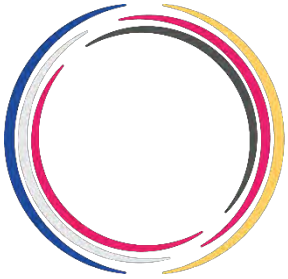

5G OPERA

5G-OPERA Deliverables 4.2 + 4.3 + 4.4

Open RAN components, interfaces and innovations developed in 5G-OPERA

1 Executive summary

Workpackage 4 of the 5G-OPERA project brought together a consortium of industry and research partners—AW2S, CEA-Leti, Ekinops, Eurecom, Firecell, Fraunhofer HHI, Kalray, NXP Germany, TU Berlin, and Xelera Technologies—to advance Open RAN (O-RAN) and 5G technologies through innovative contributions across hardware, software, and system integration.

CEA-Leti focused on Integrated Access and Backhaul (IAB) 5G systems in the FR2 band (26 GHz). Leveraging a Hardware-in-the-Loop (HIL) implementation, they explored interference management using Reconfigurable Intelligent Surfaces (RIS) and demonstrated a mixed hardware/software solution, laying the groundwork for upcoming IAB demonstrations.

Ekinops developed boundary clock support for Ethernet synchronization across O-RAN nodes using IEEE1588v2 PTP and SyncE. They adapted existing Layer 2 switch designs to transport eCPRI over fronthaul networks, enabling synchronized frequency, phase, and time distribution.

Eurecom enhanced the OpenAirInterface (OAI) project to support O-RAN and 3GPP standards, enabling hardware disaggregation, virtualization, and interface integration. Key developments include support for O-RAN 7.2, F1, E1, and E2 interfaces, alongside hardware acceleration for Layer 1 tasks and integration with Service Management and Orchestration (SMO).

Firecell collaborated with Eurecom to design and commercialize the Firecell Orion Labkit O-RAN, a user-friendly solution for private 4G/5G networks. This labkit supports COTS user equipment and enables rapid development and testing of customized O-RAN networks.

Fraunhofer HHI optimized OAI Radio Access Network (RAN) components, integrating multiple Quality of Service (QoS) classes and enhancing compatibility with various RUs and core networks. They also implemented near-real-time (FlexRIC) and non-real-time RAN controllers, providing end-to-end setups to evaluate system features and control loops.

Kalray developed a hardware acceleration solution for the Distributed Unit (DU) using its MPPA Coolidge v2 processor, featuring high-speed interfaces and 5G-specific accelerators. Their PCIe-based Kalray K300 acceleration card demonstrated significant performance improvements in Open RAN deployments.

NXP Germany enabled the implementation of proprietary CU, DU, and RU stacks, using its Layerscape® Access processors, while supporting open interfaces for integration with partner solutions like Xelera's DU acceleration. NXP's CU/DU board incorporates ARM NEON SIMD optimizations, hardware accelerators, and high-speed Ethernet interfaces. Additionally, the Glasgow board was adapted for a 2T2R Radio Unit (RU) with eCPRI compatibility, showcasing seamless integration of open-source 5G stacks with ARM-based platforms.

TU Berlin developed a radio head for the Radio Unit (RU) in NR258 band, featuring a 4x32 cylindrical antenna array with beamforming and low-PHY processing. The design adheres to the 7.2 split and integrates with other O-RAN components.

Xelera Technologies partnered with NXP to create an Open-RAN-compliant DU leveraging hardware acceleration. Their solution reduced costs and energy consumption while supporting multiple RUs and hundreds of devices per DU, demonstrating the scalability of the O-RAN architecture.

This collaborative effort demonstrates significant progress in advancing the technical capabilities, interoperability, and scalability of 5G networks, setting the stage for future innovations in the telecommunications industry.

Table of Contents

Abbreviations

IAB	Integrated Access and Backhaul
PA	Power Amplifier
NG-RAN	Next Generation Radio Access Network
CU	Centralized Unit
DU	Decentralized Unit
NR	New Radio
gNB	Next Generation NodeB
UE	user Equipment
MT	Mobile Termination
RRC	Radio Resource Control
NAS	Non-Access Stratum
SA	StandAlone
E-UTRAN EN-DC	New Radio - Dual Connectivity
EUTRA	Evolved Universal Terrestrial Radio Access
MeNB	Master NodeB
SgNB	Secondary NB
PHY	Physical layer
AP	Application Protocol
HW	Hardware
SW	Software
BF-OFDM	Block Filtered Orthogonal Frequency Division Multiplexing
RF-SoC	Radio Frequency System on Chip
FPGA	Field Programmable Gate Array
HIL	HW in the Loop
SDR	Software Defined Radio
DL	Down-link
UL	Up-link

2 Introduction

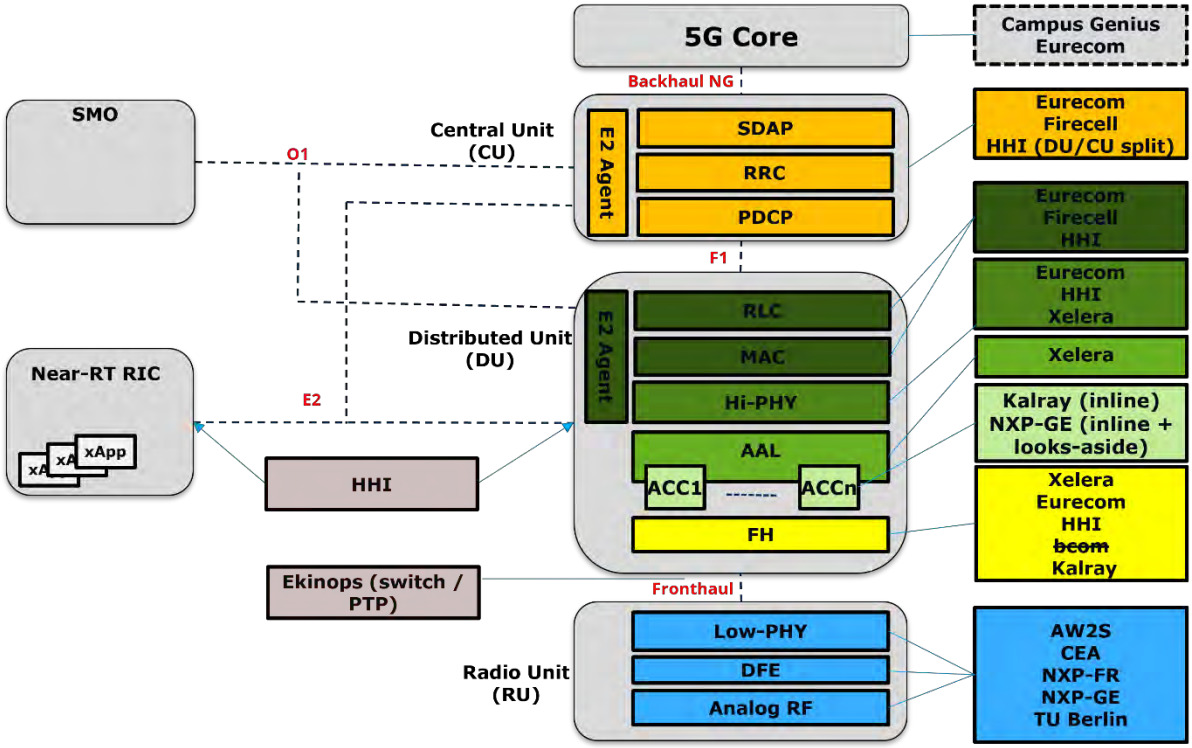

Figure 1: WP4 partner contributions

As part of workpackage 4 of the 5G-OPERA project the partners listed on the right side of Figure 1 contributed with the development of their respective Open-RAN components as shown in the same diagram. This document presents each partner's final results of the entire workpackage.

This document contains the results for the following three deliverables:

- D4.2 Development of Required Interfaces
- D4.3 Development of Open RAN components (including acceleration)
- D4.4 RAN-specific Implementation of Innovations from WP5

Since those deliverables are very much intertwined, it was decided to combine them in this single document. This is intended to be more convenient for the contributors and readers as it significantly reduces redundancy.

3 Releases

The following list represents the results generated by the different partners in workpackage 4 including the integration of results from other workpackages.

3.1 CEA-Leti

3.1.1 Objective and status

CEA aims to explore and demonstrate the capabilities of Integrated Access and Backhaul (IAB) 5G system in the FR2 frequency band (26 GHz), eventually relying on a Hardware in the Loop (HIL) implementation. The objective is to study, simulate, implement and test various functionalities of the IAB 5G system, the status of which – at the date this deliverable is written – is provided below:

Functionality	**Analytical study**	**Software only simulation**	**HIL implementation**
Time and frequency synchronization			X
Synchronization Signal Block (SSB) detection			X
Reconfigurable Intelligent Surfaces (RIS)		X [1]	
Beam discovery		X	
Digital precoding	X [2]	X [2]	
Interference measurements	X [2]	X [2]	
Beam sweeping	X*	X*	

* Patent pending

The interference measurement studies undertaken within the scope of project (reference 1) facilitated the characterization of interference between sources and between users when a Reconfigurable Intelligent Surface (RIS) is utilized for Integrated Access and Backhaul (IAB) networks. The influence of receiver-to-transmitter feedback on digital precoding was explored in reference 2. This deliverable provides a detailed description of the mixed hardware/software implementation, which is a critical step towards the upcoming IAB demonstration.

3.1.2 Equipment Used

Devices

- PC with Ethernet interfaces for TCP/IP communication
- ZCU111 (Zynq UltraScale+ MPSoC)
- Digital-to-analog converters (DAC) and analog-to-digital converters (ADC)
- Horns antennas
- 26 GHz up and down-converters.

Software

- Home-made Graphical User Interface (GUI) for waveform generation and reception (MATLAB)
- MATLAB 5G toolbox.

3.1.3 Waveform and filters

Waveform

We generate waveforms using the MATLAB 5G Toolbox, specifically targeting FR2 frequencies. The sampling frequency employed is Fs=122.88 MHz, with a corresponding bandwidth B=68.52 MHz. We generate a 5G frame with the 64 Synchronization Signal Blocks (SSBs) (case D, numerology 3 – i.e. SubCarrier Spacing (SCS)=120 kHz).

Figure 2 shows the spectrogram and the power spectral density (PSD) of the transmitted signal (first 4 ms, where the SSBs are located). The SSBs are used for beam discovery and tracking process: each assessed beam at the Tx is associated to a given SSB; the Rx measures the SSBs power and correlation magnitude and sends back to the Tx the best one identified.

The spectrogram of the transmitter waveform illustrates the time-frequency representation of the signal. The vertical axis represents frequency (MHz), while the horizontal axis represents time (ms). The color intensity indicates the power spectral density (dB/Hz), with warmer colors showing higher power levels. The periodic patterns and distinct frequency bands are indicative of the signal's structure and provide insights into the signal's temporal and spectral characteristics.

The vertical axis of the PSD figure represents the power spectral density in dB, while the horizontal axis represents frequency in Hz. The plot shows the distribution of signal power across different frequencies, highlighting the spectral components of the signal. The PSD indicates the presence of significant signal energy within the expected frequency bands, which is crucial for assessing the signal's quality and identifying any potential interference or distortions.

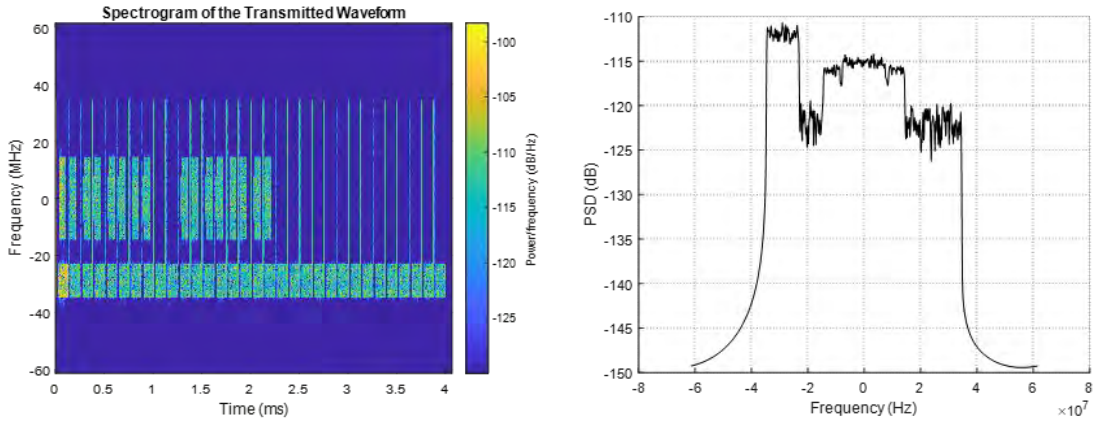

Figure 2: The spectrogram and the power spectral density (PSD) of the Tx signal

Filters

The following figures show the spectrum of the Tx and Rx filter at their sampling frequency. The spectrum of the Rx filter is shown before down-sampling.

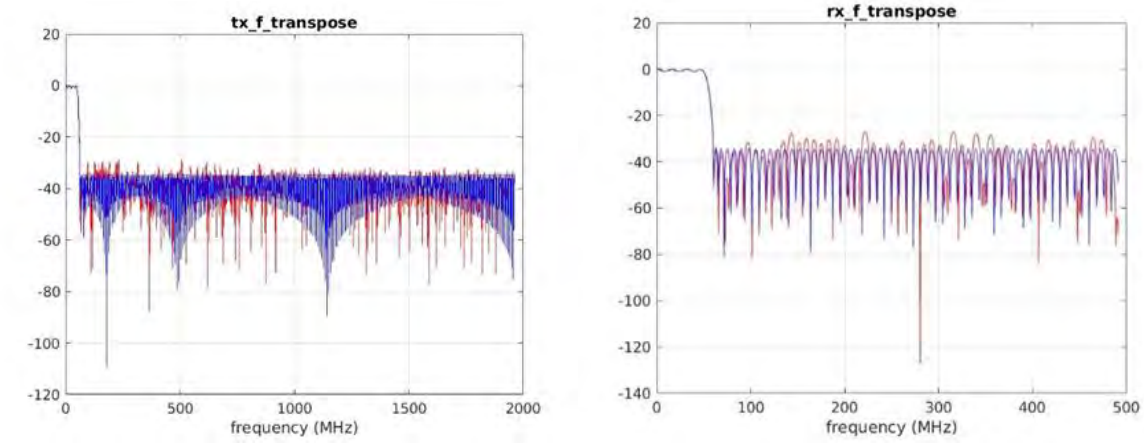

Figure 3: The spectrum of the Tx and Rx filter

3.1.4 The Hardware in the Loop (HIL) implementation

The scheme of the final targeted HIL implementation is shown in Figure 4.

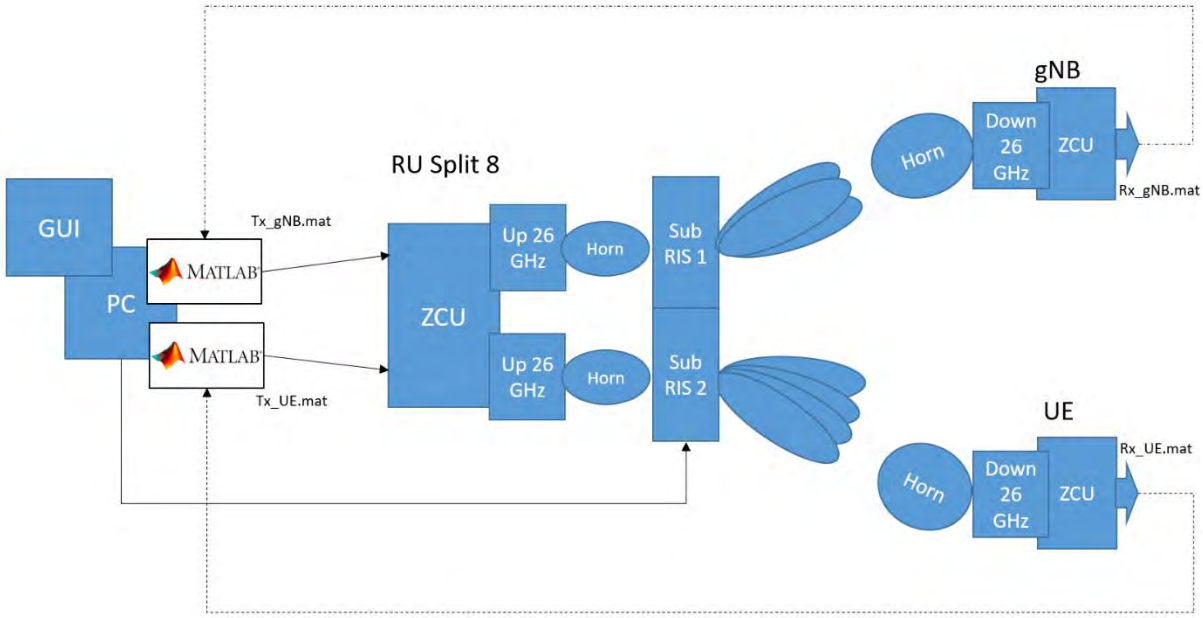

Figure 4: Final targeted HIL implementation scheme

In order to test the different equipment, we first implemented a one-Tx one-Rx system.

The transmitter setup shown in Figure 5 begins with an RF-SoC (Radio Frequency System on Chip) that handles baseband signal processing, modulation, and Intermediate Frequency (IF) upconversion. The digital IF signals generated by the RF-SoC are converted into analog signals by a Digital-to-Analog Converter (DAC). These analog signals are then upconverted to the target frequency of 26 GHz using a 26 GHz upconverter. This signal path ensures that the high-frequency 26 GHz signal is accurately generated. The receiver setup, also depicted in the image, is designed for signal reception and downconversion. The incoming 26 GHz signal is first downconverted to an IF using a 26 GHz downconverter. This IF signal is then fed into an RF-SoC, which performs the necessary demodulation and baseband signal processing. An Analog-to-Digital Converter (ADC) within the system converts the analog IF signal into digital form, facilitating further processing and analysis by the RF-SoC. The receiver configuration ensures reliable and precise signal processing, leveraging the RF-SoC's advanced capabilities to handle high-frequency signals effectively. The setup is streamlined with wired connections, ensuring signal integrity and minimizing losses.

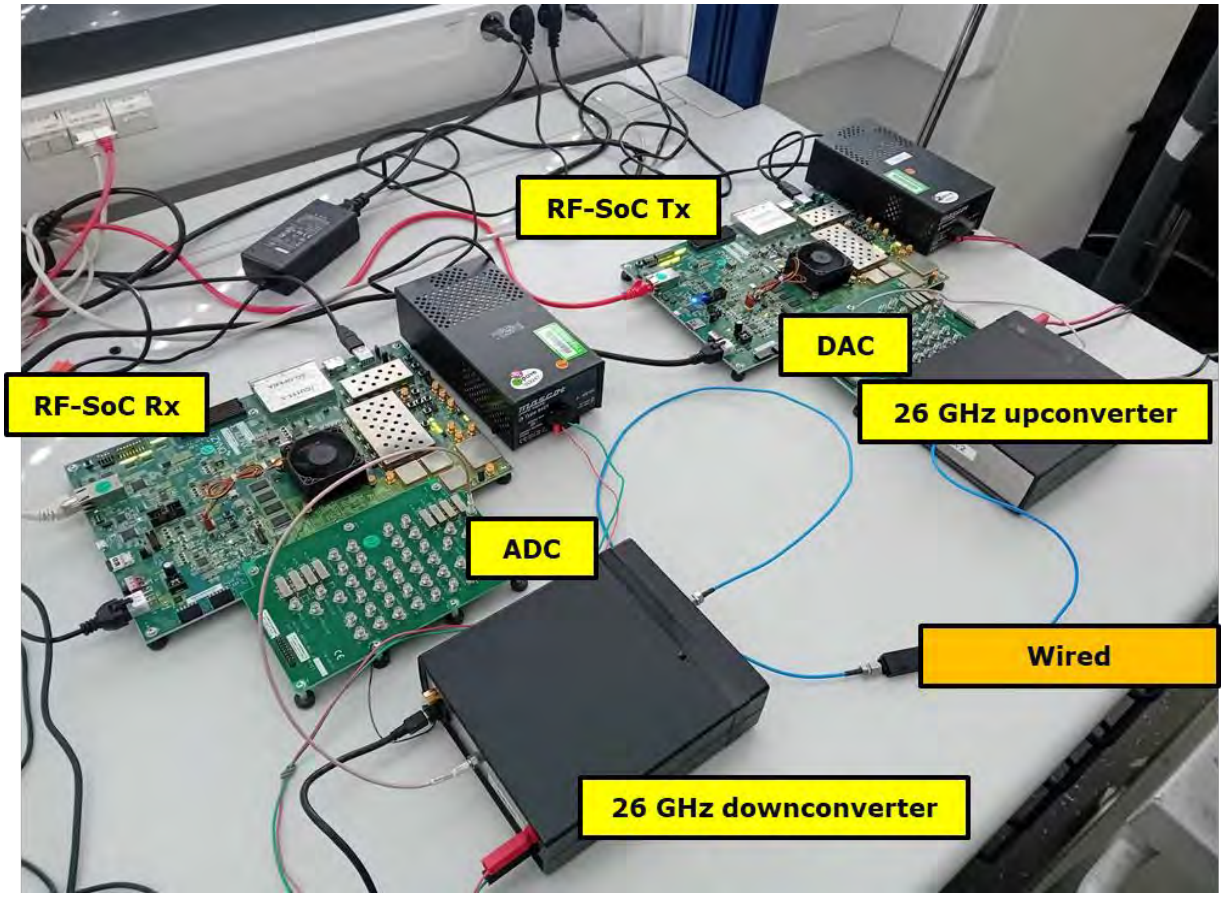

Figure 5: The first setup

Figure 6 presents key synchronization metrics. The time asynchronization of 2.793 ms and frequency asynchronization of -9.444 kHz highlight the necessary timing and frequency corrections. The detection of the best SS block (here set at index 17), with a high correlation magnitude, indicates a strong synchronization signal, crucial for robust signal alignment and decoding.

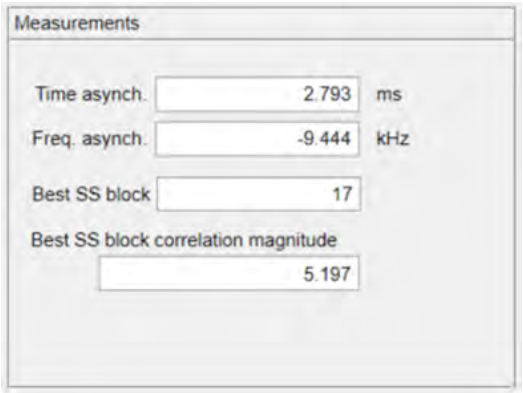

Figure 6: Time and frequency asynchronization, and the best SS block detection

Figure 7 shows the spectrogram and the PSD of the received signal.

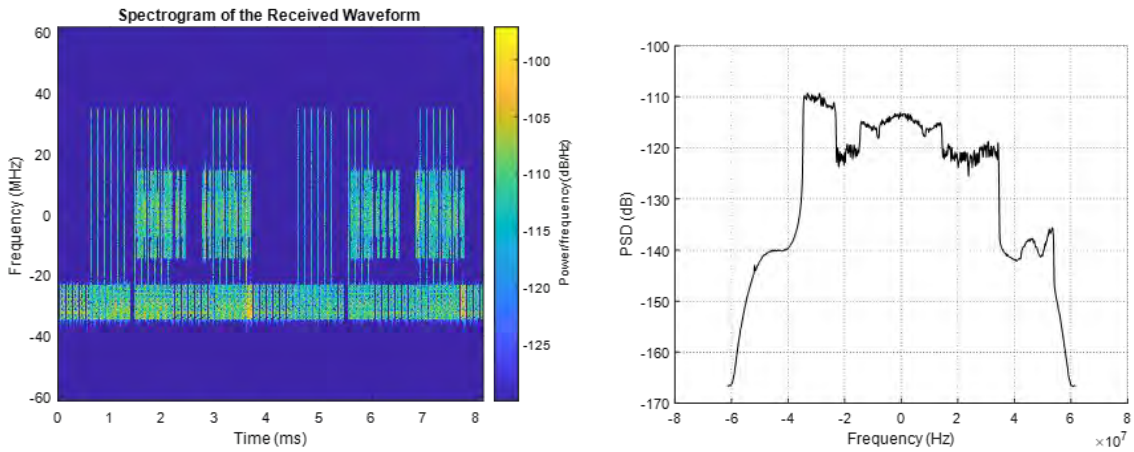

Figure 7: The spectrogram and the power spectral density (PSD) of the received signal

Figure 8 presents the equalized PDSCH constellation diagram, displaying distinct and well-clustered points indicative of a successful equalization process. The clear constellation points suggest a 16-QAM with minimal noise and interference, ensuring high-quality signal reception and accurate data demodulation.

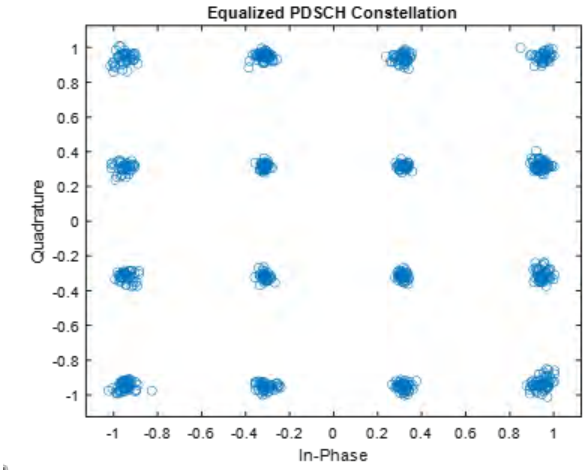

Figure 8: The received signal after equalization

In the second setup, we kept one transmitter and one receiver, and we replaced the wire with horn antennas, as shown in Figure 9.

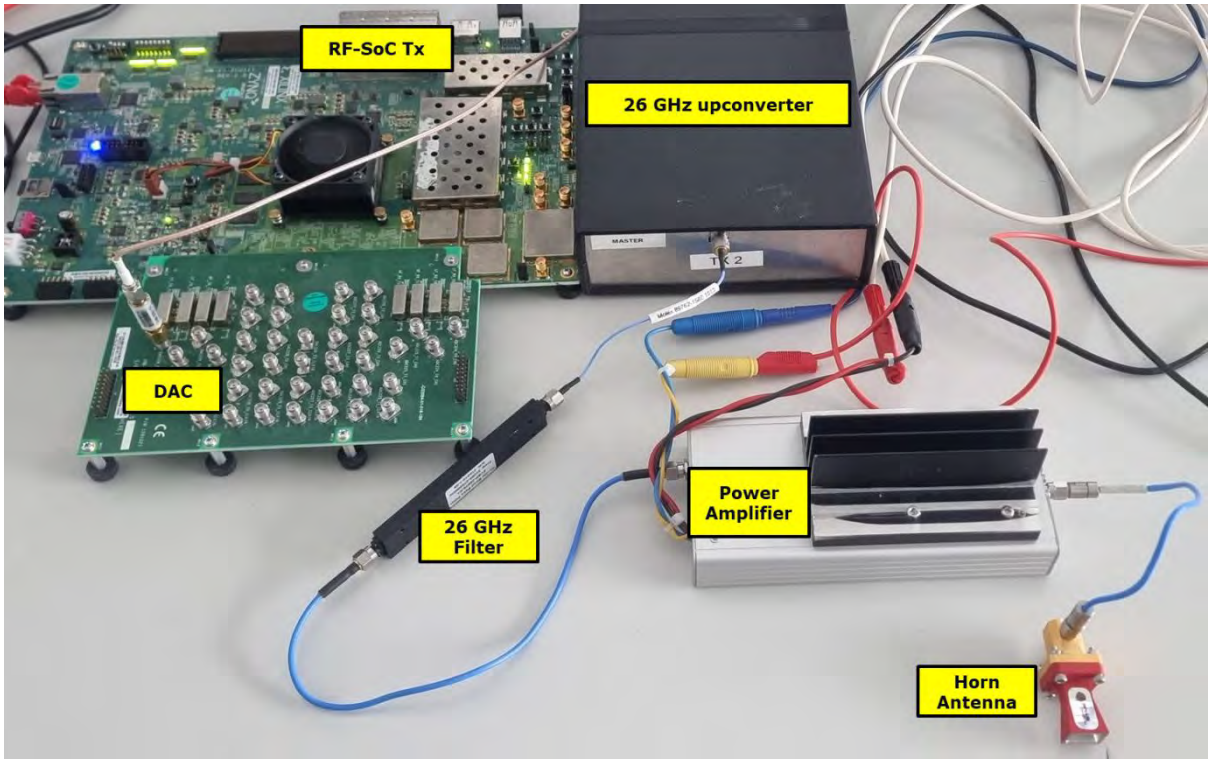

Figure 9: The transmitter (Tx) of the second setup

The transmitter setup consists of the same RF-SoC as in the previous setup. The digital baseband signals generated by the RF-SoC are converted into analog signals by the DAC. These analog signals are then upconverted to the target frequency of 26 GHz using the 26 GHz upconverter. To ensure signal purity, a 26 GHz filter is employed, which removes any unwanted frequencies, allowing only the desired 26 GHz signals to pass through. The filtered signal is then amplified by a power amplifier to achieve the necessary power level for transmission. Finally, the amplified 26 GHz signal is radiated into free space by a horn antenna.

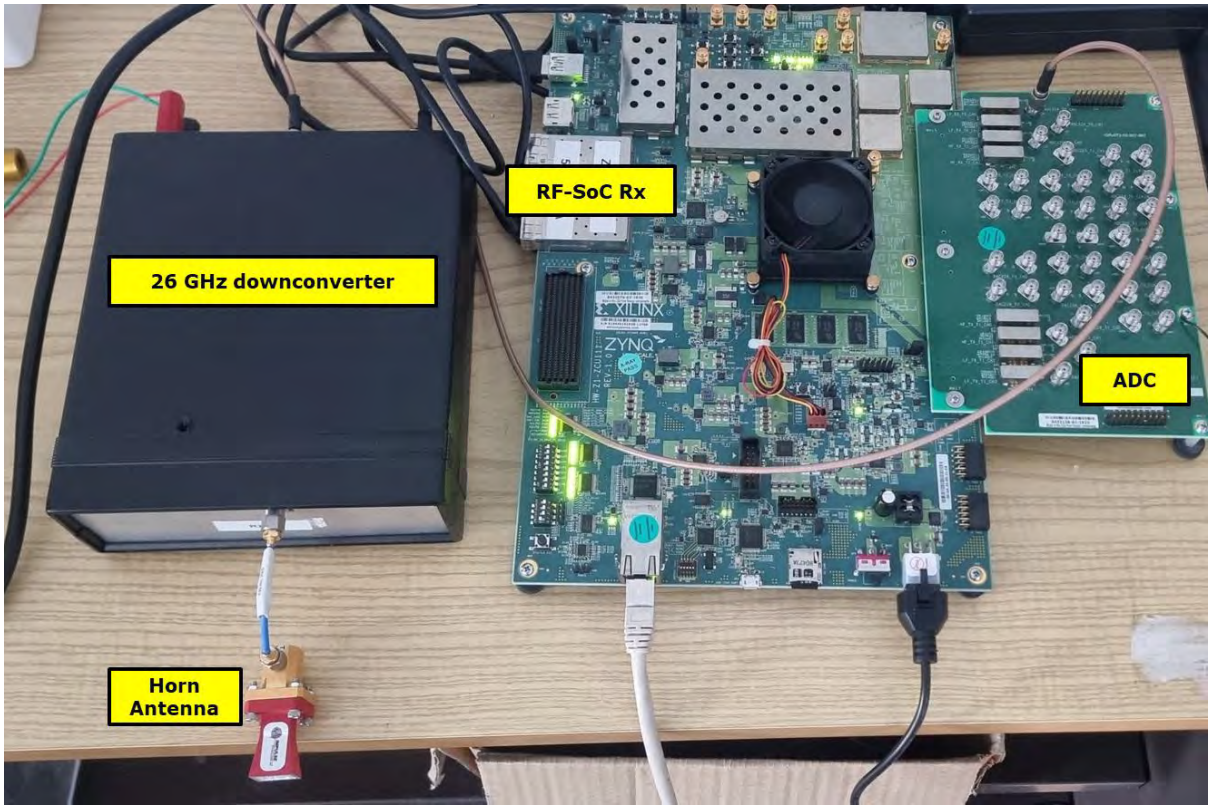

Figure 10: The receiver (Rx) of the second setup

The receiver setup mirrors the transmitter setup with components configured for signal reception and downconversion. The received 26 GHz signal is captured by a horn antenna and then downconverted to an IF using a 26 GHz upconverter. This downconverted signal is fed into the RF-SoC, which performs the essential demodulation and baseband signal processing tasks. An ADC within the RF-SoC converts the analog IF signal into digital form, facilitating further processing and analysis by the RF-SoC. This configuration allows for efficient reception and processing of the transmitted 26 GHz signals, ensuring reliable communication between the transmitting and receiving ends.

Figure 11 presents key synchronization metrics. The time asynchronization of 0.6747 ms and frequency asynchronization of 5.026 kHz highlight the necessary timing and frequency corrections. The detection of the best SS block (here set at index 0).

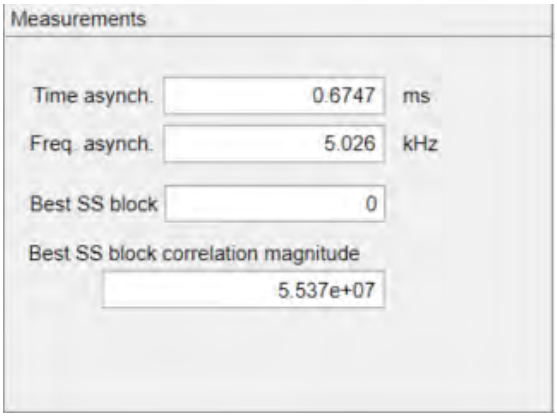

Figure 11: Time and frequency asynchronization, and the best SS block detection

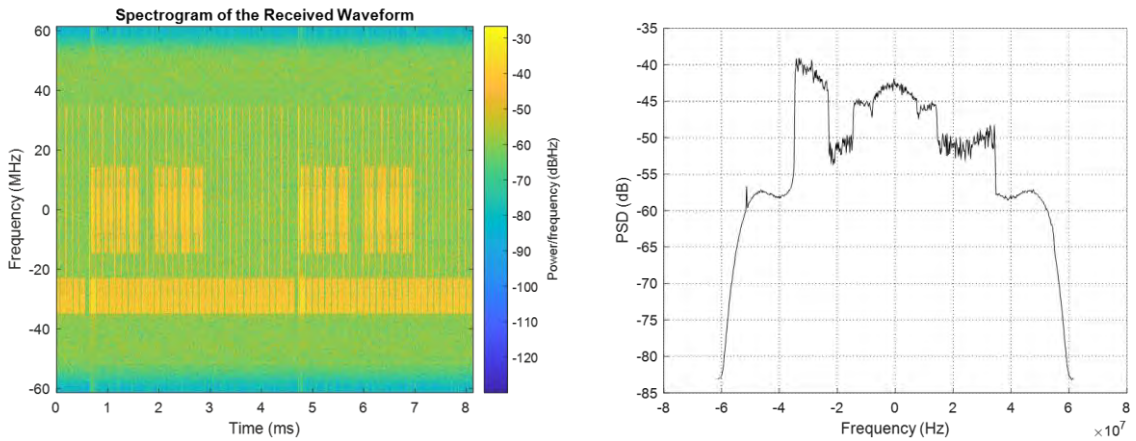

Figure 12: The spectrogram and the power spectral density (PSD) of the received signal

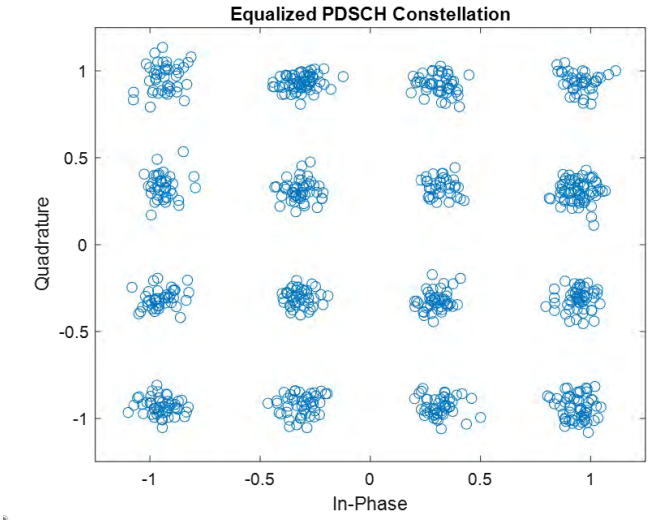

Figure 13: The received signal after equalization

In the third setup, we add a second transmitter and one receiver, as shown in Figure 14 and Figure 15.

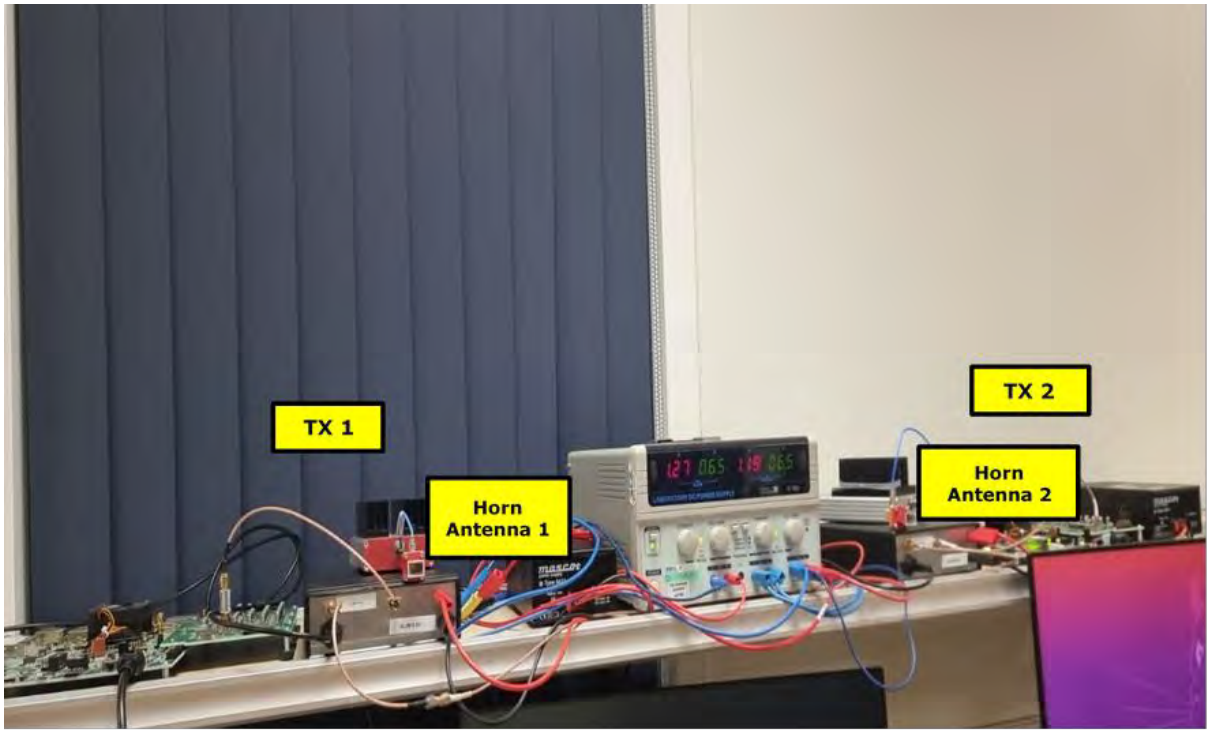

Figure 14: The transmitters (Tx1 and Tx2) of the third setup

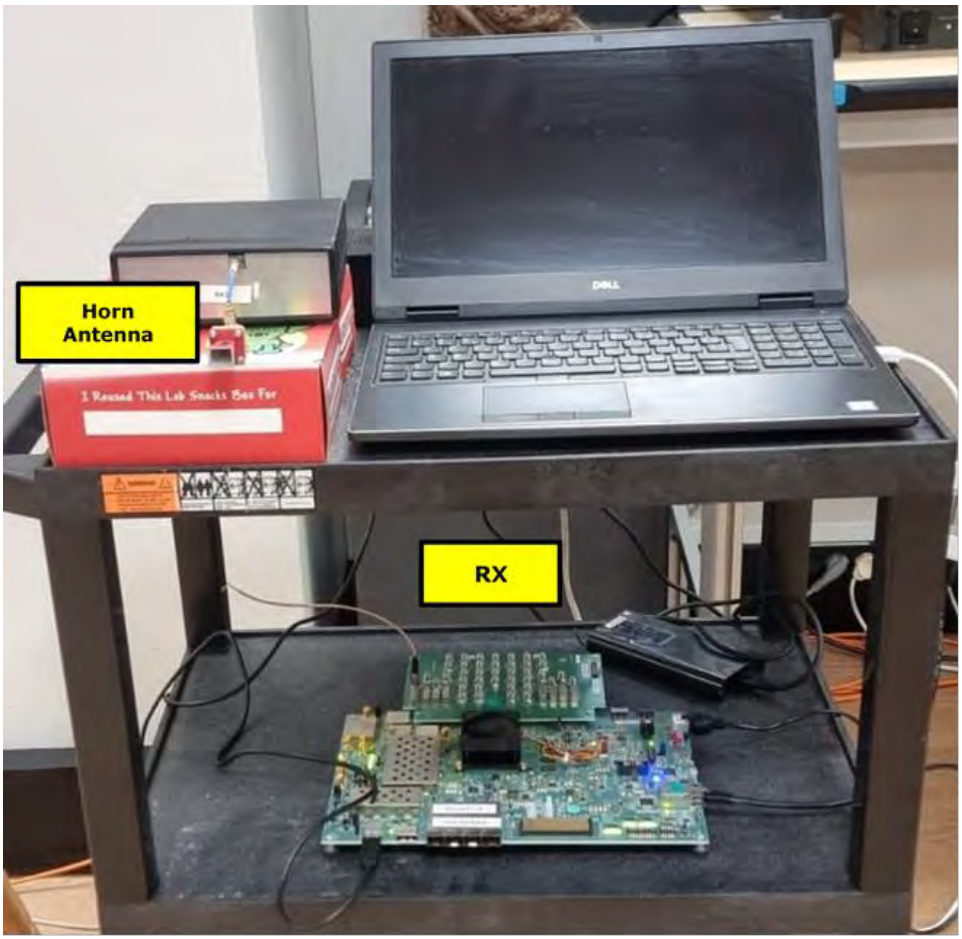

Figure 15: The receiver (Rx) in the third setup

In the next figures, we present the spectrogram and the equalized PDSCH constellation diagram of the received signal in the case of same numerology (SCS = 120KHz) for Tx1 and Tx2. We can see that, even with high interference we still can detect the best SSB (here set at index 0).

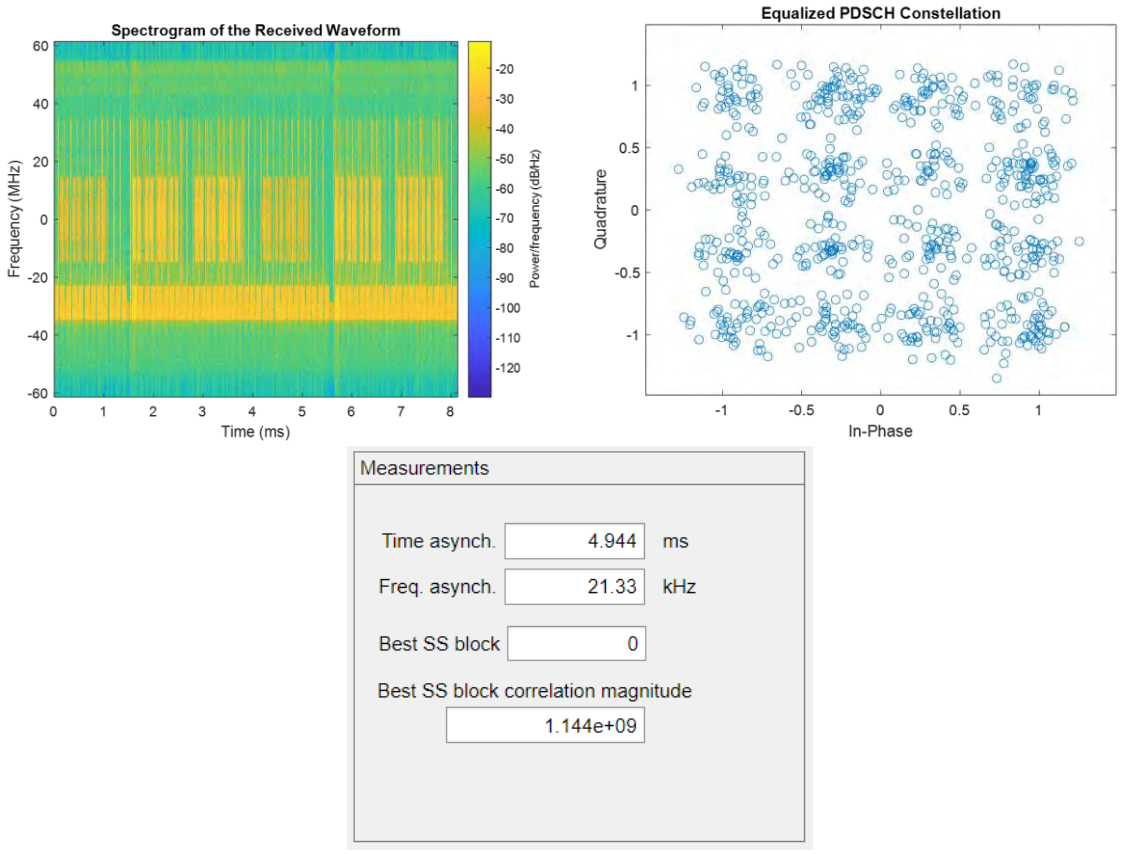

Figure 16: Time and frequency asynchronization, and the best SS block detection and equalized PDSCH of the received signal in the case of same numerology

In the case of different numerology, we present first the case where the SCS of Tx1 is 120KHz and the SCS of Tx2 is 60KHz, then the opposite (i.e. SCS = 60KHz for Tx1 and 120KHz for Tx2).

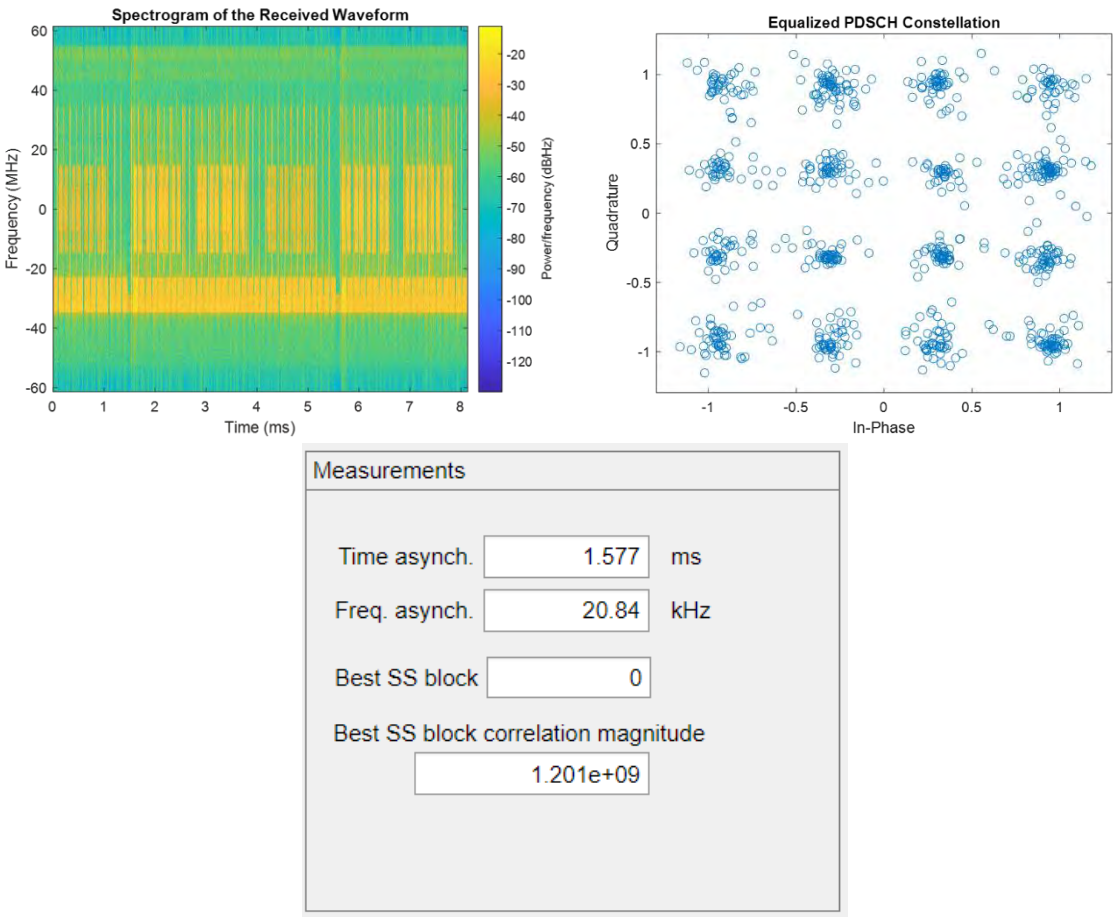

Figure 17: Different numerology Tx1 120KHz and Tx2 60KHz

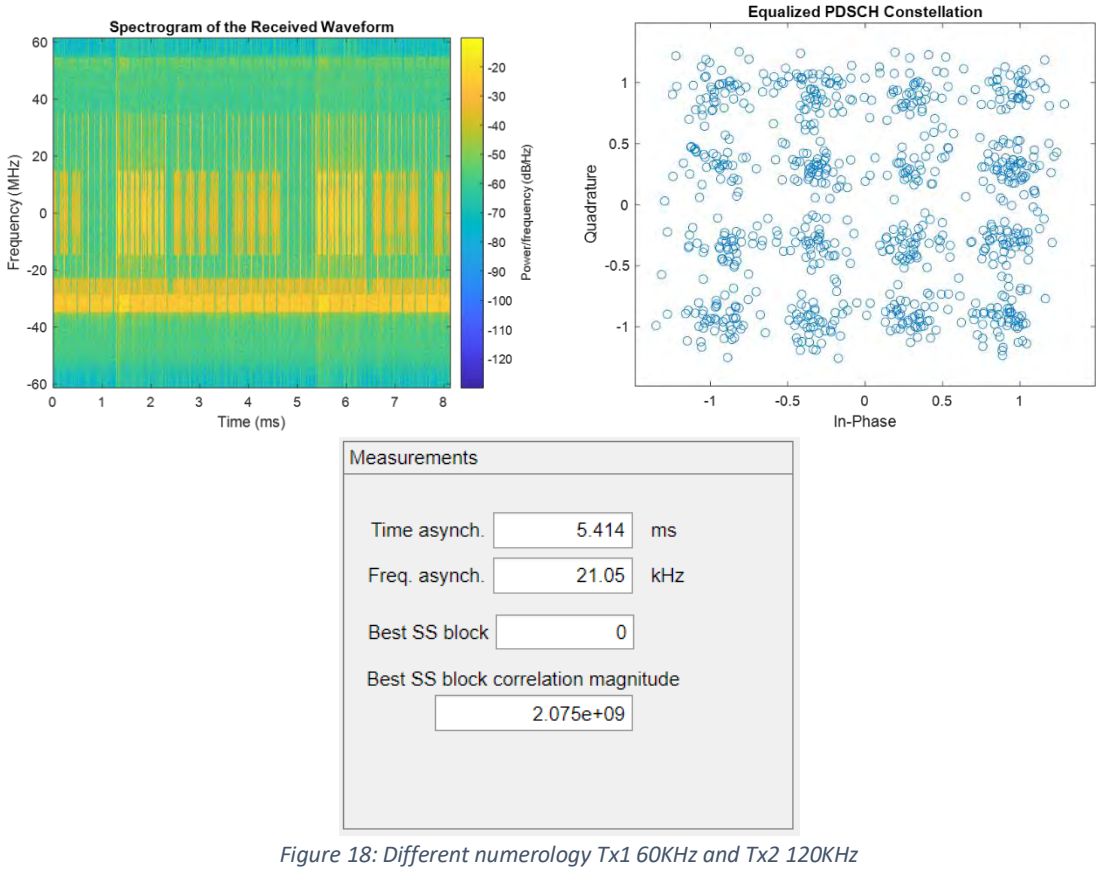

Figure 18: Different numerology Tx1 60KHz and Tx2 120KHz

3.2 Ekinops

3.2.1 Introduction

The contribution of Ekinops is the development of boundary clock support (based on IEEE1588v2 PTP/SyncE) on a L2 switch platform to synchronize all O-RAN nodes through Ethernet. The purpose is to transport eCPRI from the O-DU over a L2 fronthaul network to multiple O-RUs and use the same network for frequency, phase and time synchronization of the different O-RAN nodes (see 5.1.4 of [1]).

Considering the requirements for such implementation, Ekinops adapted an existing switch design and extended the functionality with SyncE and IEEE1588v2.

3.2.2 Hardware description

The new design is based on the existing Ekinops 1651 Ethernet Access Device and has the following specifications:

Data interfaces:

- 6x 1/10G SFP+ (**against 4 on original switch**)
- 6x 1G COMBO (UTP or SFP)
- 4x 1G SFP
- 4x 1G UTP

Other interfaces:

- Console port (RJ-45, RS-232)
- TOD in/out (RJ-45, RS422), fixed or programmable **(new)**
- 1PPS in/out (SMA-F), fixed or programmable **(new)**
- CLK (10Mhz) in/out (SMA-F), fixed or programmable **(new)**

Timing support hardware:

- Hardware-based timestamping in switch
- SyncE support **(new)**
- PLL/Buffer **(new)**
- Synchronizer **(new)**
- TCXO (4,6ppm accuracy) **(new)**

Other hardware (existing):

- 19” Chassis, 1 RU high, 250mm depth
- Dual modular power supplies (redundant,hot swappable, DC or AC)
- Reset button
- Earth connection
- Temperature-controlled fans
- Temperature support: 0°C to 45°C
- Dying gasp support (50ms)

This hardware covers the use case described in par. 5.1.4 of deliverable 4.1.

The functional diagram of the SyncE/PTP enabled switch can be found below:

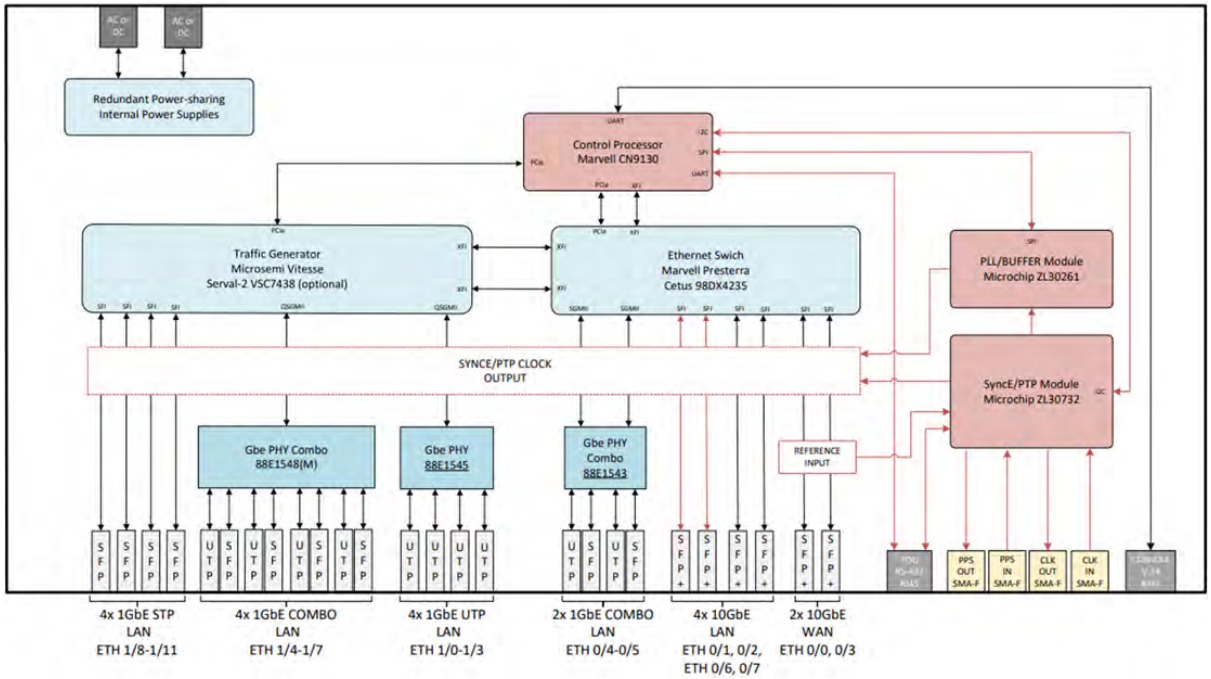

Figure 19: Functional Diagram SyncE/PTP switch for 5G O-RAN fronthaul

The functional diagram and interaction of the different timing components can be found below:

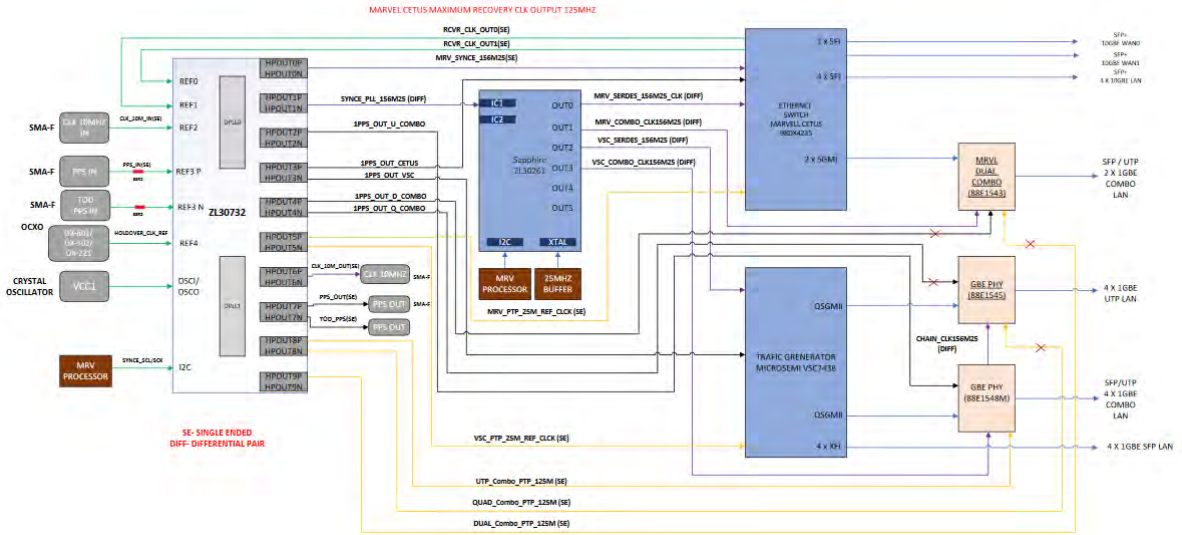

Figure 20: Functional diagram of timing recovery related components

A picture of the stuffed PCB of the SyncE/PTP switch can be found below:

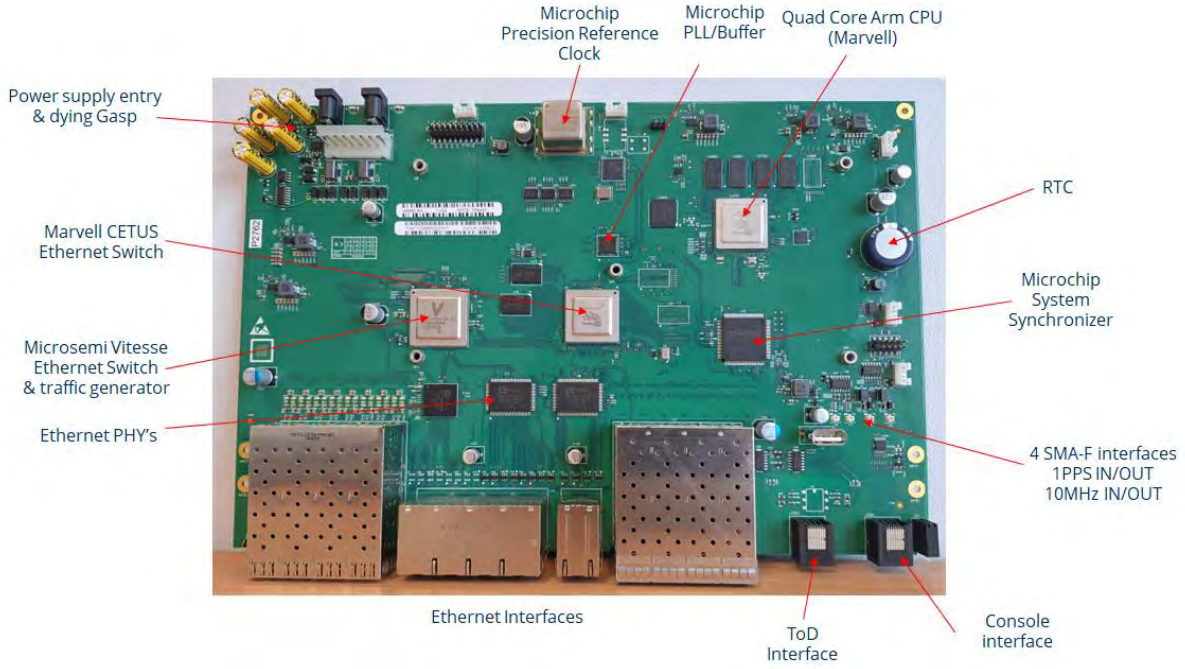

Figure 21: Stuffed PCB of the new SyncE/PTP switch

3.2.3 Timing-related software specifications

Synchronous Ethernet (SyncE)

- Support for 2 SyncE inputs (only 0/0 and 0/1)
- Timing info received from SyncE can be used as timing source and redistributed to other Ethernet ports (limited to 10G interfaces)
- Compliant with:
    - G.8261 Timing and Synchronization Aspects in Packet Networks
    - G.8262 Timing characteristics of a synchronous equipment slave clock
    - G.8264 Distribution of timing information through packet networks
- Support for Ethernet Synchronization Message Channel (ESMC)
- Possibility to configure SyncE source priority

IEEE1588v2 PTP with Boundary Clock support

- ITU-T G.8273.2 Class C
- ITU-T G.8265.1 frequency delivery profile
- ITU-T G.8275.1 phase and time synchronization profile , see ORAN-WG4-CUS.0
- Support for T-TC (Transparent Clock), T-BC/OC (Boundary Clock)
- Support for clocking scenario LLS-C3

These features are developed within the Ekinops OneOS6 networking operating software.

Both SyncE and IEEE1588v2 PTP features are configurable using the OneOS6 Command Line Interface.

The remaining part of the development of the switch is the testing of the SyncE and PTP functionality in a lab environment and deployment in a real testbed.

3.3 Eurecom

3.3.1 Introduction

EURECOM's main contribution to the 5G-OPERA project was the advancement of the OpenAirInterface (OAI) project to allow disaggregation and virtualization according to O-RAN and 3GPP standards, integration of different hardware acceleration capabilities, and general improvements of stability, key performance indicators and code quality.

The key developments in the 5G-OPERA project include

- Development of the O-RAN 7.2 interface to allow use of commercial O-RAN radio units (O-RUs) like VVDN, Benetel, Liteon, etc
- Development of F1-U and F1-C interfaces to allow disaggregation of RAN into DU and CU
- Development of E1 interface to allow further disaggregation of CU into CU-U and CU-C
- Integration of bb-dev interface in the DU to allow offload of LDPC coding and decoding to a look-aside accelerator card (T2 or Kalray)
- Integration of FAPI interface in the DU to allow offload of complete layer 1 to an in-line accelerator card (with e.g. Nvidia Aerial)
- Integration of the E2 agent and interface in DU and CU as well as several service models to allow interface with a near real-time RIC
- Integration of the O1 interface in the CU and DU to allow interaction with an SMO.

These developments are used by several partners in 5G-OPERA such as Firecell, NXP, Fraunhofer HHI, and Kalray. See the corresponding sections of this deliverable for details. In the following we report on the basic features and some performance metrics of basic OAI deployments.

3.3.2 OAI RAN features

The following text has been taken from [1].

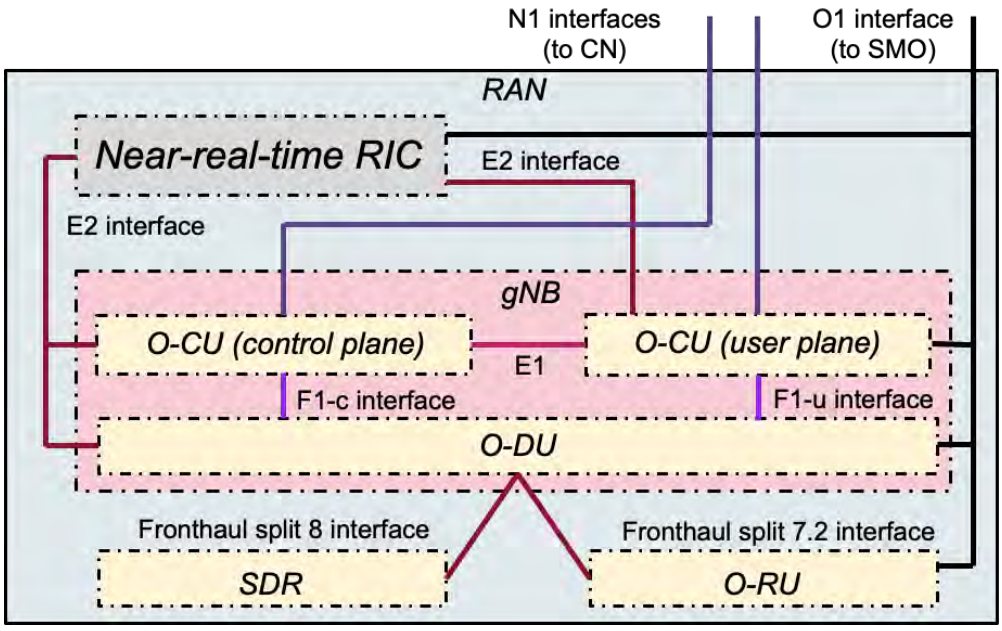

Figure 22: 5G RAN architecture including O-RAN interfaces

Figure 22 shows the basic architecture of a 5G RAN including O-RAN inter- faces. The OAI RAN project implements the CU and DU network functions together with all the necessary 3GPP and O-RAN interfaces. OAI also has an implementation of a RIC, which is described in detail in Section 2.2. The OAI RAN can be deployed with the full F1 and E1 splits, i.e., with CU-UP, CU-CP, and DU, or it can also be configured as a monolithic gNB.

The OAI DU implements RLC, MAC, and the PHY layers of the protocol stack. It can interface with SDRs like the USRP as well as commercial RUs either using eCPRI split 8 or the O-RAN split 7.2 fronthaul interface, respec- tively. OAI has been tested successfully with O-RUs from VVDN, Benetel, LiteOn, and Foxconn [2].

The PHY supports static TDD as well as FDD, subcarrier spacings of 15 kHz, 30 kHz, (FR-1), and 120 kHz (FR-2), cell bandwidths of 10, 20, 40, 60, 80, and 100 MHz, as well as 200 MHz (FR-2 only). We support 4-layer DL and 2-layer UL MIMO as well as highly efficient 3GPP compliant channel encoder and decoder (turbo, LDPC, polar).

The PHY can run entirely and in real-time on x86 architecture by exploiting its SIMD vector processing extensions. The layer 2 and above also runs on ARM architecture. The DU can also make use of different kinds of hardware accelerators, both in-line and look-aside. The in-line accelerator makes use of the FAPI interface between the PHY and the MAC layer. Currently, OAI supports the NVIDIA Aerial SDK Layer 1 [3] [4] but others will follow. The look-aside accelerator makes use of the Intel bb-dev interface and currently supports AMD Xilinx T1 and T2 cards [5] [6].

OAI implements both monolithic CU or disaggregated CUs (CU-UP and CU-CP). The CU-UP contains mainly the SDAP and PDCP layers. The CU- CP features the RRC layer, for radio resource management and UE lifecycle management. The communication between the two happens over the E1 interface. Both CU-UP and CU-CP can handle multiple DUs. In addition, F1 handover between multiple DUs is supported. The handover can either be triggered manually or based on neighbor cell measurements and events.

Both OAI DU and CU implement the O-RAN E2 interface to the near-RT RIC. The E2 agent supports service models KPM v2.03/v3.0 and RC v1.03. The E2 interface has been demonstrated successfully with the OSC RIC [7] [8], as well as with OAI's own implementation of the near near-RT RIC called flexRIC [9] [10]. OAI DU and CU also implement the O-RAN O1 interface to Operations and Management software such as ONAP [11] [12].

The OAI DU and CU are compatible with many different open-source core networks such as the Open5GS [13], free5GC [14], and of course the OAI 5G Core as well as many commercial core network solutions.

For more information about the OAI RAN project, please visit the readme on the OAI git repository [15].

3.3.3 OAI RAN KPIs

In this section we report some performance results with OAI. The configuration used is 100Mhz cell bandwidth, TDD pattern DDDFU, 256 QAM DL, 64 QAM UL. Measurements have been mostly done with USRP radios except for 4 layers DL, which is only supported with O-RUs.

DL measurements are aggregated over 2 users. In the DDDSU TDD configuration, single user DL TP is limited since we can schedule a maximum of 2 DL slots per user in one TDD period, as only 2 ACK/NACK feedback bits are available per user.

In the DL we get 400 Mbps with 1 layer and 800 Mbps with 2 layers. These values are very close to the maximum that can be achieved with this configuration. We have validated the basic operation for 4 layers as well, but the current performance is not significantly higher than 2 layers at the moment. This is currently being investigated and improved.

On the UL the maximum reported throughput for this configuration is 123Mbps for 1 layer which is very close to the theoretical maximum and 175M for 2 layers, which is about 70% of the peak throughput which is probably due to the MMSE MIMO receiver which is known to be sub-optimal.

3.3.4 Release notes

The source of the OAI RAN project can be found at https://gitlab.eurecom.fr/oai/openairinterface5g/. Docker images of CU and DU are also hosted on https://hub.docker.com/orgs/oaisoftwarealliance/repositories

3.4 Firecell

The objective of Firecell was to design and deploy a comprehensive O-RAN implementation. To achieve this, we collaborated with Eurecom to develop and integrate the various components necessary for a fully functional O-RAN solution. This section outlines the results of our efforts.

We have successfully finalized and begun the commercialization of our O-RAN product. The Firecell Orion Labkit O-RAN is an intuitive, ready-to-use solution that enables private and customized 4G/5G networks. Designed for development and testing, it supports commercial off-the-shelf user equipment (COTS UE) and delivers a seamless experience for users seeking an efficient and adaptable networking solution.

Firecell's O-RAN solution, epitomized by the Firecell Orion Labkit O-RAN, offers an innovative and accessible approach to deploying private 4G/5G networks. Leveraging the open and interoperable principles of the O-RAN architecture, this solution is designed to meet the needs of developers, researchers, and organizations seeking flexible and cost-effective wireless network deployment.

The Orion Labkit O-RAN delivers comprehensive support for modern 5G network requirements with the following capabilities:

- 5G Standalone (SA) Mode: Operates in full SA mode for seamless next-generation connectivity.
- High Bandwidth Support: Offers up to 100 MHz bandwidth with SISO and 2x2 MIMO configurations.
- Coverage: Provides a coverage range of 30 to 100 meters, suitable for various deployment scenarios.
- Performance: Achieves data rates of up to 800 Mbps downlink and 80 Mbps uplink.
- Real-Time Visualization: Monitors bit rates, resource utilization, and the number of connected devices in real time.
- SIM and eSIM Support: Fully compatible with physical SIM cards and eSIMs.

- Quality of Service (QoS): Enables configuration of higher-priority data streams to meet specific application requirements.
- REST API: Allows remote monitoring and control of the network for enhanced operational flexibility.
- Advanced RAN Architecture: Includes a Central Unit (O-CU) and Distributed Unit (O-DU) for a modular and scalable setup.
- Open-RAN Interface: Utilizes an Open-RAN 7.2 split interface between the O-DU and O-RU for interoperability.
- Precision Timing: Supports IEEE 1588 Precision Time Protocol (PTP) synchronization for accurate time alignment across the network.

Deployment Configurations

The Orion Labkit O-RAN provides flexible deployment options tailored to diverse network requirements, supporting two main architectures: the All-in-One Package and the CU/DU Split Architecture.

In the All-in-One Package, the system integrates a monolithic gNodeB (gNB) with the 5G Core Network (5G CN), combining all RAN and core functions into a single, streamlined setup. This configuration is ideal for straightforward deployments requiring a compact and cohesive solution.

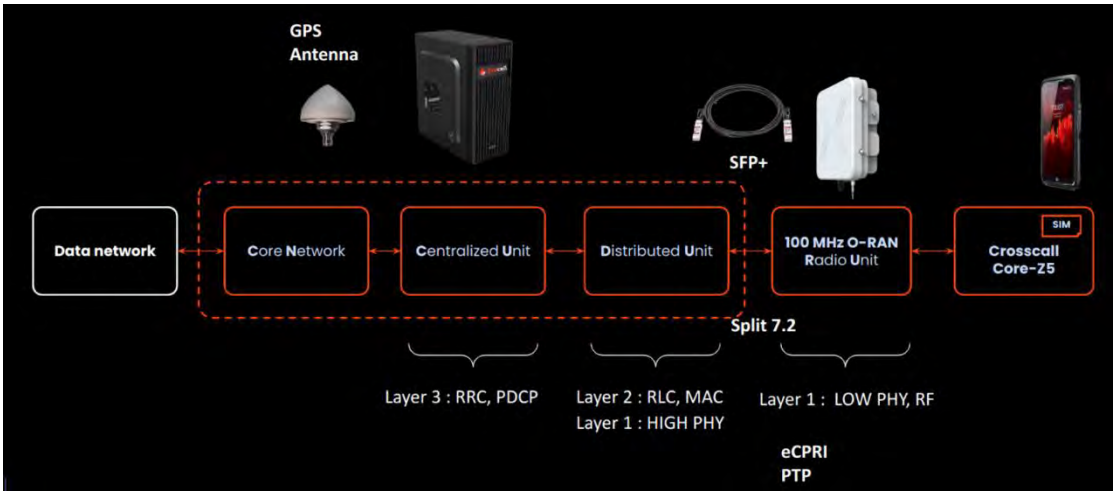

The CU/DU Split Architecture adheres to O-RAN standards by separating RAN functions traditionally housed within a baseband unit (BBU) into two distinct components. The Centralized Unit (CU) manages higher-layer functions, such as mobility and session management, while the Distributed Unit (DU) handles lower-layer functions near the radio. This separation enables scalable deployments, providing network operators with the flexibility to optimize centralization and edge processing for enhanced performance and adaptability.

To further accommodate specific network demands, the Orion Labkit O-RAN supports multiple deployment configurations within the CU/DU Split Architecture:

5G CN Only: This configuration deploys a standalone 5G Core Network connected to an independent CU, which interfaces with one or more DUs. It is suitable for networks requiring centralized core management with distributed data processing.

gNB CU Only: In this mode, a standalone CU is deployed to manage control functions and link to remote DUs and RUs, centralizing control while distributing radio and data processing units across the network.

gNB DU Only: This setup focuses on localized data processing by deploying a standalone DU. It is ideal for environments requiring ultra-low latency and rapid response times.

By offering these flexible deployment options, the Orion Labkit O-RAN empowers network operators to tailor their infrastructure to specific use cases and performance requirements, ensuring optimal functionality in a variety of scenarios.

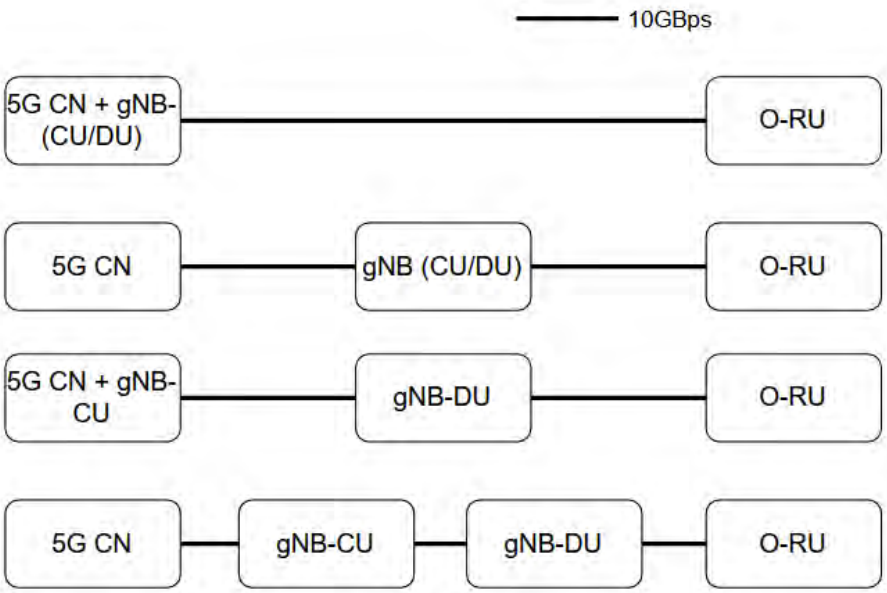

The Orion Labkit O-RAN supports the following fronthaul 7.2 APIs, enabling seamless communication and interoperability (Control Plane, User Plane and synchronisation).

3.5 Fraunhofer HHI

3.5.1 Summary

This section describes the work developed at Fraunhofer HHI within WP4, as part of Deliverables D4.2, D4.3 and D4.4, focusing on the development and interconnection of O-RAN components within the Radio Access Network (RAN), Core Network, and RAN controllers, including near- and non-RT RIC. Our work on RAN interfaces involved optimizing the OpenAirInterface (OAI) O-DU and O-CU implementations, adding support for multiple Quality of Service (QoS) classes, integrating the DU with various Radio Units (RUs) via the 7.2 split, and setting up integration of OAI CUs and DU-high with NVIDIA DU-low for PHY layer processing using NVIDIA ARC nodes. On the Core Network side, the OAI RAN components were fully integrated with the OAI-5G Core Network, as well as other core network implementations such as Open5GS, free5gc, and FOKUS. For near-real-time RAN monitoring and control, we integrated FlexRIC and O-SC RIC, along with necessary modifications to the E2 interface, implementing key features like RIC REPORT and RIC INSERT services to support mobility and resource management across the network. For service orchestration and control at longer time scales, we contributed to the implementation of the O1 interface, enabling integration with the non-RT RIC. Lastly, we describe the simple end-to-end O-RAN setups developed in our lab, providing an environment for evaluating and testing the interfaces and features implemented in WP4, as well as the initial evaluation of the monitoring and control loops developed in WP5.

3.5.2 RAN (DU/CU/RU)

The work on the RAN interfaces at Fraunhofer HHI included several modifications to the OAI-based O-DU and O-CU implementations. As previously reported in the Deliverable D4.1, the laboratory tests we performed with the OAI-based disaggregated DU and CU setup using F1 interface demonstrated worse performance compared to the monolithic gNB setup (i.e., coupled DU/CU). To address this issue, considerable effort was devoted to optimizing the system performance and introducing support for different QoS classes, which was lacking in the OAI codebase.

In the early phase of development, support was added to orchestrate the RRC Reestablishment procedures at the gNB before F1 Reestablishment takes over. To enable the QoS support in the CU-UP component, an initial version of the Service Data Adaptation Layer (SDAP) has been introduced to OAI software stack implementation. Based on the QoS Flow Identifier (QFI) that is signaled from the Core Network, the SDAP layer decides based on a mapping scheme to which Data Radio Bearer (DRB) a packet should be sent. Additionally, multiple DRB support was added to the monolithic version of the OAI gNB to enable this feature in such setup as well. The mapping scheme can be extended to support different QoS requirements, allowing appropriate handling of the Time-Sensitive Networking (TSN) and Ultra-Reliable Low Latency Communications (URLLC) classes of importance to the project.

At the MAC layer, further developments were made to support prioritized scheduling of specific users. Specifically, from the Core Network side, there is a 3GPP parameter Allocation and Retention Policy (ARP) Priority, i.e., “ARP Priority”. This parameter is used when the channel is congested, when the gNB can use it to decide on how to prioritize users. In this case, the resources at the gNB are allocated dynamically to the user based on the parameter’s value. To set this threshold the OAI’s internal telnet module is used, of course this functionality can also be extended to be controlled by an xApp. For the full support of this feature, the upper layers also had to be adjusted to allow passing the “ARP priority” information to the Medium Access Control (MAC) scheduler. For the full support of this feature, the upper layers also had to be adjusted to allow passing this information to the MAC scheduler, which was successfully accomplished. With the mechanism in place, the system can guarantee that specific resources will be allocated to a user with a given priority. Most of the additions to the OAI codebase

developed towards implementation of the features described here were already reported to OAI and merged into the public repository.

The remaining work on the RAN interfaces at Fraunhofer HHI was focusing on the integration of the DU with different RUs on the front-haul. Specifically, the Benetel Outdoor O-RU RAN650 [16] and the LiteON FlexFi Indoor O-RU [17] have been integrated with the OAI DU using the 7.2 split.

Additional efforts were dedicated to integrating the OAI CUs and DU-high (MAC/RLC) with the NVIDIA DU-low (high-PHY) system [18]. In this configuration, the DU and CU manage signal processing and data plane functions, while PHY layer processing is handled by a system of eight NVIDIA ARC nodes running on Gigabyte E251-U70 servers. The communication between the DU-high and DU-low is facilitated through the 5G Functional API (FAPI), as defined by the Small Cell Forum (SCF).

3.5.3 Core network

The OAI RAN components (CU and DU) are fully integrated with the OAI-5G Core Network in our setup, enabling the implementation of a 3GPP-compliant 5G Standalone (SA) system. This integration allows us to provide a robust and realistic environment for testing and validating the performance and functionalities of 5G SA networks in alignment with the 3GPP specifications. In addition to the OAI-5G Core Network [19], we have also successfully tested a series of core networks like including the Open5GS Core Network [13], free5gc [14] and FOKUS core [20], further enhancing the flexibility and scope of our testing environment for the development of a comprehensive end-to-end evaluation testbed (in WP6).

3.5.4 Near-RT RIC

Following the identification and comparison of the open-source implementations of the Near-RT RIC in the Deliverable D4.1, we have focused on the FlexRIC implementation from the Mosaic5G project [21] and the reference implementation by the O-RAN Software Community (O-SC RIC) [22]. To support integration of the DU/CU with multiple near-RT implementations, we had to adapt the E2 interface which enables interaction with xApps, enabling the testing of network optimization and the evaluation of various xApp concepts developed in WP5.

FlexRIC

The E2 interface for FlexRIC supports Key Performance Metric (KPM) and RAN Control (RC) service models, which can be exploited by RAN monitoring and management xApps. While KMP service was successfully reported in the Deliverable 4.1, full exploitation of the RC service model required implementation of additional features at the E2 nodes in the OAI gNB and the FlexRIC codebases, according to the O-RAN alliance WG3 technical specifications [23]. In particular, additional features have been implemented for the RIC REPORT services using the Report Style Type 1, Message Copy, which enables reporting of the complete copies of Radio Resource Control (RRC) messages. These include the respective management of RIC subscriptions, RIC action definition (Format 1) and event trigger (Style 1). To illustrate the implemented mechanisms, the Fig. F1 shows the flowchart of the messages exchanged between the xApp and the gNB to realize the RIC REPORT service. During the runtime, when an RRC message is generated in the downlink or received in the uplink, the gNB checks whether the corresponding RRC message ID has been previously subscribed to and, if so, sends a copy of this message to the RIC via a RIC Indication. In addition to the copied RRC messages, the UE ID is reported alongside.

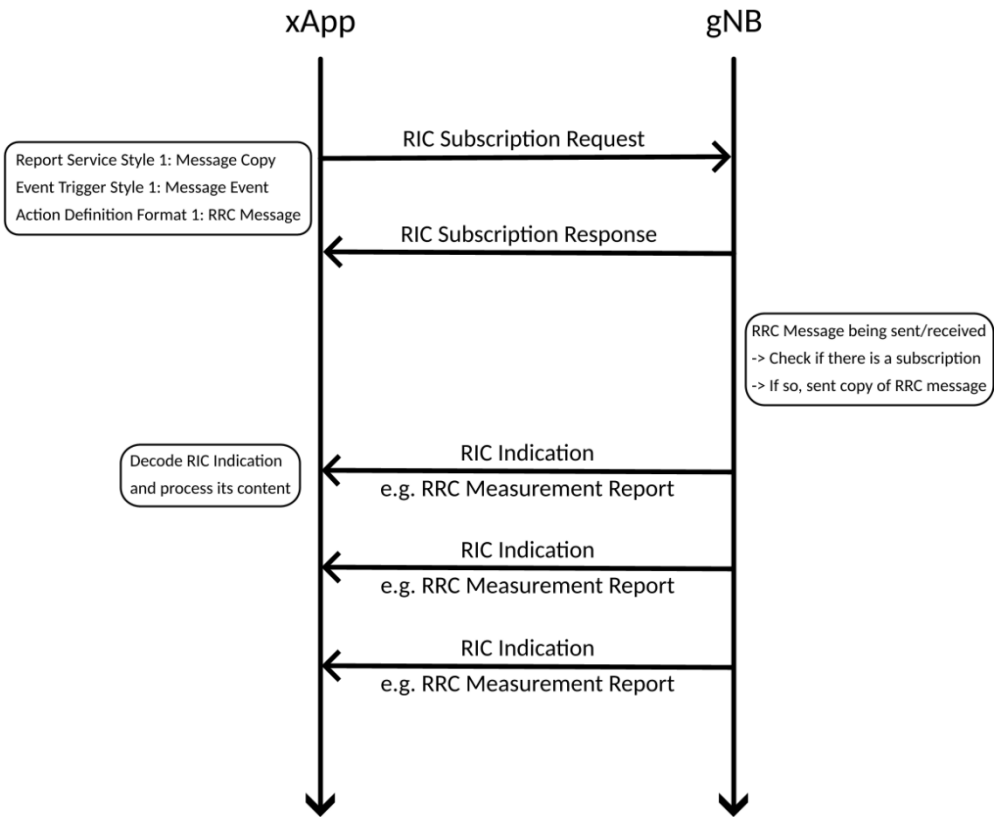

Figure 23: Mechanism for the RIC REPORT service of the RC service model using a subscription to RRC measurement report messages as an example

The implemented features allow xApps to subscribe to various types of RRC messages in the following channels:

- UL_DCCH_NR_RRC
- UL_CCCH_NR_RRC
- PCCH_NR_RRC
- DL_CCCH
- DL_DCCH.

On the near-RT RIC side, the xApp subscribes to the RRC messages of interest and receives the corresponding message copies. In addition to the subscription management, the message contents over the E2 interface are ASN encoded and, therefore, the necessary ASN decoding procedure has been implemented at the FlexRIC xApp side. Furthermore, the features added within the RAN were also added within the FlexRIC emulator to facilitate xApp development using these features. These functionalities were tested with specifically developed xApps, deployed within the OAI-based OTA lab setup.

In addition to the REPORT service features, at Fraunhofer HHI we have also implemented preliminary steps of the RIC INSERT service for the OAI gNB. The RIC INSERT service extends the abilities of an xApp to shape the behavior of the RAN with a completely different mechanism in which the RAN detects a previously defined trigger, suspends the associated procedure and prompts the Near-RT RIC for a RIC Control Request. The implemented preliminary features belong to the INSERT Service Style Type 3,

Connected Mode Mobility Request, with the indication for Handover Control Request and its corresponding RIC Action Definition (Format 3), i.e., Event Trigger (Style 2).

With these features in place, an xApp can actively control the handovers, transferring the allocated resources from one DU-RU path to another one. In such a scenario, the gNB sends a RIC Indication message for handover control. The Near-RT RIC and the xApp respond to this message with a RIC control request, thereby actively influencing the handover process. The remaining steps necessary for the full implementation of the RIC INSERT service are implementation of the suspension of the handover process when the event trigger is detected and the corresponding response to the RIC control request. These new features will be merged into the public OAI repositories, with the implementations related to RRC message copy being scheduled for merging at the beginning of 2025.

In addition to providing support for mobility management in the RAN, these newly implemented features also provide essential management mechanisms for RAN control loops that require transfer of data streams between different RAN components. As such, they are the key enablers for energy saving and load balancing xApps developed in WP5.1.

O-SC RIC

While FlexRIC was in the primary focus of near-RT RIC integration activities at Fraunhofer HHI, considerable attention was also devoted to enabling interaction between DU/CU and the OSC-RIC over the E2 interface. To this end, modifications were made to the OAI O-DU to support message exchanges with the near-RT RIC. The implemented changes allow the O-DU to both send and receive messages from both KPM and RC service models. This functionality was successfully tested within our lab setup with the use of KPI monitoring and traffic prediction xApps, developed within WP5.1. The next development steps in this direction involve integration of the OSC-RIC with the updated OAI repository and implementation of a use case that utilizes RC messages.

3.5.5 SMO and Non-RT RIC

To provide support for Service Management and Orchestration (SMO) and non-real-time management of the RAN, Fraunhofer HHI has worked on the O1 interface. In collaboration with Highstreet Technologies and the OAI Software Alliance, Fraunhofer HHI has contributed to the integration of the O1 interface into the OAI software framework. This implementation utilizes the O1 adapter to connect an ONAP-based SMO system to the OAI CU/DU, with the following features being supported:

- Performance metric data reporting
- VES-based alarms and notifications
- RAN reconfiguration via the ONAP SDNC

The O1 adapter has been successfully tested within the lab test setup at Fraunhofer HHI.

3.5.6 E2E O-RAN evaluation system

To be able to test and evaluate the implemented interfaces and features described in the previous sections, we have set up a simple end-to-end O-RAN test system in our lab at Fraunhofer HHI. This setup provides a preliminary integration of the OAI-based RAN with disaggregated CU and DU, OAI core network, and FlexRIC acting as the near-RT RIC. An extended setup for testing the features implemented to support the near-RT RIC controlled handover described in Section 2.1.4 is illustrated in Figure 24. The setup includes two OAI DUs connected to a single OAI CU, with each DU being connected to a dedicated RU (USRP B210). This enables a handover over the F1 interface between the CU and the DUs, from one DU-RU path to the other.

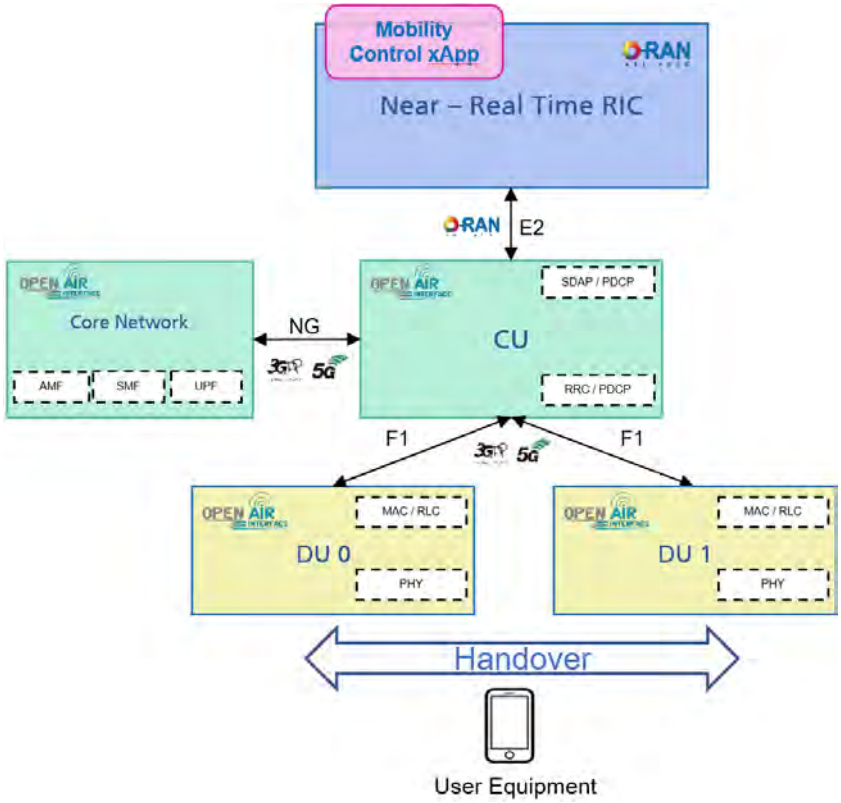

Figure 24: Overview of the test system for the Handover xApp

For evaluation purposes, a simple xApp is developed to run on top of the FlexRIC platform, which combines monitoring and control functionalities over the E2 interface. The xApp uses the monitoring functionalities of the KPM service model and combines it with monitoring using the newly added RRC message copy mechanism via the RC service model. Specifically, the xApp subscribes to the RRC measurement report to continuously monitor the RSRP value of the serving cell as measured by the UE. If the measured RSRP drops below a pre-defined threshold, the xApp pre-emptively triggers a handover command through the RIC Control Service. In addition, the xApp is capable of detecting a continuous decrease in the RSRP value, in which case it also triggers the handover pre-emptively before the connection quality becomes undesirably low.

For visualization of the results at run-time and to enable quick insight into the performance, we have also developed a Grafana-based monitoring interface which can be adjusted to the case being investigated. An example layout used with the xApp for testing of the handover functionality described above is shown in Figure 25.

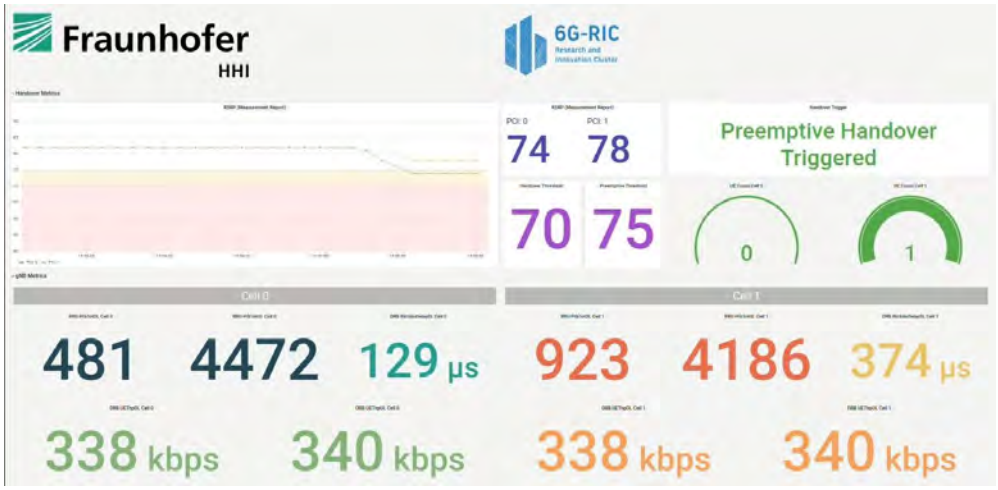

Figure 25: Grafana-based monitoring interfaces (example configuration for the Handover xApp)

In addition to the fully OAI-based setup from Figure 24, we have also prepared another end-to-end testing setup which incorporates NVIDIA ARC and OAI-based components together with the O-SC RIC (E release) for near real-time RAN monitoring, control, and optimization via E2 interface. This configuration includes the OAI core, CUs, and DU-high (MAC/RLC), as illustrated in Figure 26. In this setup, signal processing and data plane functions are handled by the OAI DU and CU, while physical (PHY) layer processing is managed by a system of eight ARC nodes operating on Gigabyte E251-U70 NVIDIA servers. The system utilizes a 5G Functional API (FAPI) interface for communication between the DU-High and DU-Low, as specified by the Small Cell Forum (SCF) [24]. Additionally, the setup integrates a Foxconn RPQN 4T4R RU with Ethernet and optical fronthaul connectivity.

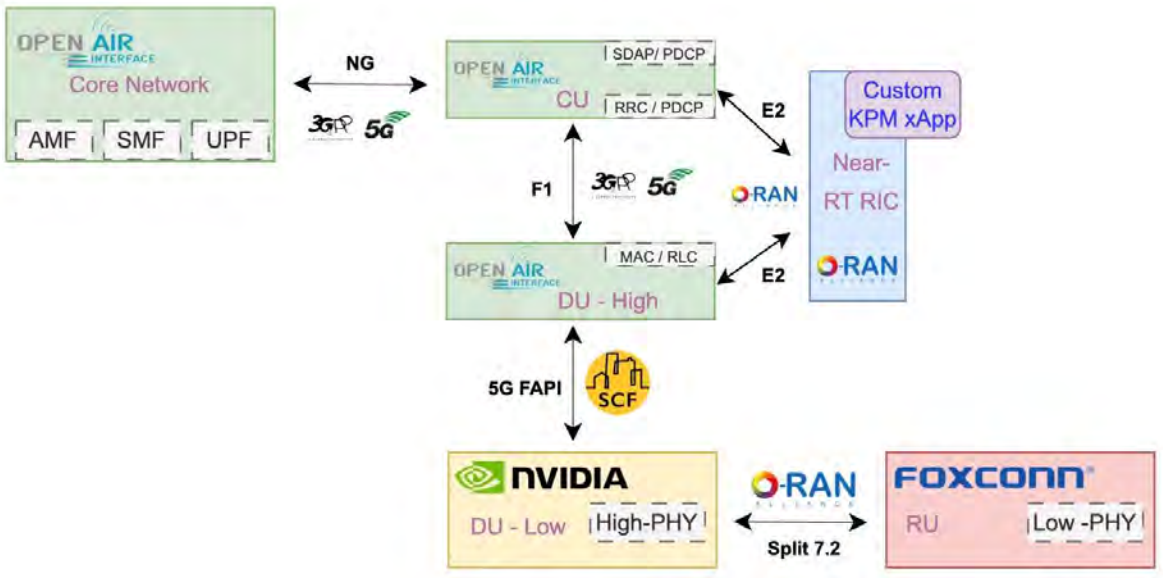

Figure 26: End-to-end architecture for NVIDIA-OAI based E2E setup

3.6 Kalray

The purpose for Kalray for the 5G-OPERA project was to develop a hardware acceleration solution for the Distributed Unit function, based on its MPPA processor architecture. The third generation of MPPA processor called Coolidge v2 was released during the project and was used for the implementation of the solution. Coolidge v2 corresponds to the generation that answers the requirements to support Open RAN solutions.

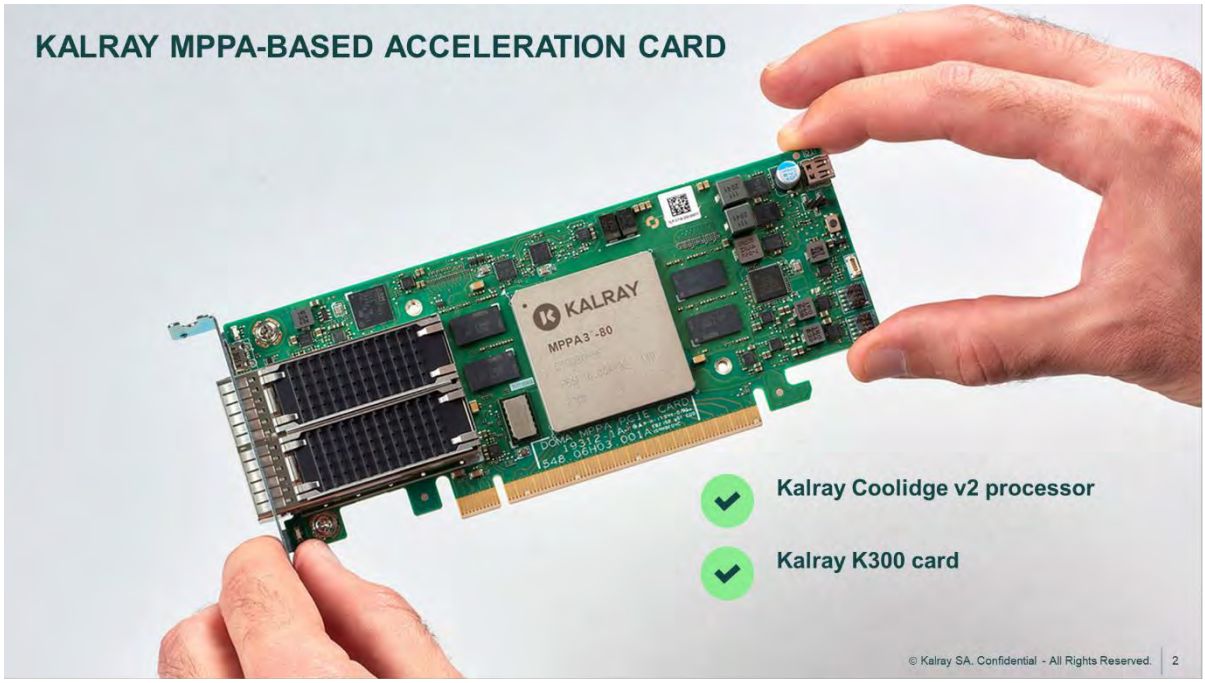

The main characteristics of the Coolidge V2 processor relevant for acceleration of Open RAN functions include:

- 80 VLIW 64-bits cores at 1.1GHz that can be used to program Layer-1 High-PHY functions on the MPPA processor using standard languages such as C/C++, coupled with 80 coprocessors for matrix operations and Artificial Intelligence inferencing
- Embedded hardware accelerators supporting 5G FEC and cryptography. For 5G FEC, the hardware bloc includes de-interleaving, rate matching / dematching, LDPC encoding / decoding, Code Blocks segmentation / desegmentation, CRC attachment / check, as well as HARQ. For cryptography, the accelerator includes support for the specific protocols required for 5G such as Snow3G, ZUC, or Kasumi.
- High speed interfaces including 2x 100G Ethernet and PCIe Gen4 16 lanes
- Support for TSN and synchronization features including PTP and SyncE

The PCIe acceleration card that was designed and developed during the 5G-OPERA project is the Kalray K300 that is based on a single MPPA chipset per board, which main specifications can be found in the following diagram. The Kalray K300 was used to realize the different 5G-OPERA tests as well as the integration with the Firecell solution.

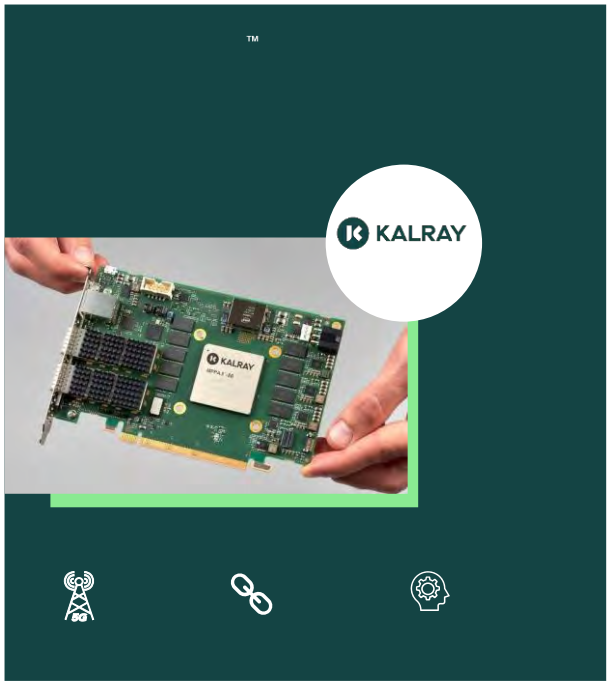

Processors	1x MPPA® Coolidge v2 @ 1.1GHz
Number of cores	80 independent cores
Memory	16GB DDR4 SDRAM @ 3200MT/s 40MB on chip memory
I/O	PCIe Gen4 16-lanes
Ethernet	4x 10/25G SFP28 interfaces
Synchronization	IEEE 1588 PTP andSyncE support
LDPC Performances	5G FEC: Decode: 12Gbps, Encode: 52Gbps
Cryptography	Cryptographic Accelerator designed to offload the server host processor to improve the speed of IPsec ESP, IPsec AH, SRTP, SSL, TLS, DTLS and MACsec, support for wireless algorithms (Snow3G, Kasumi, ZUC)
Crypto Performances	Up to 116Gbps encryption and 128Gbps decryption
AI Performances	FP16: 25 TFLOPS, INT8: 50 TOPS
Form Factor	Full Height Full Length
TDP	20W – 50W typical depending on use case
Thermal	Heatsink, passive or optional fan

The acceleration mode that was developed by Kalray for the 5G-OPERA project was FEC Lookaside. This means that the 5G FEC function is offloaded from the DU server CPU onto the Kalray acceleration card and Kalray Coolidge v2 processor. This mechanism was implemented by Kalray using the DPDK BBdev standard. DPDK BBDev (Baseband Device) is a framework within the Data Plane Development Kit (DPDK) project designed to accelerate the processing of baseband operations in wireless communication systems, including 5G NR. BBDev provides a standardized API for hardware and software baseband processing, enabling efficient integration for a variety of hardware accelerators such as FPGAs, ASICs, GPUs, etc. The Open Air Interface application supports 5G Forward Error Correction (FEC) offloading via DPDK BBdev, which facilitated the integration of the Kalray K300 with the OAI software using DPDK BBDev.

The software development performed by Kalray mainly consisted in the creation of a DPDK driver and DPDK library enabling the OAI software to offload 5G FEC onto the Kalray K300, as well as a Kalray SDK release AccessCore RAN containing the DPDK parts, but also the complete toolchain allowing to make use of the Kalray card. When instructed to offload FEC, the OAI software loads the DPDK library that was indicated during compilation, and that includes the offloading function. The DPDK driver called bbdev_mppa3_ep enables the usage of the FEC hardware accelerator from the CPU host via the DPDK interface. The Kalray DPDK implementation was first tested with a standalone OAI application, to ensure that the DPDK interface and the FEC hardware accelerator were operational.

This setup was in a second time used to measure the maximum throughput that can be achieved with the Kalray K300 as acceleration solution. Performances of up to 52Gbits/s in Encode and 12Gbits/s in Decode could be demonstrated by Kalray during the 5G-OPERA project. It shall be noted here that such figures cannot be achieved with the intended setup for the project, given that they exceed by far the amount of FEC traffic that could be generated by a single 4T4R radio with 100MHz bandwidth. The figures are however still important to compare the Kalray K300 acceleration card with alternative market products, and to confirm its suitability for radio deployment also aiming at mobile operator macro networks.

The configuration that was used by Kalray to perform end-to-end testing with the Kalray K300 acceleration card was based on the Firecell O-RAN Labkit solution. The latter consists in a pre-

integrated system that includes a VVDN radio unit (100Mhz) supporting the ORAN Split 7.2, an x86 server hosting the Firecell applications based on OAI and covering the gNB and Core functions, as well as a user terminal. The radio unit is directly connected to a port of the server NIC card (Intel E810), and the synchronization signals is derived from a GPS source. Modifications made by Kalray to the standard setup included the integration of the Kalray K300 card in the Firecell server, and recompilation of the OAI gNB software so that it could make use of the FEC offloading capabilities offered by the Kalray K300 acceleration card.

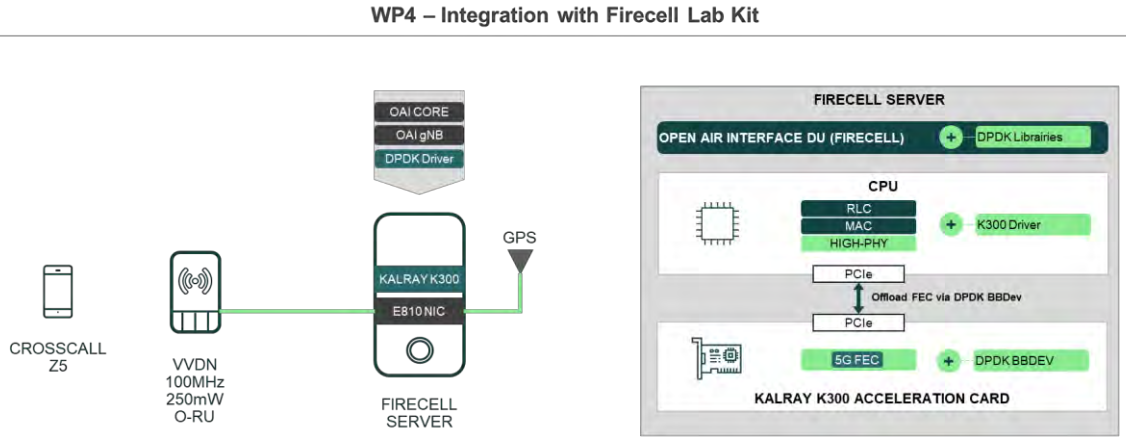

The list of software components used with Firecell O-RAN Labkit solution is provided in the following table:

Software component	Version
Kalray AccessCore RAN (ACR)	0.3
DPDK	24.07
Linux	Ubuntu 22.04
Firecell Software	3.2


KALRAY
THE POWER OF MORE

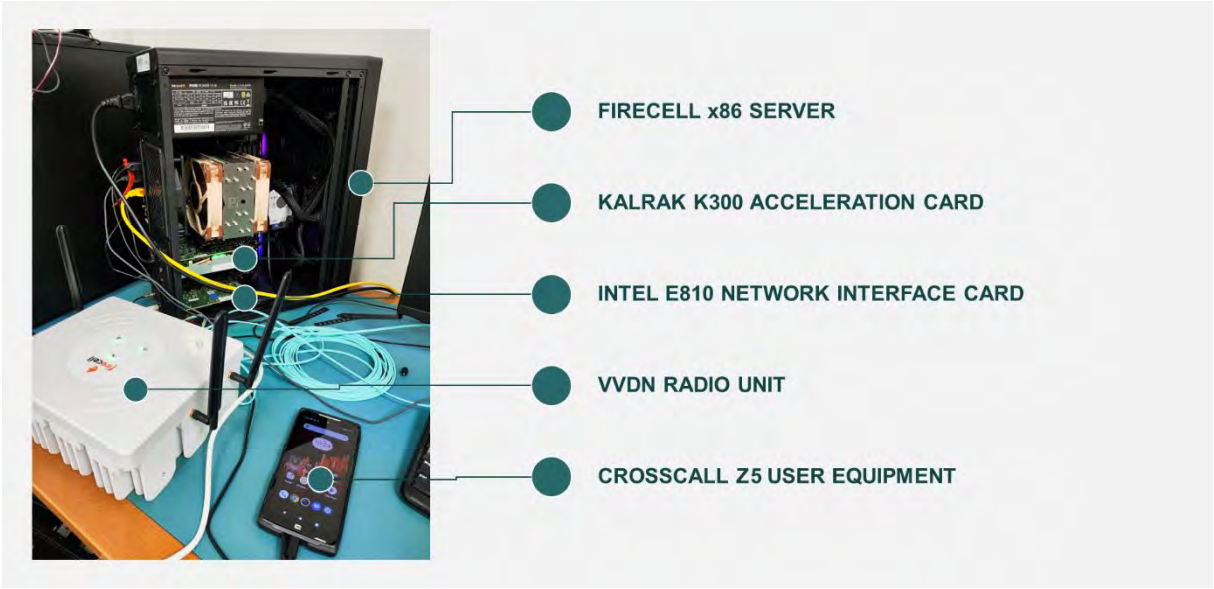
FIRECELL x86 SERVER
KALRAK K300 ACCELERATION CARD
INTEL E810 NETWORK INTERFACE CARD
VVDN RADIO UNIT
CROSSCALL Z5 USER EQUIPMENT

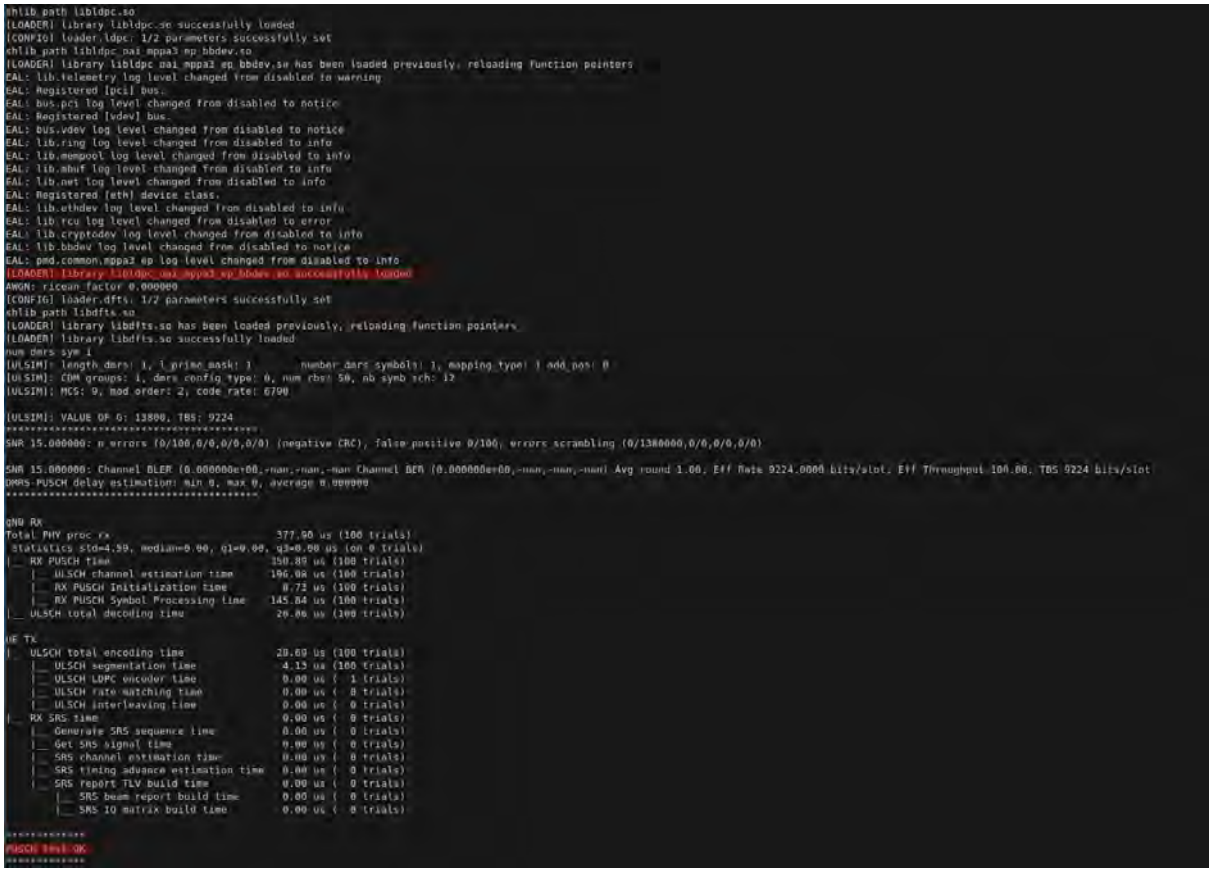
[LOADER] library libldpc.so successfully loaded
[CONFIG] loader.ldpc: 1/2 parameters successfully set
[LOADER] library libldpc_oai_mppa3_ep_bbdev.so successfully loaded
[LOADER] library libdfts.so has been loaded previously, reloading function pointers
[LOADER] library libdfts.so successfully loaded
[ULSIM]: length_dmrs: 1, l_prime_mask: 1 number_dmrs_symbols: 1, mapping_type: 1 add_pos: 0
[ULSIM]: CDM groups: 1, dmrs config type: 0, num rbs: 50, nb symb sch: 12
[ULSIM]: MCS: 9, mod_order: 2, code rate: 6790
[ULSIM]: VALUE OF G: 13800, TBS: 9224
SNR 15.000000: n_errors (0/100,0/0,0/0,0/0) (negative CRC), false positive 0/100, errors scrambling (0/13800000,0/0,0/0,0/0)
SNR 15.000000: Channel BLER (0.000000e+00,-nan,-nan,-nan Channel BER (0.000000e+00,-nan,-nan,-nan) Avg round 1.00, Eff Rate 9224.0000 bits/slot, Eff Throughput 100.00, TBS 9224 bits/slot
DMRS-PUSCH delay estimation: min 0, max 0, average 0.000000
gNB RX
Total PHY proc rx 377.90 us (100 trials)
Statistics std=4.59, median=0.00, q1=0.00, q3=0.00 us (on 0 trials)
RX PUSCH time 150.89 us (100 trials)
ULSCH channel estimation time 196.08 us (100 trials)
RX PUSCH Initialization time 0.72 us (100 trials)
RX PUSCH Symbol Processing time 145.84 us (100 trials)
ULSCH total decoding time 20.06 us (100 trials)
UE TX
ULSCH total encoding time 28.69 us (100 trials)
ULSCH segmentation time 4.15 us (100 trials)
ULSCH LDPC encoder time 0.00 us (1 trials)
ULSCH rate matching time 0.00 us (0 trials)
ULSCH interleaving time 0.00 us (0 trials)
RX SRS time 0.00 us (0 trials)
Generate SRS sequence time 0.00 us (0 trials)
Get SRS signal time 0.00 us (0 trials)
SRS channel estimation time 0.00 us (0 trials)
SRS timing advance estimation time 0.00 us (0 trials)
SRS report TLV build time 0.00 us (0 trials)
SRS beam report build time 0.00 us (0 trials)
SRS IQ matrix build time 0.00 us (0 trials)
PUSCH test OK

3.7 NXP Germany

3.7.1 O-RAN 5G SA Network

NXP Germany reviewed the capabilities of its Software Defined Radio (SDR) platforms and customized selected HW and SW components to implement an end-to-end (E2E) 5G Standalone (SA) private network based on O-RAN architecture. That system is intended to provide radio access in 5G campus and other non-public networks. Proprietary CU, DU as well as RU SW stacks of the system were enabled by NXP's Layerscape ® Access processors, following open interfaces principles to allow for integration with project-partner's solutions like the DU- acceleration by Xelera technologies, even network components running opensource solutions.

The main features of the system are listed below. Details of the CU, DU and RU 5G SA implemented by NXP are presented in following subsections.

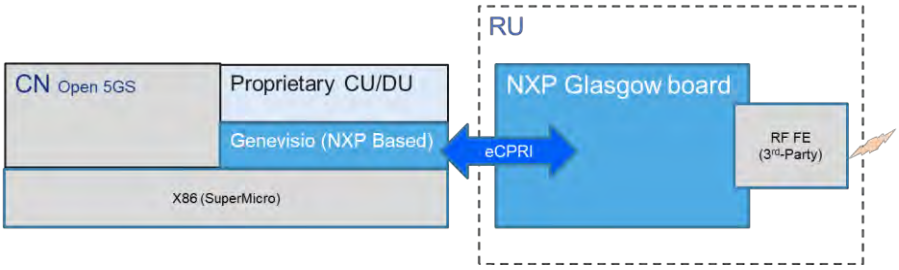

Figure 27: General concept NXP Based 5G SA Network

System Features

- O-RAN Split 7-2aF
- 5G Band: n78
- Achieved throughput: 0.94 -1.07 Gbps DL, avg ~80 Mbps UL
- Achieved Estimated latency: Avg ~10 ms (UE to 5GC), Avg ~20 ms (UE to UE)
- BW: 100 Mbps
- IEEE 1588 time syncing
- Fronthaul Interface: eCPRI

3.7.2 O-RAN Centralized Unit (O-CU), Distributed Unit (O-DU) and 5G Core (5GC)

O-RAN CU (O-CU) encompasses Transport/S1, PDCP and RRC/Control plane processing. The distributed unit (O-DU) is responsible for MAC/RLC and High-PHY processing. Use of optimized (ARM NEON) SIMD kernels allow for low power consumption, while still supporting look-aside acceleration for forward error correction, as well as DSP processing acceleration. Layerscape ® Access chips include both VSPA DSP cores and hardware accelerators for LDCP/Polar, among other functions. Extended details on NXP DU acceleration can be reviewed in section 5.1.3.3 of Deliverable 3.1

Figure 28: NXPLX2160 and LA12xx processors enable the CU/DU acceleration Board (By Genevisio).

Some technical characteristics of the CU/DU board are: can provide multiple 25GbE eCPRI interfaces, could work seamlessly with a CU running in the MEC servers for the lower layer split (LLS), then ensuring ultra-low latency, 1pps & GPS timing synchronization, PCIe x8 form factor. PCIe can be used as communication interface, however in the presented information, that bus is used mainly for power. The CU/DU board is hosted by the CN server (x86 architecture) and fronthaul is implemented on 10 Gbit Internet (external), as in Figure 29.

Figure 29: CU/DU Board (NXP Layerscape® Access processors) And 5GC Server compact implementation. Backhaul deployed over 10Gbit Ethernet.

Opensource Compatible Acceleration Enablement

NXP provided the HW platforms based to enable the project partner Xelera to implement DU acceleration. Close tech support and SW enhancements were deployed, resulting in suitable updated versions of the involved NXP BSP. The achieved implementation show how opensource 5G stacks can be integrated with ARM-based architectures to implemented state-of-the-art 5G acceleration. Further details are presented in the respective sections by Xelera Technologies.

3.7.3 NXP-based O-RAN Radio Unit (O-RU)

NXP Glasgow board (featuring LX2160 and LA12xx) was adapted and integrated with 3rd-Party RF front ends to deploy the proof of concept for an FR1 2T2R Radio Unit. Proprietary L1 software enables the RU functionality, and open interfaces allow for the integration of this O-RU with O-RAN compatible DUs.

Figure 30: NXP Glasgow board customized to deploy a 5G O-RAN RU.

The RU complies with 3GPP split 7-2a. LX2160 16-core processor supports the eCPRI interface of this RU through 10 Gbit/25 Gbit Ethernet ports. Regarding power, it performs 10 dBm per Tx Channel. Section 2.1.5 in Deliverable 3.2 extends on the architecture and other characteristics of the NXP-based RU.

The table summarizes main configurations and SW features of the implemented O-RU.

Feature	Implementation
Specification Version	ORAN.WG4.CUS.0-v07.01
TDD Configuration	DDDSU DDDSU (10:2:2)
Nominal sub-carrier spacing	30k
PRACH preamble format	B4
Category	A
Network delay determination	1588
Jumbo Frames	Yes

Table 1: NXP O-RU configuration (summary)

3.7.4 Network Side TSN Contributions

NXP participated in the research activities and proposal of different approaches to TSN implementation over O-RAN 5G networks. With the objective of enabling faster management of TSN user plane operations, NXP proposed the offload of UPF function from the Core Network to a dedicated networking hardware LS1028 TSN switch. In further steps, instances of the Network-Side TSN-Translator (NW-TT) could run in such switch.

To allow for UPF offloading, an instance of OAI UPF was built in the TSN switch. NW-TT is not the scope of NXP in 5G-OPERA. However, NXP provides guidance and support to project partners to use the adapted platform for implementing the NW-TT.

Figure 31: NXP LS1028A-RDB to implement UPF off-load.

In the followed approach, a custom version of NXP's Rea-Time Edge operating system is built using Yocto SW project to include the right OAI libarries in the operating system. The platform A full guideline for building and installing OAI UPF on the NXP LS1028 TSN switch is available for the project partners.

3.8 TU Berlin

3.8.1 Introduction

The contribution of TU Berlin primarily focused on the development of the radio head for the Radio Unit (RU). The developed radio head comprises an antenna module operating in the NR258 band (Frequency Region 2). The antenna features a cylindrical topology with a 4 x 32 element configuration (4 vertical and 32 in azimuth) and beamforming circuitry based on the microwave chip MMW9014K, provided by the project partner NXP. The RU includes the radio head (antenna and analog frontend) and the Low-PHY baseband processing unit, designed according to the 7.2 split architecture. The developed radio head interfaces with the Low-PHY, for which TU Berlin is also responsible for developing the corresponding interface.

3.8.2 Radio Head

Antenna

The radio head comprises an antenna array and an analog frontend (beamforming circuitry). The antenna element is part of a 4x32 circular array, with fixed elevation and 360° azimuth coverage. The radiating elements are patch antennas operating at 26 GHz, fed by ground apertures. The fractional bandwidth (FBW) of each individual radiator exceeds 12%.

The cylindrical array consists of 32 radiating segments, each containing 4 dual-polarized radiators (forming a 4x1 linear array) with an FBW of 10%. Power is fed to each patch radiator via one-to-four power dividers and through dog bone-shaped ground apertures. The segments are assembled in a cylindrical arrangement, with the spacing between adjacent segments not exceeding 0.8λ in free space.

Each segment is connected to the analog frontend via SMPM cables for each polarization, linking the antenna segments with the analog frontend circuitry.

For initial testing and measurements, only half of the cylindrical array is being designed and manufactured. The complete array will be developed following the successful testing of the first prototype. The first prototype has been submitted for manufacturing at Contag AG and will be ready for integration early next month (Jan 2025).

Beamformer Board

The analog frontend is based on the Beamformer Integrated Circuit (BFIC) MMW9014K, provided by NXP. This chip is a 4-channel, dual-polarized analog beamforming IC that offers a complete analog signal processing path, including amplification in both directions (receiving and transmitting) and phase shifting. Amplification and phase shifting are controlled via two distinct serial digital interfaces: a high-speed differential bus (LVDS) or a CMOS Serial Peripheral Interface (SPI).

The chip operates within the 5G FR2 frequency range (24.25 GHz–27.5 GHz) and features two RF input power ports (RFin) and eight RF output power ports (RFout), with four outputs dedicated to each polarization. For a fixed beam in elevation—requiring constant phase in rows and varying phases in columns—a full cylindrical array requires 8 BFICs, while a half-cylinder array requires 4 BFICs. The BFICs are connected in a daisy-chain configuration and are controlled using both differential and single-ended interfaces.

Since the chips operate in the 5G mmWave domain, the output signals from the digital-to-analog converters in the baseband (BB) domain must be upconverted from 6 GHz to 26 GHz. To achieve this, two 4-channel 26 GHz upconverters (UPCs) from Auden are utilized to supply input signals to the BFICs. The RFout ports are connected to the feed networks (power dividers) in the segments via SMPM SMD connectors and 12-inch SMPM cables.

The beamformer board has been manufactured at Contag AG and is ready for testing in the coming weeks.

3.8.3 Low-PHY Interface

Originally, the Low-PHY interface was to be provided by NXP Germany as part of their contribution to WP4. However, it was decided to develop a Low-PHY emulator at TU Berlin, as NXP Germany could only provide hardware for FR1, without the possibility of frequency extension.

The Low-PHY control plane has been successfully implemented and emulated on an FPGA. It will be linked to the radio head as soon as testing and characterization are completed successfully.

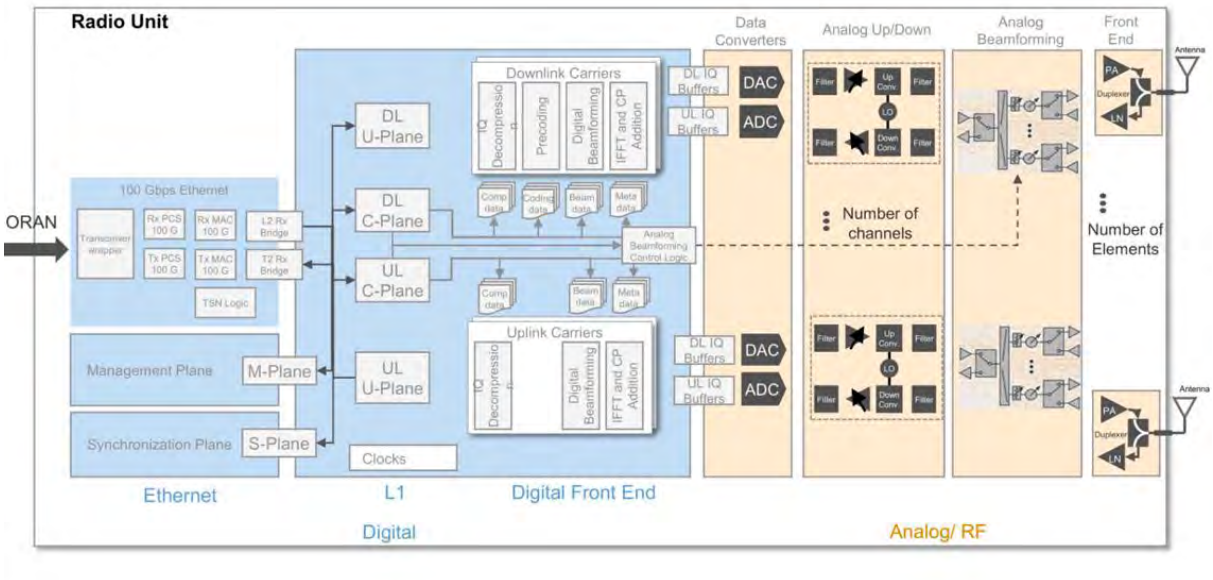

Figure 32: Low-PHY interface based on the 7.2 split architecture as it has been implemented by the TU Berlin

3.9 Xelera Technologies

3.9.1 Introduction

As part of the 5G-OPERA project, and in partnership with NXP Semiconductors Germany, Xelera Technologies has developed an Open-RAN compliant Distributed Unit (DU) that makes use of hardware acceleration. By offloading compute intensive workloads onto the inline accelerator, the DU is more cost effective, due to lower total hardware costs and lower energy consumption, and has a higher performance enabling the use of multiple Radio Units (RUs) and hundreds of devices connected to a single DU.

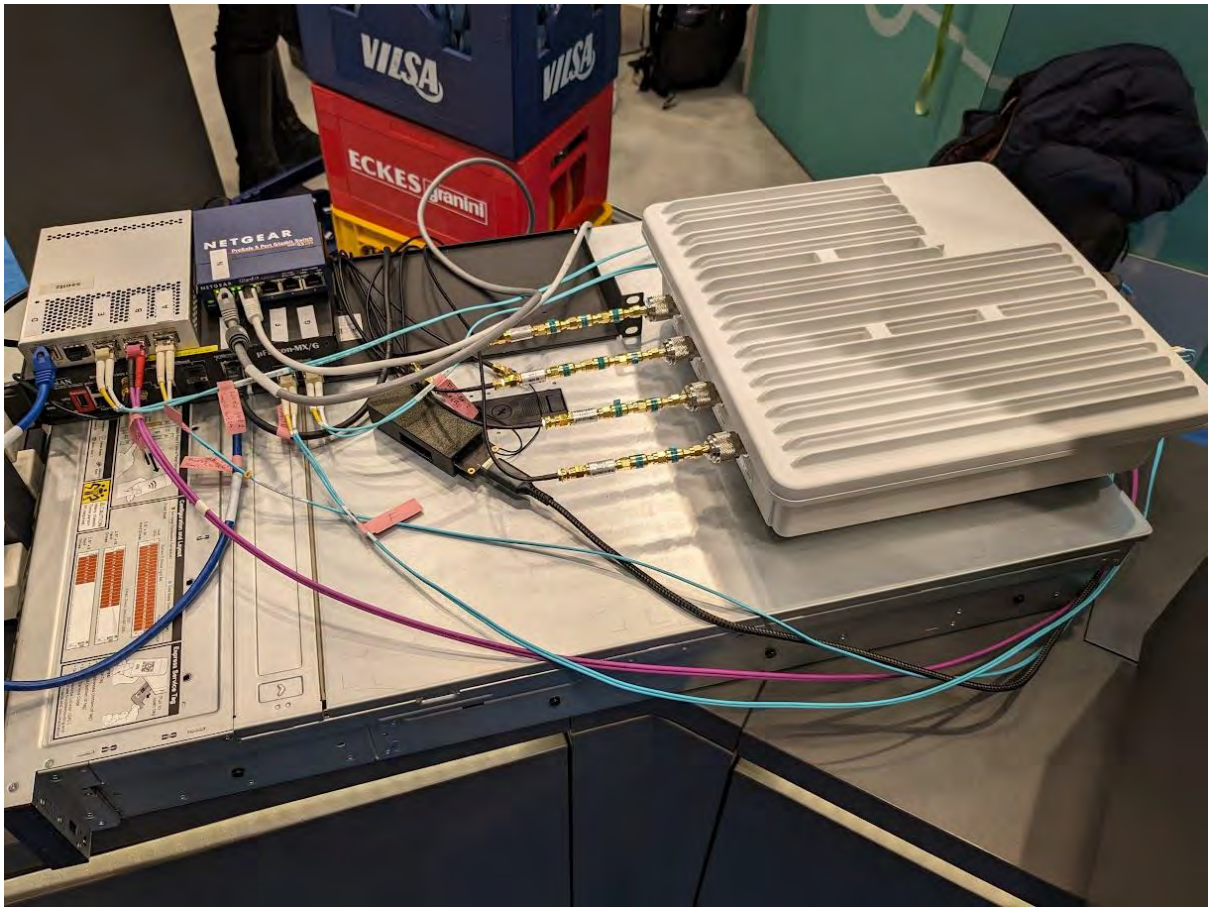

Figure 33: Example setup at Hannover Messe 2024 including COTS server with GeneVisio NPE015 hardware accelerator card, PTP switch and Benetel RAN650 RU

3.9.2 Objectives

Performance

The NXP LA1201 Programmable Baseband Processors on the hardware accelerated DU significantly increase the system performance and enables the use of multiple RUs and hundreds of devices.

Cost efficiency

The usage of hardware acceleration reduces the required hardware resources for the host system. Instead of a powerful appliance only a small PCIe card is required. This lowers the total system costs.

Energy efficiency

By offloading most functionality onto the energy efficient NXP LX2160 Multicore Processors on the hardware accelerated DU the power consumption of the system can be reduced by up to 50% compared to running the same software on a state-of-the-art server.

Flexibility

Implementing the DU to be fully Open-RAN compliant grants a high flexibility with Split Option 7-2 at Fronthaul and Option 6 (FAPI) or Option 2 (F1) at Midhaul.

3.9.3 Features

- Off-the-shelf PCIe device that can easily be integrated into existing IT infrastructure
- Hardware accelerator based on NXP Layerscape® technology
- Fully Open-RAN compliant
- Supports Fronthaul Split Option 7-2
- Supports Midhaul Split Options 6 (FAPI) and 2 (F1)
- Supports up to 4 Radio Units via 4 x 10/25GbE SFP28 Fronthaul ports
- Supports 100+ user devices
- Supports throughput of 75+ Mbit/s uplink and 1+ Gbit/s downlink

3.9.4 Components

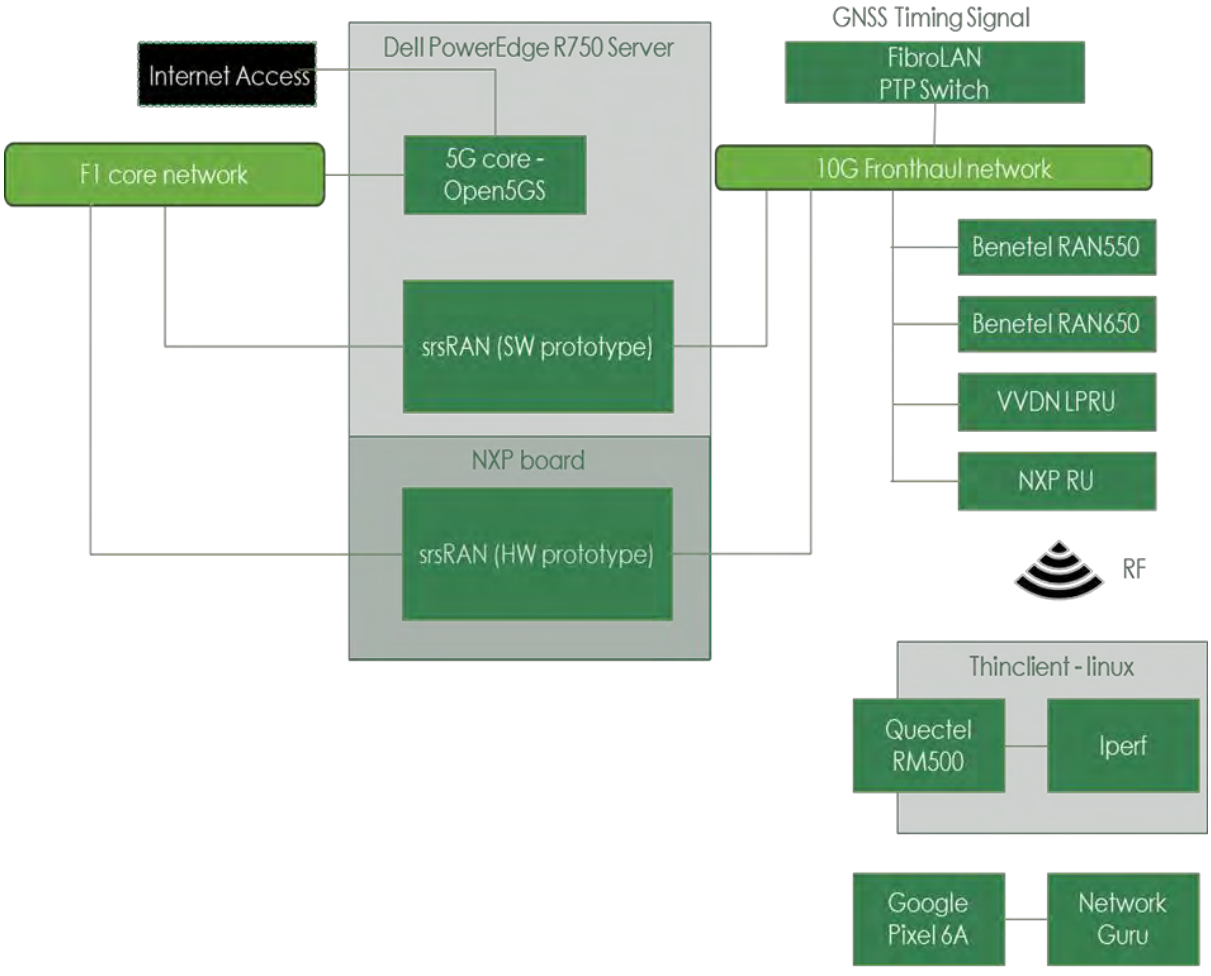

Figure 34: End-to-end test setup for software and hardware prototypes

srsRAN-bases software prototype

The software prototype is running without the hardware acceleration component on a COTS server from DELL and is used as reference system for the development of the accelerated flavour. It mainly consists of various open-source projects that require integration and configuration, and some additions by Xelera to be able to be deployed on the target system.

Integrated open-source components

- srsRAN_Project 24.04 [25]
- DPDK 23.11 [26]
- Open5GS 2.7.2 [27]
- LinuxPTP 4.2 [28]

Software specifications

- 3GPP release 17 aligned
- FDD/TDD supported, all FR1 bands
- All bandwidths (e.g. 100 MHz TDD, 50 MHz FDD)
- 15/30 kHz subcarrier spacing
- All physical channels including PUCCH Format 1 and 2, excluding Sounding-RS
- Highly optimized LDPC and Polar encoder/decoder for ARM Neon and x86 AVX2/AVX512
- All RRC procedures
- All MAC procedures
- Split 7.2 support using in-house OFH library
- Support for QAM-256
- 4x4 MIMO
- Slicing
- NTN GEO support
- CU-DU Split

GeneVisio NPE015-based hardware prototype

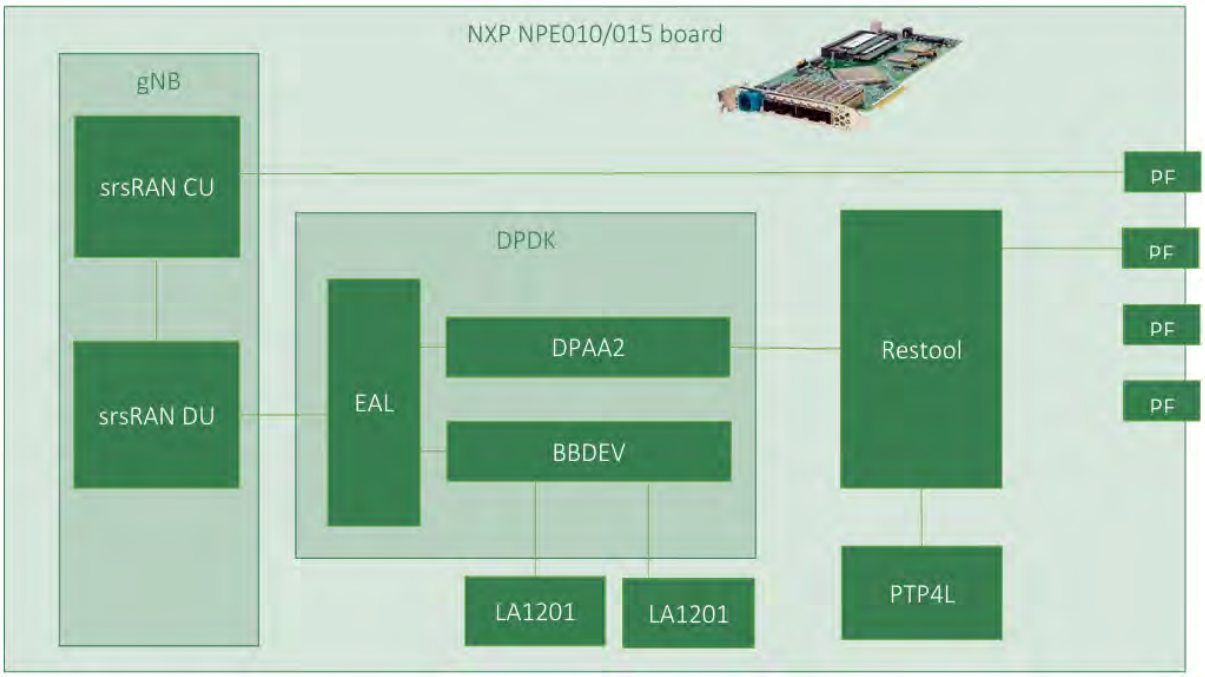

Figure 35: GeneVisio NPE015 hardware accelerator PCIe card

The hardware prototype is a progression of the srsRAN-based software prototype that is running on the hardware acceleration board utilizing its special accelerator chips to improve performance at low

costs. It consists of various open-source projects and firmware supplied by the hardware vendor and acceleration-specific additions by Xelera.

Hardware specifications

- Processor: Chipset NXP LX2160A, Max. Speed 2.2GHz, 16 Cores
- System Memory: SO-DIMM Memory 2 slots, up to 32GB per slot
- Baseband Processor: Chipset 2 LA1201
- GNSS: Module u-blox ZED-F9T
- I/O Interface: Ethernet 4 SFP28, 10/25GbE support, PCI Express Gen 3 x8, Timing and Synchronization USB-C, 1PPS in/out; 10Mhz in/out; ToD
- Power: Max. Consumption 75W
- Size: Dimension 111.15 x 241 x 17.14 mm (Full Height, 3/4 Length)
- Environmental: Passives heatsink cooling, operating temperature 0 ~ 55℃
- Certification: FCC, CE, VCCI, UL

3.9.5 Benchmarks

The following benchmarks have been conducted with the first prototype. Since optimizations are still ongoing the current numbers do not represent the target values but only the contemporary state.

	UE Device	DL Throughput	UL Throughput
2x2 20MHz	Quectel RM500Q-GL OTA	101M	12M
	Quectel RM500Q-GL Cabled	N/A	N/A
	Google Pixel 6a OTA	98M	7M
2x2 100MHz	Quectel RM500Q-GL OTA	607M	79M
	Quectel RM500Q-GL Cabled	607M	79M
	Google Pixel 6a OTA	607M	79M
4x2 100MHz	Quectel RM500Q-GL OTA	780M	75M
	Quectel RM500Q-GL Cabled	1.1G	79M
	Google Pixel 6a OTA	544M	48M

Table 2: End-to-end benchmark using Xelera DYRAN, Benetel RAN650 and different UEs

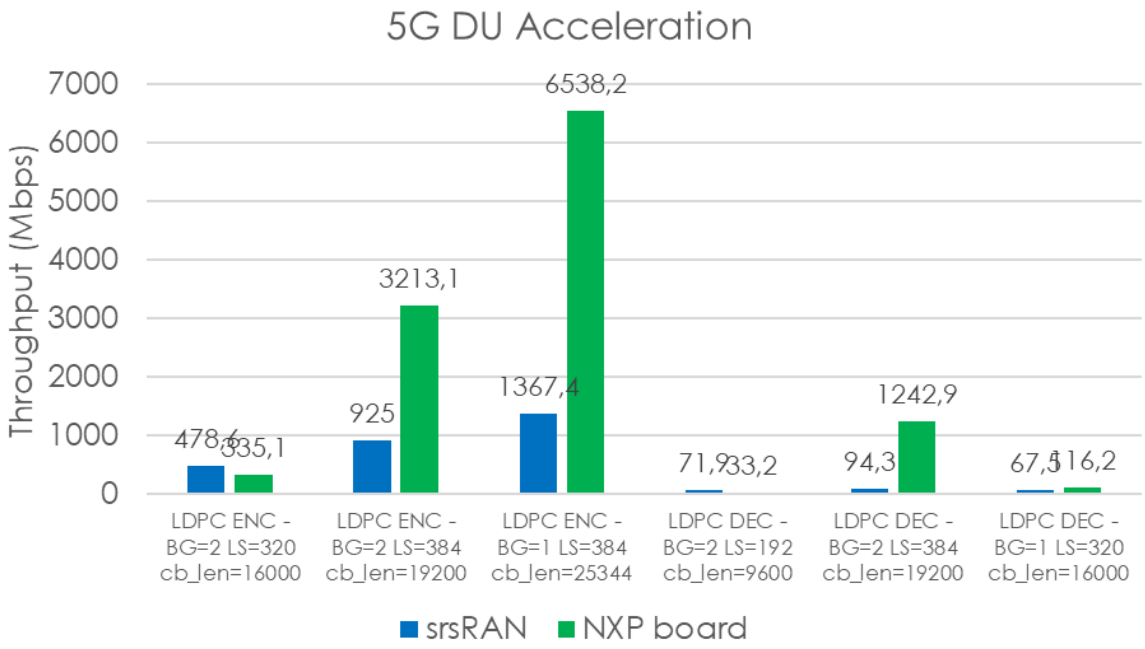

Figure 36: LDPC-only benchmark comparing srsRAN-LDPC on high-end COTS server vs. LDPC on NXP baseband processor using sample applications; BG: Base graph, LS: Lifting size, Cb_len: Code block length

3.9.6 Release

The DU development is proprietary and not officially released yet. For early access, please contact Xelera Technologies [29] directly.

4 References

[1] 5G-OPERA WP4, "5G-OPERA Deliverable 4.1: Architecture and Interfaces specification," 2022.

[2] K. C. N. Z. R. &. B. Y. Tahkoubit, "Impact of Channel Estimation Feedback on 5G NR Integrated Access and Backhaul," *2024 IEEE EuCNC/6G Summit,* 2024.

[3] K. C. N. D. D. &. D. J. B. Tahkoubit, "Integrated access and backhaul with RIS: analyzing inter-source and inter-user interference," *2023 IEEE PIMRC,* 2023.

[4] T. M. I. G. M. P. R. K. e. a. F. Kaltenberger, "Driving Innovation in 6G Wireless Technologies: The OpenAirInterface Approach," *Computer Networks,* 2024.

[5] OpenAirInterface Software Alliance, "End-to-end over-the-air 5G demo with OAI CU-DU and LITEON O-RU," in *O-RAN F2F Meeting*, Osaka, Japan, 2024.

[6] I. K. F. K. N. H. R. S. d. S. A. K. C. D. S. B. J. M. J. T. M. M. P. D. K. D. Villa, "An Open, Programmable, Multi-vendor 5G O-RAN Testbed with NVIDIA ARC and OpenAirInterface," *Proc. of the 2nd Workshop on Next-generation Open and Programmable Radio Access Networks,* 2024.

[7] OpenAirInterface Software Alliance, "NVIDIA ARC with OAI L2+ for next generation innovation," in *O-RAN F2F Meeting*, Osaka, Japan, 2023.

[8] OpenAirInterface Software Alliance, "End-to-end 5G demo featuring looka- side acceleration of LDPC encode/decode and O-RAN F1 and 7.2 splits," in *O-RAN F2F Meeting*, Phoenix, AZ, USA, 2023.

[9] OpenAirInterface Software Alliance, "Showcasing fronthaul 7.2 and looka- side hardware acceleration of 5G," in *O-RAN F2F Meeting*, Athens, Greece, 2024.

[10] M. I. N. N. R. Schmidt, "Improving QoE using OSC near-RT RIC and OAI 5G RAN leveraging O-RAN E2 KPM and RC SMs," in *O-RAN F2F meeting*, Phoenix (AZ), United States, 2023.

[11] OpenAirInterface Software Alliance, "Improving QoE using OSC near-RT RIC and OAI 5G RAN leveraging O-RAN E2SM-KPM and E2SM-RC," in *O-RAN F2F Meeting*, Phoenix, AZ, USA, 2023.

[12] M. I. N. N. R. Schmidt, "FlexRIC: an SDK for next-generation SD-RANs," *Proceedings of the 17th International Conference on Emerging Networking EXperiments and Technologies,* p. 411–425, 2021.

[13] OpenAirInterface Software Alliance, "Improving QoE using O-RAN com- pliant near-RT RIC, KPM v03.00 and RC v01.03 SMs in an online multi-player game," in *O-RAN F2F Meeting*, Osaka, Japan, 2023.

[14] M. Kasparik, "Non-real-time control of openairinterface using the O-RAN O1-interface," in *Joint OSC/OSFG-OAI Workshop: End-to-End Reference Designs for O-RAN*, Boston (MA), United States, 2023.

[15] OpenAirInterface Software Alliance, “5G RAN management through an SMO/non-RT RIC via the O1 interface,” in *O-RAN F2F Meeting, October 2023*, Phoenix, AZ, USA, 2023.

[16] Open5GS, “Open5GS website,” Open5GS, [Online]. Available: https://open5gs.org. [Accessed December 2024].

[17] free5GC, “free5GC website,” free5GC, [Online]. Available: https://free5gc.org/. [Accessed December 2024].

[18] OpenAirInterface Software Alliance, “OpenAirInterface git repository,” OpenAirInterface Software Alliance, [Online]. Available: https://gitlab.eurecom.fr/oai/openairinterface5g/ . [Accessed 2024].

[19] Benetel, “5G Split 7.2x Outdoor O-RU product-website,” Benetel, [Online]. Available: https://benetel.com/ran650. [Accessed December 2024].

[20] LiteON, “LITEON, FlexFi All-in-one small cell,” LiteON, [Online]. Available: https://liteon.vinova.sg/180-5G-Total-Solutions. [Accessed December 2024].

[21] NVIDIA, “NVIDIA Aerial CUDA-Accelerated RAN website,” NVIDIA, [Online]. Available: https://developer.nvidia.com/aerial-cuda-accelerated-ran. [Accessed December 2024].

[22] OpenAirInterface, “OAI-5G Core Network website,” OpenAirInterface, [Online]. Available: https://openairinterface.org/oai-5g-core-network-project/. [Accessed December 2024].

[23] The Fraunhofer Institute for Open Communication Systems FOKUS, “Fraunhofer FOKUS Open5GCore website,” The Fraunhofer Institute for Open Communication Systems FOKUS, [Online]. Available: https://www.open5gcore.org/. [Accessed December 2024].

[24] Eurecom, “FlexRIC project gitlab,” Eurecom, [Online]. Available: https://gitlab.eurecom.fr/mosaic5g/flexric. [Accessed December 2024].

[25] O-RAN Alliance, “O-RAN SC Non-RT RIC project website,” O-RAN Alliance, [Online]. Available: https://docs.o-ran-sc.org/en/latest/. [Accessed December 2024].

[26] O-RAN Alliance, “O-RAN Working Group 3 Near-Real-time RAN Intelligent Controller E2 Service Model (E2SM), UE Related,” O-RAN Alliance, 2024.

[27] Small Cell Forum, “Small Cell Forum website,” Small Cell Forum, [Online]. Available: https://www.smallcellforum.org/. [Accessed December 2024].

[28] srsRAN, “GitHub: srsRAN Project,” srsRAN, [Online]. Available: https://github.com/srsran/srsRAN_Project/tree/release_24_04. [Accessed December 2024].

[29] DPDK Project, “ DPDK Downloads,” DPDK Project, [Online]. Available: https://fast.dpdk.org/rel/dpdk-23.11.tar.xz. [Accessed December 2024].

[30] Open5GS, “GitHub: Open5GS,” Open5GS, [Online]. Available: https://github.com/open5gs/open5gs/tree/v2.7.2. [Accessed December 2024].

[31] The Linux PTP Project, “The Linux PTP Project,” The Linux PTP Project, [Online]. Available: https://linuxptp.sourceforge.net/. [Accessed December 2024].

[32] Xelera Technologies GmbH, “Xelera Website,” Xelera Technologies GmbH, [Online]. Available: https://www.xelera.io/. [Accessed December 2024].