents/JT-NM_Tested_Catalog_ST2110_Full-Online-2020-05-12.pdf