Deliverable 5.1

Use Case Identification and Demonstration Plan

December 2023

Co-funded by the European Union

Contractual Date of Delivery:	**August 31, 2023**
Actual Date of Delivery:	**December 15, 2023**
Editor(s):	**Josep Xavier Salvat, Jose A. Ayala Romero (NEC)**
Author(s)/Contributor(s):	**Josep Xavier Salvat, Jose A. Ayala Romero (NEC)**
	German Castellanos, Revaz Berozashvili, Simon Pryor (ACC)
	Vladica Sark, Darko Cvetkovski, Jesús Gutiérrez (IHP)
	Guillermo Bielsa (TSA)
	Miguel Catalan-Cid, Esteban Municio (i2CAT)
	Joss Armstrong, Jimmy O'Meara, Joseph McNamara (LMI)
	Israel Koffman, Baruch Globen (RunEL)
	Richard Mackenzie (BT)
	Ory Eger (PW)
	Anna Umbert (UPC)
	Mir Ghoraishi (GIGASYS)
Work Package	**WP5**
Target Dissemination Level	**Public**

This work is supported by the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101097083, BeGREEN project. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or SNS-JU. Neither the European Union nor the granting authority can be held responsible for them.

Revision History

Revision	Date	Editor / Commentator	Description of Edits
0.1	12.05.2023	Josep Xavier Salvat (NEC)	ToC made ready
0.2	27.09.2023	Guillermo Bielsa (TSA)	Architecture draft
0.3	04.10.2023	Miguel Catalán -Cid (i2CAT)	IP PoC Draft
	12.10.2023	German Castellanos (ACC)	CU-HW PoC Draft
	13.10.2023	Josep Xavier Salvat (NEC)	Sensing PoC draft
0.4	16.10.2023	Vladica Sark (IHP)	Added text on sensing in the sub-6 GHz and mmWave domain
0.5	18.10.2023	Richard Mackenzie (BT)	Added information about Adastral Park PoC integration
0.6	18.10.2023	Israel Koffman (RunEL)	PA blanking PoC draft
	19.10.2023	Ory Eger (PW)	DU HW Accelerator PoC draft
0.7	23.10.2023	Josep Xavier Salvat (NEC)	Added KPI mapping section
0.8	25.10.2023	Josep Xavier Salvat (NEC)	Editing effort and first review
0.9	08.11.2023	Anna Umbert (UPC)	Final Review
0.91	08.11.2023	Esteban Municio (i2CAT)	Final Review
0.92	09.11.2023	Josep Xavier Salvat (NEC)	Final changes to the document
0.93	27.11.2023	Jesús Gutiérrez (IHP) Mir Ghoraishi (GIGASYS)	Added missing information, rewrite sections and improve quality
0.94	04.12.2023	Jesús Gutiérrez (IHP) Mir Ghoraishi (GIGASYS)	Final review and proof reading
1.00	15.12.2023	Simon Pryor (ACC)	Submission to the EC

Table of Contents

List of Figures

List of Tables

List of Acronyms

3GPP	3rd Generation Partnership Project
5GC	5G Core
5GNR	5G New Radio
A/D	Analogue to Digital
AI	Artificial Intelligence
AoA	Angle-of-Arrival
API	Application Programming Interface
B5G	Beyond 5G
BLER	Block Error rate
BPF	Berkeley Packet Filter
COTS	Commercial Off-The-Shelf
CP	Control Plane
CPE	Customer-premises Equipment
CU	Central Unit
CW	Cell Wrapper
DoW	Description of Work
D/A	Digital to Analogue
DU	Distributed Unit
E2E	End-to-End
EE	Energy Efficiency
FPGA	Field Programmable Gate Array
gNB	Next Generation Node B
GPU	Graphics Processing Unit
GTP	GPRS Tunnelling Protocol
HW	Hardware
IDFT	Inverse discrete Fourier transform
IP	Internet Protocol
ISAC	Integrated Sensing and Communication
ISM	Industrial Scientific and Medical
KPI	Key Performance Indicator
L2	Layer 2
LDPC	Low Density Parity Check codes
LNS	Linux Network Stack
LO	Local Oscillator
LoS	Line-of-Sight
LTE	Long Term Evolution
MAC	Medium Access Control
ML	Machine Learning
MMSE	Minimum Mean Square Error
mmWave	Millimetre Wave
MU-MIMO	Multiuser MIMO
NATS	Neural Autonomic Transport System
near-RT	near Real-Time
NFV	Network Function Virtualisation
NLoS	Non-Line-of-Sight
non-RT	Non-Real-Time

OFDMA	Orthogonal Frequency-Division Multiple Access
O-RAN	Open RAN
PDCP	Packet Data Convergence Protocol
PHY	Physical layer
PC	Personal Computer
PoC	Proof-of-Concept
RAN	Radio Access Network
RF	Radio Frequency
RFSoC	Radio Frequency System on Chip
RIC	Radio Intelligent Controller
RMS	Root Mean Square
RRM	Radio Resource Management
RU	Radio Unit
SDR	Software Defined Radio
SLA	Service Level Assurance
SMO	Service Management Orchestrator
SNR	Signal-to-Noise Ratio
SRS	Sounding Reference Signal
SU-MIMO	Single User MIMO
UE	User Equipment
UPF	User Plane Function
WP	Work Package
XDP	eXpress Data Path

Executive Summary

The primary aim of BeGREEN D5.1 is to provide the guidelines for project's proof-of-concept (PoC) that will demonstrate project innovations. This document provides the scope and description of these PoCs and related activities that will be performed in Work Package 5 (WP5). The outline of the integrations, final demonstration plans, and detailed project plan for each PoC will be provided in BeGREEN D5.2.

The document starts with reviewing BeGREEN architecture introduced in BeGREEN D2.1 [1], and mapping each PoC to the use cases defined in BeGREEN D2.1 [1]. BeGREEN architecture is proposed based on O-RAN architecture with the introduction of new features to serve BeGREEN innovations, and BeGREEN PoCs are setup alighned with BeGREEN architecture. The document also provides details of testbeds that will be used for PoCs, and the description of BT laboratories testbed is Adastral Park that will be used for final demonstration.

BeGREEN PoCs described in this document are as follows:

- PoC1: BeGREEN Intelligent Plane; that is one of the main components of the BeGREEN architecture. Its main functionality is to expose AI/ML functions to rApps and xApps, enabling the creation of intelligent closed-loop automation empowered by AI/ML, aiming atimproving energy efficiency. Its features will be tested by including it in different PoCs, and at least one AI/ML-based rApp/xApp targeting energy efficiency at the RAN will be showcased.
- PoC2: Sensing assisted communications using RIS; The sensing-assisted communications PoC aims to demonstrate how user cell positioning data coming from sensing services can be used for energy savings in future mobile networks. The main idea of the demo is to dynamically change the reflection properties of a RIS depending on the sensed user location.
- PoC3: Energy efficient CU and O-RAN RIC; This PoC will provide an alternative implementation (acceleration) of CU and O-RAN RIC based on ARM architecture to demonstrate enhanced performance and reduce energy consumption in these components.
- PoC4: Energy-efficient DU implementation using hardware acceleration; This PoC is on implementation of DU high-PHY algorithms, using hadware acceleration techniques, to reduce power consumption compared to legacy implementation. Specific targets are Low Density Parity Check codes (LDPC) decoder and Sphere Decoder.
- PoC5: RU power amplifier blanking; This PoC is on a RU power consumption optimisation technique called "power amplifier (PA) blanking". The PA blanking algorithm reduces the RU power consumption by turning off (blanking) the RF PA, which is the highest power-hungry component in the RU (especially in high power RUs > 1 Watt) at times when there is no data to be transmitted by the RU.

Chapter 4 of this document is dedicated to details of the PoCs, including descriptions, associated use cases (reference to use cases defined in BeGREEN D2.1 [1]), and a high-level plan for each PoC implementation. In addition, a short description of the BeGREEN final demonstrations in Adastral Park, by identifying PoCs that are going to be integrated into this testbed are provided. Finally, the outline of the KPIs that each PoC will cover, and the associated measurements are provided in Chapter 5.

The key contributions and the associated outcomes of this deliverable are the following:

- PoC desriptions (including testbed descriptions) and high-level implementation plans;
- PoC mapping to reference use cases introduced in BeGREEN D2.1, and associated requirements;
- PoC mapping to the project Key Performance Indicators (KPIs).

1 Introduction

The BeGREEN solution promises extensions to the baseline 3GPP and Open RAN (O-RAN) architectural blueprint that target energy efficiency improvement and reduced energy consumption in all RAN elements. These extensions can be summarized as: 1) provisioning of an Intelligent Plane, 2) provision of additional functions and methods that are directly applicable to elements like the Radio Unit (RU), Distributed unit (DU) and Centralised Unit (CU), 3) implementation of RAN functions in a different architecture or hardware (e.g. acceleration), and 4) additional elements such as Reconfigurable Intelligent Surfaces (RIS) and relays, and functionalities that the network can accommodate such as Integrated Sensing and Communication (ISAC).

One of the main goals of BeGREEN Work Package 5 (WP5) in the project is to demonstrate the energy-saving enablers that are developed in WP3 and WP4. This will be done by preparing the proofs-of-concept (PoC) defined by the project. These PoCs ae designed and implemented with accurate alignment with the BeGREEN architecture. A set of these PoCs will be integrated into British Telecom (BT)'s Adastral Park testbed for the final demonstration of the project.

On the one hand, to fulfil the main objectives of WP5, Task 5.1 (T5.1) specifies the use cases against which the solutions developed by BeGREEN can be demonstrated and tested. Also, T5.1 includes the measurements needed to achieve the KPIs defined in T2.1 and a plan for designing the required testbeds, including the timescales. On the other hand, T5.2 will create a detailed plan for developing and implementing each use case, including defining the responsibilities for the provision of components and the definition of interfaces. Finally, T5.3 is on project's final demonstrations that is deployment of some of the PoCs developed onto the final testbed demonstration at BT Laboratories, plus providing additional PoCs deployed in other partners' premises.

This document, BeGREEN D5.1, starts presenting an overall view of the BeGREEN architecture in Chapter 2. Chapter 3 introduces testbeds used for PoC implementation and summarizes the main technologies available on each testbed. Chapter 4 presents the description and use cases corresponding to each PoC and their initial planning. Chapter 5 lists PoCs that are expected to be demonstrated at Adastral Park. Finally, Chapter 6 presents the mapping of project Key Performance Indicators (KPIs) to the PoCs.

2 BeGREEN Architecture

BeGREEN PoCs are designed accurately aligned with BeGREEN's architecture. This chapter summarises the initial BeGREEN architecture as it was first introduced in BeGREEN D2.1 [1].

2.1 BeGREEN architecture description

BeGREEN architecture builds upon the 3GPP/O-RAN architectural framework, which ensures compatibility between all the technological enablers and components designed in BeGREEN project and the 3GPP/O-RAN building blocks [2].

The main goal of the O-RAN Alliance is to promote an open and interoperable RAN architecture for 5G and beyond. As such, the O-RAN architecture involves RAN functions virtualisation, i.e., vRAN, and disaggregation, to promote innovations as it allows vendors to devise innovative solutions for various RAN components rapidly. To optimise the RAN performance, O-RAN defines two different controllers, namely the near-Real Time (RT) Radio Intelligent Controller (RIC) and the non-RT RIC, with different time-scale control loops, i.e., between 10ms and 1s for the near-RT RIC, and higher than 1s for the non-RT RIC. The near-RT RIC enforces RAN policies that require relatively fast adjustment of the network parameters while the non-RT RIC enforces RAN policies for the reconfiguration of network configuration parameters that do not require to be adjusted with such a high periodicity.

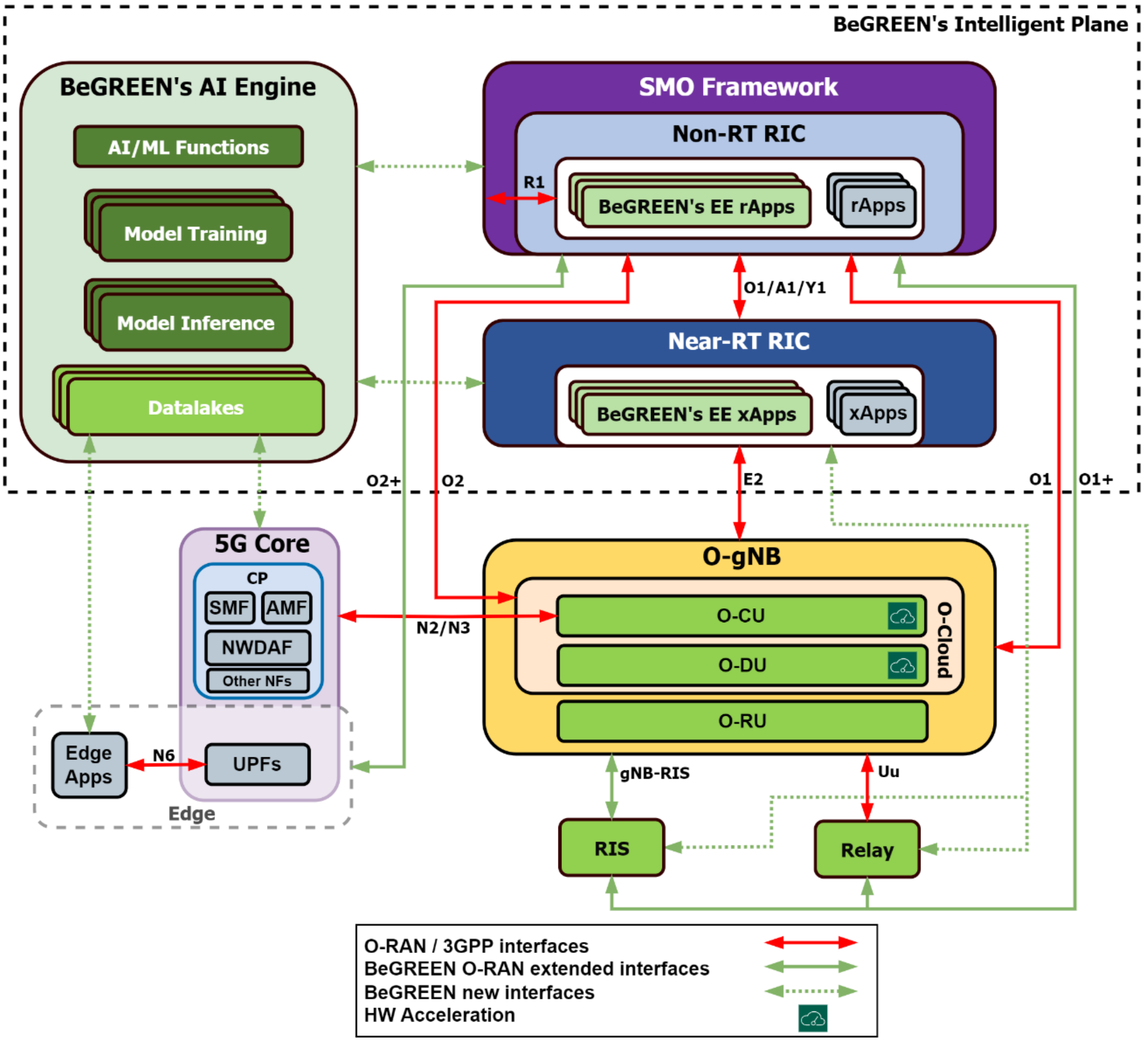

Figure 2-1 BeGREEN proposed architecture.

Figure 2-1 shows the BeGREEN architecture, which is a refined and condensed version of that presented in deliverable D2.1 Figure 4-12. Figure 2-1 provides a clearer view of what the BeGREEN components that are highlighted in green. The architecture is centred around two key ideas:

- First key idea is the introduction of the Intelligent Plane to the BeGREEN architecture. The Intelligent Plane includes rApps and xApps deployed in the O-RAN RICs equipped with the Intelligent Plane's AI Engine, which enables them to create, train, and deploy AI/ML models, i.e., the AI Engine manages the lifecycle of the AI/ML models. Intelligent Plane's AI Engine ingests the E2 nodes and O-Cloud monitoring data in different datalakes. The AI Engine model training component uses this data to train AI/ML models that will be exposed to rApps and xApps. Adding the Intelligent Plane to the O-RAN architecture fuels the adoption of AI/ML-powered rApps/xApps.
- Secondly, BeGREEN integrates Relays and RISs into its architecture. Relays can enhance signal coverage and capacity by processing and retransmitting incoming signals between base stations (BSs) and user equipment. It helps to extend network coverage and maintain seamless connectivity in areas with signal gaps or high traffic. Similarly, RIS also improves signal quality and coverage of radio waves. However, RIS consists of a surface covered in reconfigurable elements that can adapt to optimise signal paths. Using both technologies, mobile operators can provide a more comprehensive and high-quality service to their users, ensuring better coverage, capacity, and signal quality, increasing overall energy efficiency. For more details on the architecture, the reader is referred to BeGREEN deliverable D2.1 [1].

To measure the effectiveness of the different PoCs developed and their relationship with the project KPIs, measurements from the components involved in those PoCs will be collected and used to evaluate several of BeGREEN KPIs. Chapter 6 of this document describes BeGREEN PoCs mapping to the project KPIs.

In what follows, we outline the baseline monitoring in O-RAN underlying the BeGREEN's project. BeGREEN enhanced the O-RAN mechanisms, adding different centralised datalakes that keep all the measurements.

O-RAN monitoring data has two main sources: 1) O-Cloud platform, and 2) E2 nodes. First, the data coming from the O-Cloud platform is divided into two groups, i.e., O-Cloud operations measurements, and O-Cloud performance measurements. The metrics from the O-Cloud operations refer to deployment and infrastructure operations telemetry which monitors the deployed instances and infrastructure health. O-Cloud performance monitoring data includes operational information related to O-Cloud resources. Performance information provides O-Cloud operators with a sense of how well the system is behaving. Performance measurements are typically captured periodically using performance monitoring jobs, which add the measurements to a database. Second, xApps collect data from E2 nodes such as the eNB, gNB, or just a part thereof, e.g., CU orDU. The data collection of the vRAN components is done through the E2 service models of the vRAN functions, which define the type and frequency of data, e.g., physical layer measurements, sensing information, etc.; reporting through the E2 interface. E2 node metrics through the E2 interfaces are collected using xApps that push the data to databases.

In addition, BeGREEN architecture also collects measurements from sources other than RAN. It collects measurements from external components such as the 5G Core (5GC) and AI edge-computing applications. This enables AI/ML models deployed within rApps/xApps to use a broader set of data analytics leading to energy efficiency enablers that can enforce policies in the User Plane Function (UPF) according to traffic predictions or that can tune the RAN configuration from the perspective of AI edge services deployed.

BeGREEN architecture pushes all metrics to AI Engine datalakes. The measurements collected are used to evaluate PoC KPIs.

2.2 PoC mapping to reference use cases

Five PoCs will be showcased that are detailed in Chapter 4. Table 2-1 provides a clear mapping of the reference use cases (defined in D2.1 [1]) to the PoC within the BeGREEN project. This mapping is essential for understanding how the project's conceptual ideas are being translated into practical, real-world applications. By aligning the reference use cases with the PoC, the reader can better grasp the project's progress and the specific scenarios in which it will make a meaningful impact. This table serves as a valuable reference tool to bridge the gap between theory and implementation, ensuring that the BeGREEN project moves efficiently towards its goals.

Table 2-1 Reference Use Cases to BeGREEN's PoCs Mapping

		PoC 1	PoC 2	PoC 3	PoC 4	PoC 5
		Intelligent Plane	Sensing- assisted communications	Energy efficient implementation of CU and RIC	DU HW accelerator	PA blanking demo
Physical layer reference use cases	B5G energy efficiency enhanced RAN through relay nodes					
	Densification of the radio access architecture		X			X
	General mMIMO incentives		X		X	
	Reconfigurable Intelligent surfaces (RIS)		X			
System-level reference use cases	Energy efficiency in vRAN deployments with shared computing infrastructure			X		
	Joint Orchestration of vRAN and Mobile Edge AI Services				X	
	Traffic-aware management of NFV user-plane functions					
	RIC driven energy-efficient RU on/off control	X		X		X

3 Testbed Descriptions

3.1 5G SNPN integration laboratory

In Accelleran (ACC) labs, in Antwerp, Belgium, ACC has developed a lab environment for Open 5G Standalone networks targeting private Standalone Non-Public Network (SNPN) and shared RAN neutral hosting. ACC dRAX product development is a cloud-native, O-RAN-aligned product, comprising a loosely coupled CU (so called control user plane separation, CUPS), a distributed near-RT/non-RT RIC, and a RAN-oriented Service Management Orchestrator (SMO)/Dashboard with xApp/rApp software development kit (SDK). The latter serves for building, integration, validation, release management and support. Furthermore, ACC labs support ongoing R&I projects and commercial deployments. ACC also provides commercial O-RAN professional integration services to the industry.

On the scope of BeGREEN, the ACC 5G integration labs consist of two Dell servers incorporating multiple instances of the Core (Open 5GS) connected to diverse instances of the dRAX (CU/RIC/CW) hosted on virtual machines (VMs). These servers also host data buses and interfaces such as Neural Autonomic Transport System (NATS) and Kafka to deliver information between xApps and gateway interfaces for live setups of simulation environments like the RIC Tester. Moreover, the servers host Effnet[1] DU instances that allow connectivity with different types of Radio Unit (RU), including CableFree, Benetel 650 and Benetel 550 radios. These radios are connected to splitters and programmable attenuators to modify the radio environment. Finally, these are connected to shielded boxes or directly to the User Equipment (UE) antenna, depending on the type of test desired.

3.2 Integrated sensing and communication (ISAC) testbed

Within the BeGREEN project, Integrated Sensing and Communication (ISAC) plays an essential role as an enabler for improving the energy efficiency of the next-generation mobile networks. The sensing functionality to be added to the network should be integrated together with the wireless communications system and reuse the same existing radio hardware. Furthermore, the sensing functionality should not impose an additional burden on the existing wireless communications system, i.e., with respect to increased usage of the wireless medium or bandwidth, since this will diminish the gains in terms of energy efficiency.

To showcase the sensing functionalities and advantages of the ISAC system, BeGREEN partners will collaborate in the development of a testbed and showcase different energy efficiency use cases leveraging the sensing functionalities. The developed testbed will be focused on the sensing functionality of the ISAC system. It will use the transmitted data communication frames for sensing purposes. Two different ISAC systems will be demonstrated, one working at mmWaves (60 GHz) and other at Sub-6 GHz.

3.2.1 mmWave testbed

The mmWave testbed operates at 60 GHz ISM band. The main motivation for using this band is its available large channel bandwidth, as bandwidths of 500 and 2000 MHz are available. This enables excellent range resolution for sensing. Additionally, due to the small wavelength, the antenna size is relatively small allowing for a large number of antennas to be mounted in a relatively small antenna array, thus enabling a better angular resolution in sensing scenarios.

Figure 3-1 depicts a block diagram of the ISAC demonstrator at mmWaves with three different entities. MATLAB or Python scripts are running on a personal computer (PC) for preparing the transmitter waveforms and for processing the received waveforms. The block in the middle is a Field Programmable Gate Array

[1] https://www.effnet.com/products/protocolstack-nw/

(FPGA)-based platform that includes a high-speed analogue to digital (A/D) and digital to analogue (D/A) converters, which generate electric waveforms from the data received from the PC. These waveforms are further fed into an analogue frontend, which modulates these signals on a carrier frequency in the mmWave band. The Radio Frequency System on Chip (RFSoC) board is an AMD/Xilinx product shown in Figure 3-2 and its block diagram is given in Figure 3-3. Additional adapter board, not shown in the figure, is attached to this board to enable interfacing the board to the analogue frontend board shown in Figure 3-4. The analogue frontend board consists of a 60 GHz analogue frontend module from Sivers Semiconductors and a custom interfacing board developed by BeGREEN partner Innovation for High Performance Microelectronics (IHP). The developed interfacing board is needed to provide the necessary power supplies and clock signals and to bring the necessary signal to SMA connector which can be easily interfaced. The complete system, showing both RFSoC board as well as the 60 GHz analogue frontend, connected, is shown in Figure 3-5.

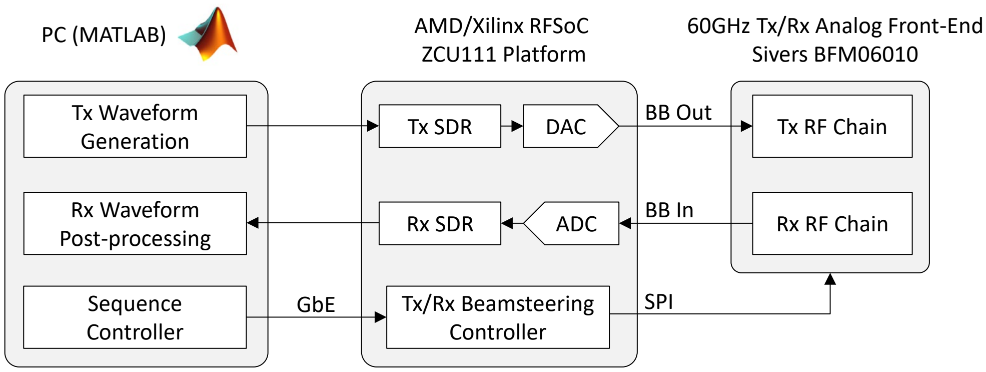

Figure 3-1 Block diagram of the hardware used for sensing in the mmWave band

Figure 3-2 AMD/Xilinx RFSoC board used for waveform generation and acquisition (Source: Xilinx)

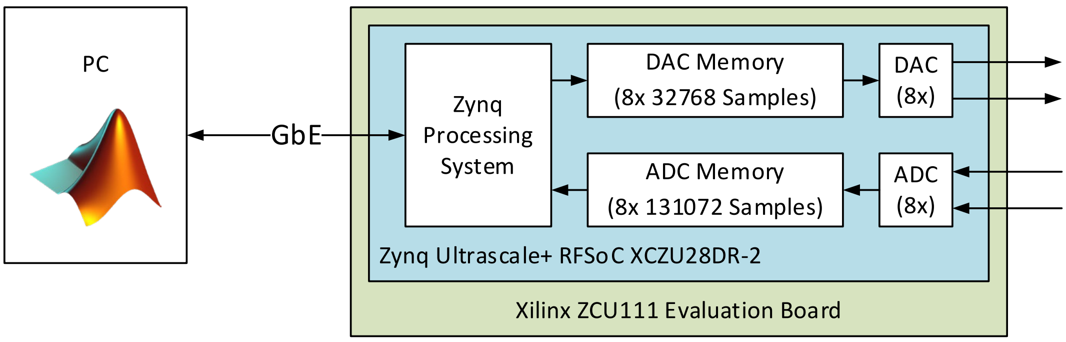

Figure 3-3 Block diagram of the used RFSoC board from AMD/Xilinx

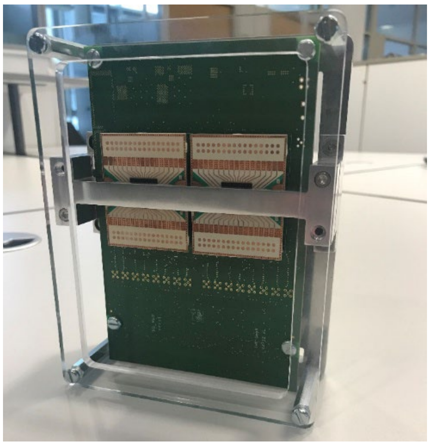

Figure 3-4 60 GHz analogue frontend to be used for ISAC testing

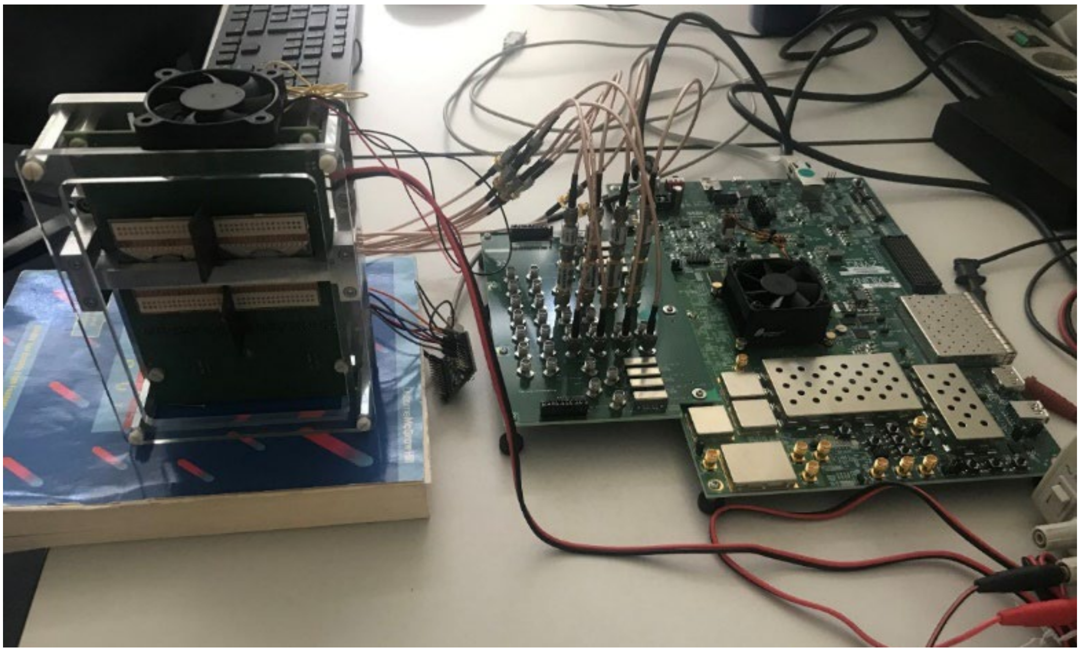

Figure 3-5 RFSoC board connected to the 60 GHz analogue frontend

3.2.2 Sub-6 testbed

The second testbed to be developed is a sub-6 GHz demonstrator. This band is interesting for the equipment vendors since the sub-6 GHz technology is already mature, cheap, and readily available. Nevertheless, the available bandwidth is barely exceeding 320 MHz and due to the larger wavelengths, it will be impractical to have an antenna array with a large number of antenna elements. Thus limiting the angular resolution. Having these limitations in mind, the sub-6 GHz sensing solution can still be of great interest which is the main motivation in BeGREEN for its development.

The sub-6 GHz testebed is consisted of multiple software-defined radios (SDR) from Ettus Research. The preferred model is USRP N321, due to its relatively large bandwidth of 200 MHz, carrier frequencies up to 6 GHz, and the possibility to share the same local oscillator (LO) with multiple devices. Additionally, precise timing synchronisation between multiple devices can be performed. These two functionalities, i.e., LO and timing synchronisation, allow multiple devices to be used as a single device with multiple antennas. Each of these antennas have completely independent RF chains as well as digital chains, i.e., A/D and D/A converters, which enables use of different MIMO signal processing algorithms.

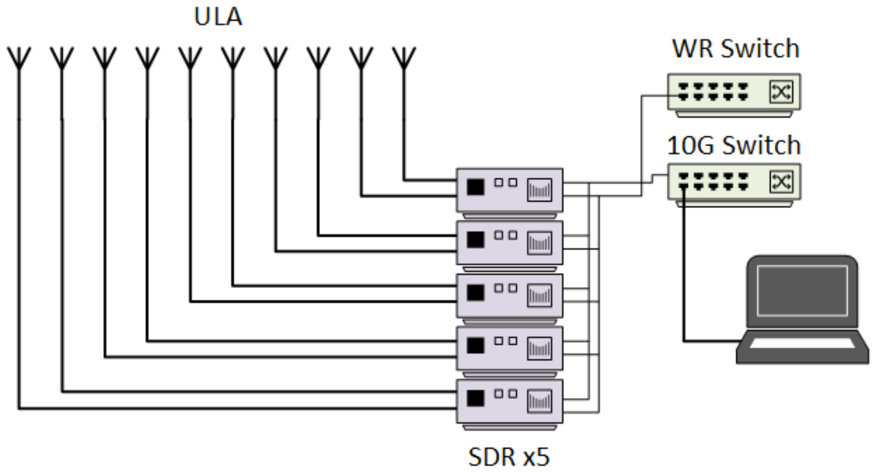

Figure 3-6 Sub-6 GHz sensing system architecture

To demonstrate BeGREEN's sub-6 GHz sensing algorithm, a testbed consisting of 4 to 5 SDRs will be built, each SDR having 2 independent RF chains. This setup will have the possibility to connect an antenna array of up to 8 to 10 antennas. The intention is to build a uniform linear array (ULA) made of patch antennas. This antenna array should have a bandwidth of a few hundred of MHz, which will be sufficient to cover the bandwidth supported by the radios. Finally, these SDRs will have timing and LO synchronisation, making all the channels from the different radios coherent. In Figure 3-6, a simplified architecture of the sub-6 GHz ISAC testbed is shown.

3.2.3 ISAC testbed deployment

Both testbeds will be deployed initially in an anechoic chamber available at the premises of IHP. A photo of the anechoic chamber is given in Figure 3-7. The anechoic chamber has substantial dimensions of 7 × 4 × 2,5 meters, allowing for different objects to be placed in it. These objects will be the obstacles to be used for testing the sensing system.

Additionally, the KPIs of ISAC (Table 6-3) will be evaluated in the anechoic chamber. Further, after characterising the developed system, tests in indoor environment will be performed. These tests include people, objects, and corner reflectors. The main objective is to test the system in realistic environments.

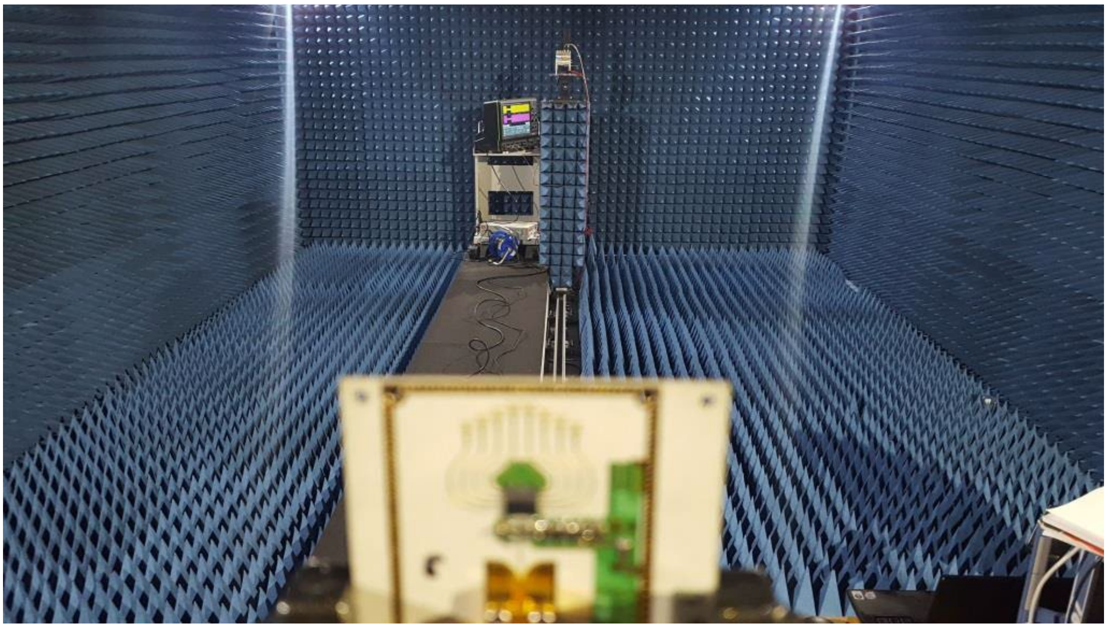

Figure 3-7 Anechoic chamber at IHP

Furthermore, tests with the RIS provided by the BeGREEN partner NEC will be performed. The initial tests will be performed in the anechoic chamber. Since this is a sub-6 GHz RIS, indoor tests would not give results that can show the gains achieved with the use of the RIS. Therefore, depending on the obtained results from the anechoic chamber, additional outdoor tests can be performed.

3.3 DU acceleration testbed

Figure 3-8 shows the DU accelerator testbed at Parallel Wireless (PW). The testbed consists of a DU, a RU, a channel simulator/emulator, and different UEs. First, any DUs under test can be deployed in an x86-based platform, an ARM-based platform and a platform containing a Graphics Processing Unit (GPU), e.g., the NVIDIA Jetson AGX Orin. Second, the carrier frequency ranges supported are below sub-6 GHz. The channel emulator can be used to create the relevant propagation channels (rural or urban) by filtering the signal using a multipath channel impulse response. For example, a Propsim[2] channel emulator can be used to map all UE Tx antenna to all UE Rx antennas. The way this emulator implements the multipath coefficient is by generating several taps, each with a different mean power level. Then, the value of the tap changes as a function of time, for emulating different channel realisations. The delay spread is defined as the Root Mean Square (RMS) delay of the channel.

This testbed allows us to test different MIMO scenarios and to emulate correlations between the different Tx antennas and the different Rx antennas. Thus, this testbed is well suited to test different levels of spatial diversity and challenge the DU under test.

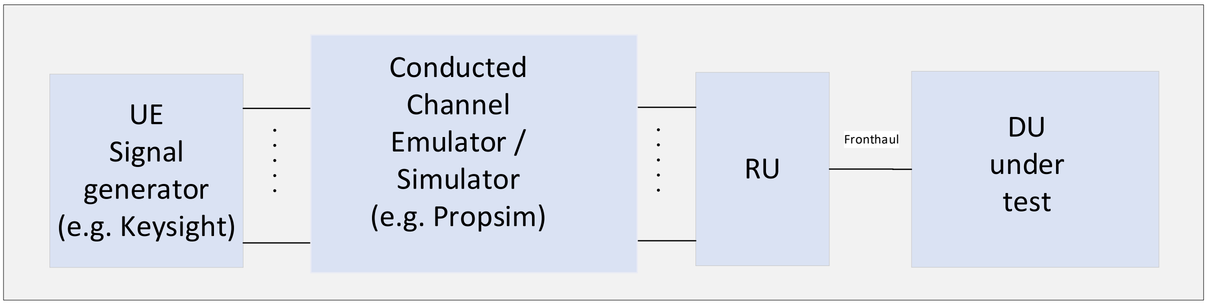

Figure 3-8 DU accelerator testbed

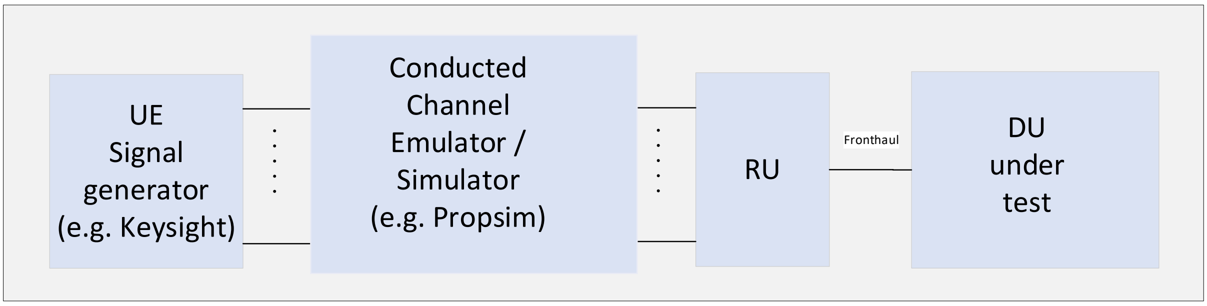

Figure 3-9 DU Hardware accelerator PoC setup

[2] https://www.keysight.com/us/en/products/channel-emulators/propsim-platforms.html

3.4 BT laboratories testbeds in Adastral Park

British Telecom (BT) research and development campus in Adastral Park, Martlesham Heath (near Ipswich), UK, includes laboratories and test facilities aiming to provide a range of deployment scenarios for O-RAN radio access technologies. The main over-the-air testbed consists of six cell site locations spread across the campus with an inter-site distance of approximately 200m. Figure 3-10 shows the locations of the cell-site in a picture of the Adastral Park campus.

Each cell site consists of a lamp post upon which O-RUs can be fitted, GPS antennas if required and a cabinet that can be connected to the server room. The cabinet can host any power supplies required by the O-RU and has a rack-space upon which any edge processing can take place. Fibres from all six sites link back to the server room. Both AC and DC power distribution is available, as well as headend fibre connectivity to cell site locations. Figure 3-11 shows the previous described cell site.

Figure 3-10 Adastral Park Outdoor Testbed

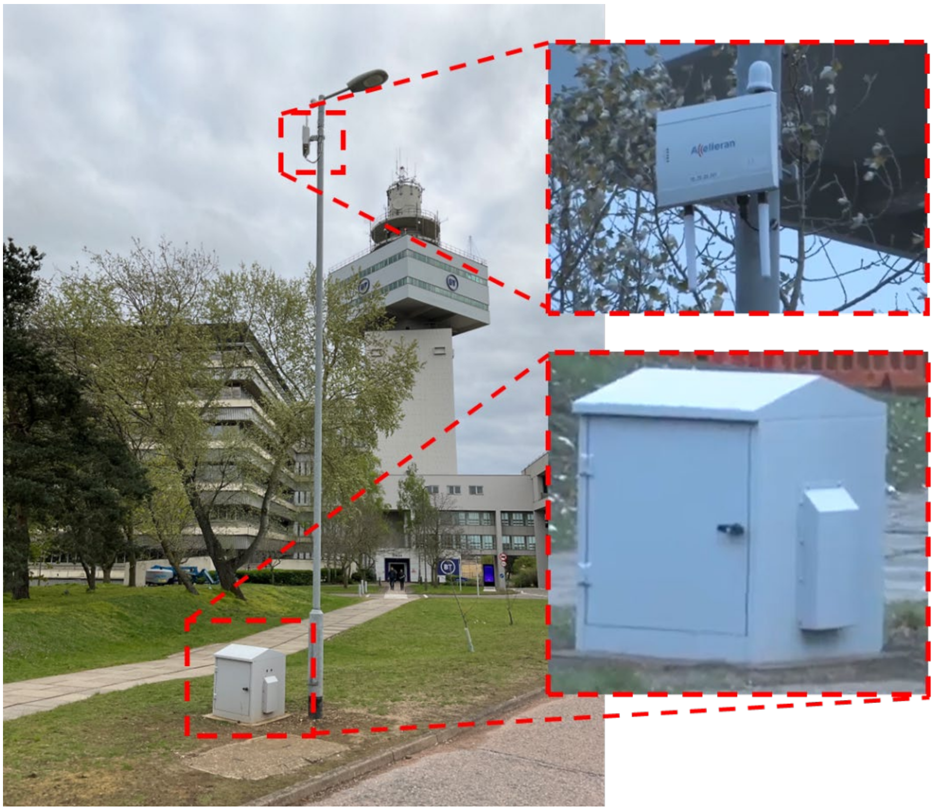

Figure 3-11 Cell Site Design

The server room consists of a fully air-conditioned room where baseband or 19-inch rack-mounted equipment can be installed or hosted. Internet connectivity and secure remote access facilities can be provided to any equipment installed. The TeraVM[3] RIC tester can be deployed in the server room. This E2 node emulator can allow for automated testing of RIC platforms and installed xApps/rApps.

Adastral Park also has two roof-top macro sites, which can be included in the over-the-air testbed if required. They also have power supplies and fibre connections back to the same server room.

For testing, handsets could be provided, as well as network scanners. This can allow for drive tests and walk tests. For testing with handsets, the availability and associated logging tools will vary based on radio configuration and, in particular, the spectrum bands used.

Spectrum licenses will be applied for as and when required. At the time of writing, the testbed currently uses a 40 MHz license in n77 spectrum. Advanced notice, ideally several months, is requested to allow for the spectrum license application process. Ideally, spectrum plans should be shared as early as possible so that preliminary investigations can be done to indicate if such a license is likely to be available.

[3] TeraVM Application Emulation and Security Validation, Viavi Solutions, https://www.viavisolutions.com/en-us/products/teravm

4 BeGREEN Proof of Concepts

BeGREEN PoCs are setup alighned with BeGREEN architecture, and by having in mind novel solutions developed in the BeGREEN technical WPs. PoCs are aiming to serve as experimental evidence for selected project innovations, demonstrating the seamless integration of the different envisioned building blocks as energy-saving enablers in future mobile networks. This experimental validation will confirm the fundamental elements and discoveries of the project.

This chapter describes BeGREEN PoCs that will be deployed for demonstration and validation. These PoCs are designed to test technologies and/or functionalities that play a role in BeGREEN's architectural framework as deecribed in Figure 4-1, and demonstrate specific use cases. Each PoC includes an in-depth description of the main use case that will be shown and the initial plan on how to build the PoC.

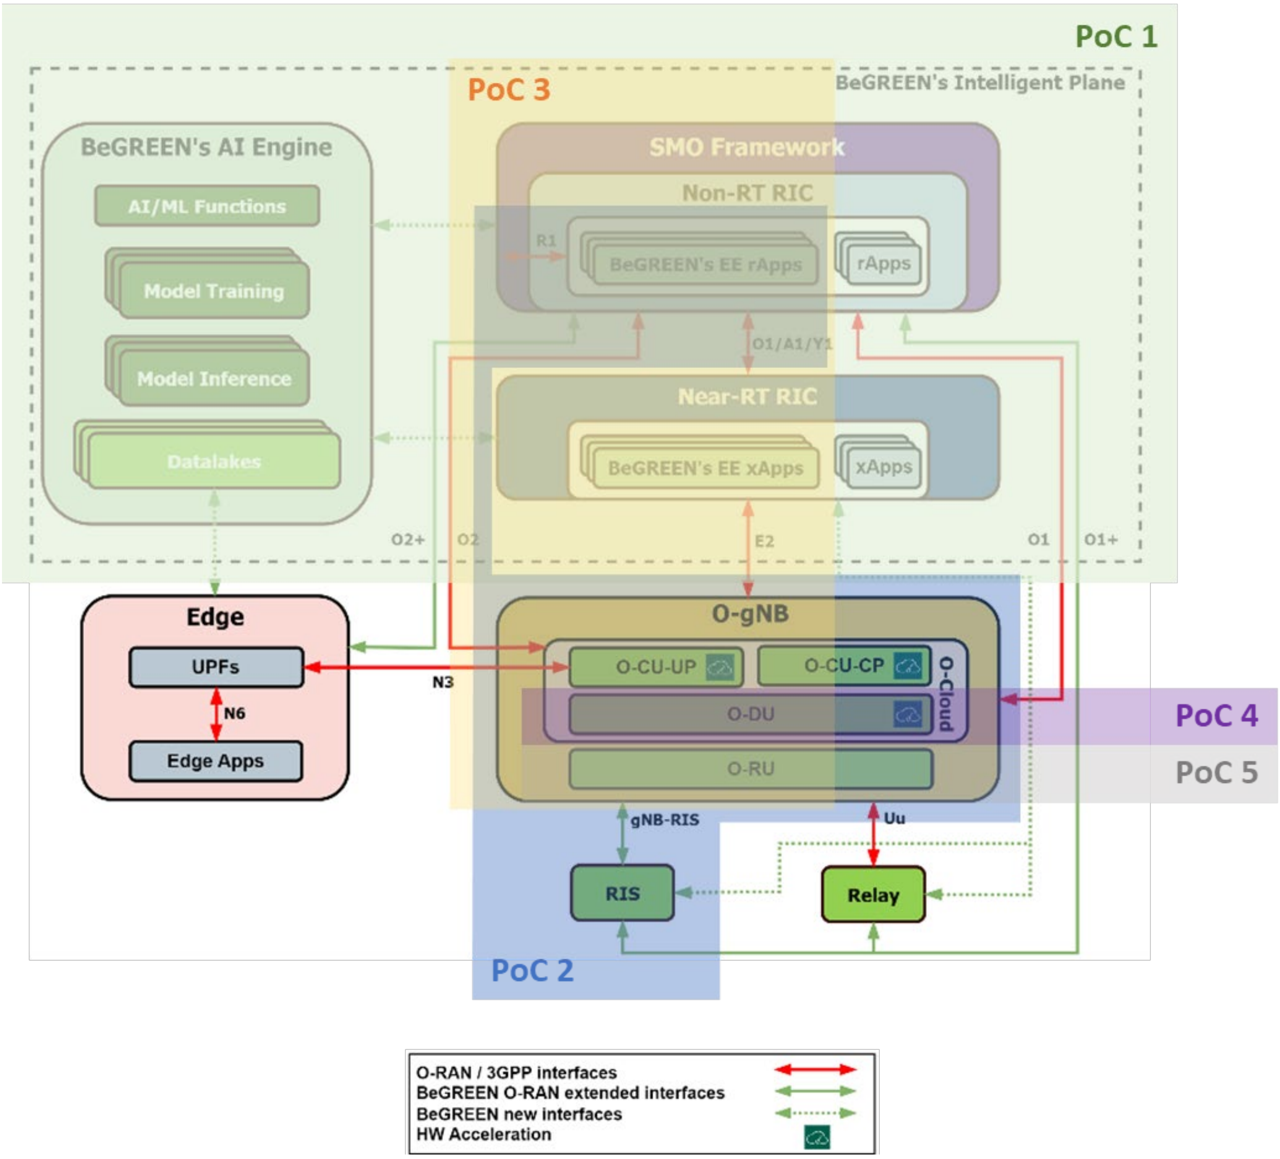

Figure 4-1 Mapping of each of the PoCs to BeGREEN's architectural framework

4.1 PoC1: BeGREEN Intelligent Plane

The Intelligent Plane is one of the main components of the BeGREEN architecture. Its main functionality is to expose AI/ML functions to rApps and xApps, enabling the creation of intelligent closed-loop automation empowered by AI/ML. This automation is aimed at improving energy efficiency. To showcase the features of the Intelligent Plane, several BeGREEN partners will build a PoC testing its features and its different components and at least one AI/ML-based rApp/xApp targeting energy efficiency at the RAN. The next sections describe in detail the PoC, use cases, and the initial planning.

4.1.1 PoC1 description

Figure 4-2 illustrates the components involved in the Intelligent Plane PoC. For a detailed description of the components of the Intelligent Plane, refer to deliverables D2.1 [1] and D4.1 [4]. The PoC considers an O-RAN baseline deployment and adds BeGREEN's Intelligent Plane on top of it. The Intelligent Plane enables the creation of AI/ML-enhanced rApps/xApps. In this PoC, the RAN will be emulated using Viavi's TeraVM to provide the scale and flexibility needed to evaluate the designed rApps/xApps in a broad range of different scenarios. In detail, this will allow testing of the designed optimisations in complex RAN scenarios with multiple cells and users, users' mobility, and different traffic patterns. The final demonstration will include at least one rApp/xApp. The initial planning will include a RAN application to decide the RU on/off switching using AI/ML models provided by the Intelligent Plane. However, we open the possibility of including other xApps/rApps, such as Traffic Steering or Load Predictors applications. In what follows, we outline a description of the main components and subcomponents of the Intelligent Plane PoC and their role in the demonstration:

- **Non-RT RIC**: The non-RT RIC's main role is to host the rApps and expose to them the needed services to implement the intelligent control loop (e.g., telemetry, RAN control, AI/ML). The Intelligent Plane's Datalake integration with O-RAN will enable gathering RAN data from the E2 nodes emulated with TeraVM (raw or processed data). This data will be used to train AI/ML models that will be deployed in the rApps to perform RAN optimisations. Although the final rApps to be demonstrated will depend on WP4 outputs, the final rApps will include one of the following operations:
    - rApp RU A1-control: The rApps developed will be able to provide policies to the near-RT RIC so that they can be enforced in deployed RUs using E2-control according to a global non-RT view of RAN resources.
    - rApp RU O1-control: The rApps developed will be capable of directly deciding the on/off switching of the RUs through the O1 interface according to RAN status and load predictions exposed by other rApps. They could also cooperate with a Traffic Steering xApp to avoid service interruptions or cell saturation when deciding to switch off or on a cell.

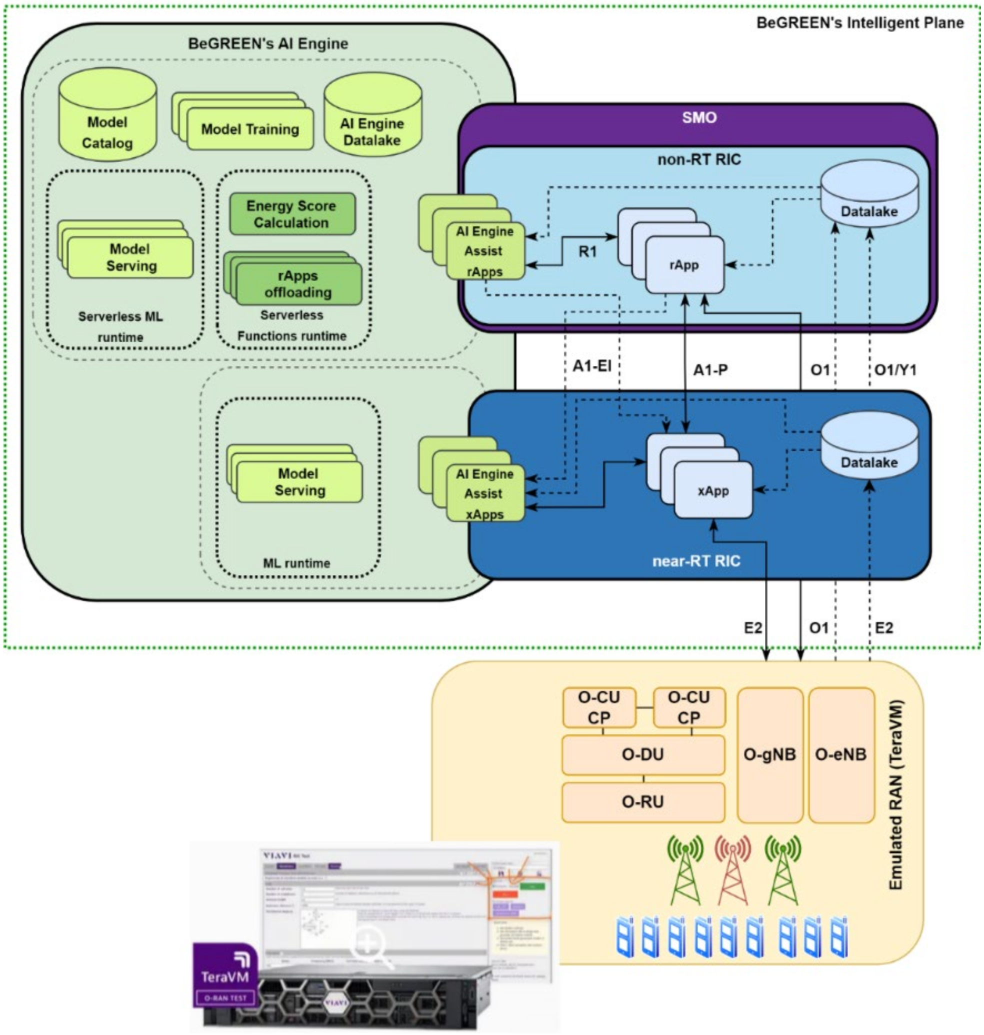

Figure 4-2 Intelligent Plane demonstrator - initial architecture

- **Near-RT RIC**: The Near-RT RIC provides a fast connection with the CU/DU metrics and a fast telemetry gateway to expose radio layer telemetry to the xApps/rApps environment. The data can be used as input to model inference in the developed rApps/xApps. The xApp applications hosted by the Near-RT RIC can support data aggregation, manipulation, and fast pre-processing to control the RAN. The final xApps developed in the project and included in the PoC will include the following operations:
    - xApp RU E2-control: Control the RU using the E2 interfaces to change different configurations such as transmitted power and cell status. It will also support metrics collection using E2 service models.
    - xApp Energy Savings interconnection control: Creation of an Energy Saving xApp based on the policies defined on the non-RT environment that control the energy of the RU equipment based on localization of UE and RU entities. This xApp should support smart handover and traffic management xApps to provide higher Energy Efficiency without degrading the QoS of the connected UEs.
- **AI Engine**: The AI Engine will be the platform that hosts the components that allow training and serving the AI/ML models used by the rApps and xApps, using a loosely coupled serverless approach. The concrete implementation of the AI Engine and the AI/ML functions as external components, or as part of the SMO, the non-RT RIC and/or the near-RT RIC is still under definition in WP4.
- **Emulated RAN**: The TeraVM RIC Test provides a platform for emulating RAN measurements to assess

the efficiency of the RIC and the xApp/rApps created on it. TeraVM, a virtualized testing tool, offers swift adaptation of inputs for testing purposes. It can support E2 and O1 interfaces, to control the radio environment and send commands over it for a proper interaction with a northbound layer. The RIC tester can support a large volume of cells and UEs, allowing intelligent algorithms to work in a realistic and large environment for testing.

The components described above will be integrated and evaluated into the Intelligent Plane demonstration according to the following use cases:

- **PoC1 - use case 1 – Intelligent Plane baseline:** This first use case will entail the validation of the Intelligent Plane baseline architecture comprehending the AI Engine, the RICs and the RAN emulator. First, the basic O-RAN operation will be evaluated, such as supporting the different required interfaces enabling RAN node control and monitoring (e.g., E2, O1, A1…). Then, we will consider the integration of the AI Engine by implementing and evaluating the serverless environment. Finally, we will demonstrate the support of AI/ML workflows including model training and inference, and its exposure to rApps/xApps.
- **PoC1 - use case 2 – RU on/off control through the Intelligent Plane:** This use case will validate the utilization of the Intelligent Plane to empower rApps/xApps targeting Energy Efficiency. The main focus will be RU on/off control through an xApp according to the monitored status of the network (e.g., user location and traffic demands). In addition, the xApp operations will allow configuration through A1 policies, which a rApp will make according to AI/ML-based predictions. Finally, as mentioned in the description, the evaluation of an rApp performing on/off switching through O1 will be also considered.

These two use cases will be evaluated by means of different KPIs and functional validations. On one hand, the proper execution of both use cases will validate the implementation of the Intelligent Plane framework, including the RICs, the datalakes, the AI Engine and related functions. Additionally, in the case of the evaluated ML models, we will consider model performance indicators such as accuracy or precision. On the other hand, use case 2 will be evaluated measuring the energy consumption and the achieved throughput, characterising the performance of the developed xApps/rApps in terms of energy efficiency and QoS.

4.1.2 PoC1 initial planning

Table 4-1 describes PoC versions of the Intelligent Plane demo and the envisioned planning. The use case 1 that focuses on the Intelligent Plane baseline, will be demonstrated with V5 of the PoC. Then, next versions of the PoC will be devoted to the implementation of the RU on/off control through the Intelligent Plane functionalities, to be demonstrated with V8.

Table 4-1 Initial Planning of PoC1

PoC 1 Version	Description	Functionalities	Required Developments	Month
V1	RICs integration with TeraVM	Integration of RICs and TeraVM interfaces (i.e., O1, A1, E2, TeraVM APIs)	Leverage/extend interfaces developed in previous projects to support BeGREEN features.	M13
V2	RICs integration with AI Engine	RICs interfaces to AI Engine	AI Engine components. Exposure of the AI/ML functions to rApps/xApps.	M15
V3	RAN monitoring operations	RAN data from TeraVM is exposed to AI Engine and xApps/rApps	BeGREEN datalakes and related interfaces and functions.	M15
V4	RAN control operations	xApp/rApp doing basic RAN control in TeraVM	rApp/xApps implementing RAN control logic through O-RAN interfaces.	M17
V5	AI Engine operations Demonstration of Use case 1 – Intelligent Plane baseline	ML model inference/training using AI Engine. Exposure of ML model outputs to rApps/xApps.	ML model to be trained and served. rApp/xApps assisting ML workflows. Exposure of ML model outputs to rApps/xApps through R1.	M19
V6	Integration of rApp/xApps devoted to RU on/off control	RU on/off control through rApps/xApps is integrated in the Intelligent Plane plus TeraVM.	rApps/xApps performing the RU on/off control in TeraVM.	M26
V7	Automated RU on/off control through Intelligent Plane	RU on/off control automated and optimized through Intelligent Plane	Logic to manage RU on/off control integrated in the Intelligent Plane. Evaluation through TeraVM.	M28
V8	Demonstration of Use case 2 – RU on/off control through the Intelligent Plane	Demonstration of a scenario comprehending the main outputs of Intelligent Plane.	Definition of the scenario for the final demonstration.	M30

4.2 PoC2: Sensing assisted communications using RIS

ISAC is one of the envisioned technologies that mobile networks beyond 5G will include, enabling new solutions in a broad range of applications. Sensing data enables a new plethora of data, including environmental, location, transportation, or even healthcare data, which will fuel new mobile services. The sensing-assisted communications PoC aims to demonstrate how user cell positioning data coming from sensing services can be used for energy savings in future mobile networks. The main idea of the demo is to dynamically change the reflection properties of a RIS depending on the sensed user location.

4.2.1 PoC2 description

We consider a scenario where an obstacle separates two adjacent RUs that cover different areas. Both RUs have ISAC support, i.e., the RUs can sense the environment, and it is possible to retrieve sensing metrics such as the position of the users in both cells. Furthermore, we have a RIS between both RUs, which we can control and change its reflection properties. Figure 4-3 shows the scenario of the PoC. This scenario enables us to show the following use cases:

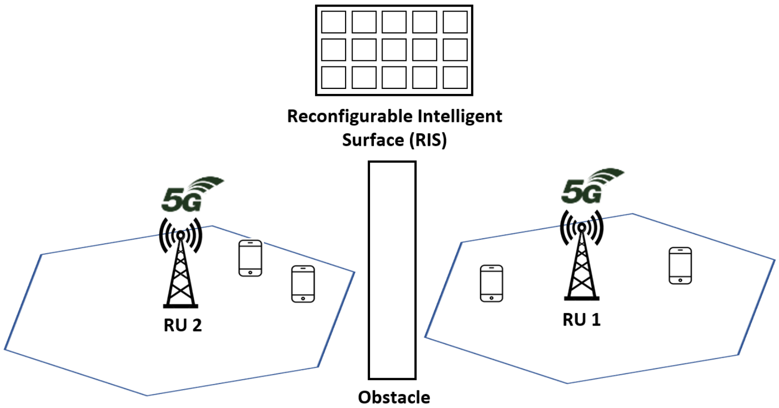

Figure 4-3 Sensing assisted communications PoC scenario

- **PoC 2 - use case 1 - standalone sensing:** Before a complex sensing demo including RIS is set up, an initial testing of the sensing system will be performed. These tests will be performed by setting up a sensing demo including only the sensing system without the RIS. This demo is additionally needed since the used RIS is working in the sub-6 GHz band and will not be able to cover the mmWave band, i.e., cannot be used in conjunction with the developed mmWave sensing system. The demos will be initially set up in an anechoic chamber. They will be used to verify the initial functionality, as well as to demonstrate the capabilities of the sensing system. A sketch of the demonstrator, i.e., the test setup, in the anechoic chamber is shown in Figure 4-4. The setup consists of a sensing transceiver and a few corner reflectors. The corner reflectors to be used are shown in Figure 4-5 and they have well known radar cross section. The positions of the corner reflectors with respect to the transceiver are known. The initial test that will be conducted should show how good the system is able to detect the corner reflectors as well as their position. These tests should be conducted both for the sub-6 GHz and mmWave sensing systems. Additionally, tests outside of the anechoic chamber, i.e., in indoor and eventually in outdoor environments will be performed.
- **PoC2 - use case 2 - RU on/off:** In this use case, we show how an RU can be turned on/off, leveraging the position of its users. As traffic demand patterns change between the day and night, we can turn off one of the RUs offloading the users in a specific area and use the RIS and the second RU to extend the coverage. In this way, we could cover the active users with one RU using a lower energy footprint. Figure 4-6 shows use case 2, where the users of RU 2 are located in an area that the RIS can cover. Thus, we can turn off RU 2 and use the RIS to extend the coverage to that area. However, in this use case, while the RU will not transmit any data, it still has its sensing capabilities.
- **PoC2 - use case 3 - beam steering/tracking**: In use case 3, we extend use case 2, adding mobility to the UE that we offloaded from RU 2 to RU 1. We show how users' position-sensed data can be used to do beam tracking, reconfiguring the reflection properties of the RIS. Moving users can still be covered, and the RU can maintain its sleeping state, leading to substantial energy savings. Figure 4-7 shows use case 3 where we extend use case 2 and add mobility to the users of RU 2 that are offloaded to RU 1.

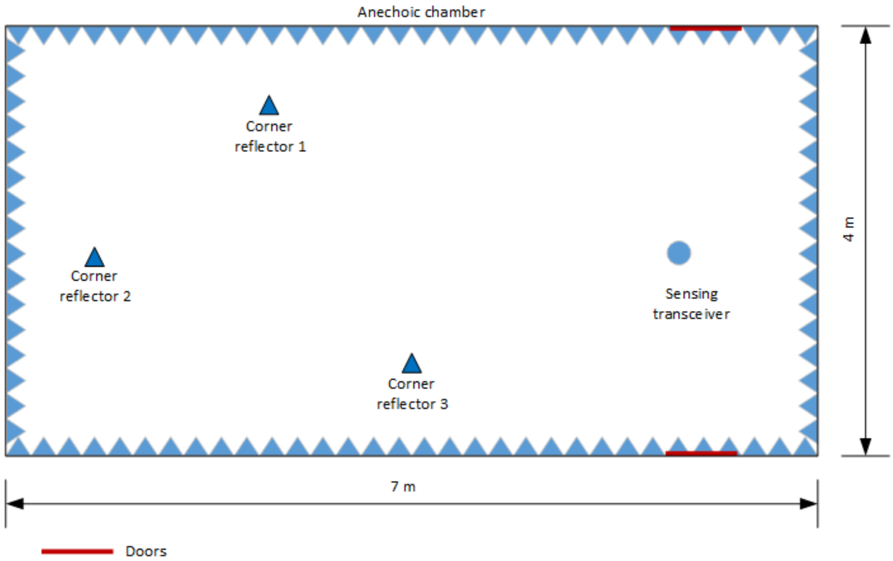

Figure 4-4 Test setup for the sensing demo

Figure 4-5 Corner reflector to be used for tests in anechoic chamber

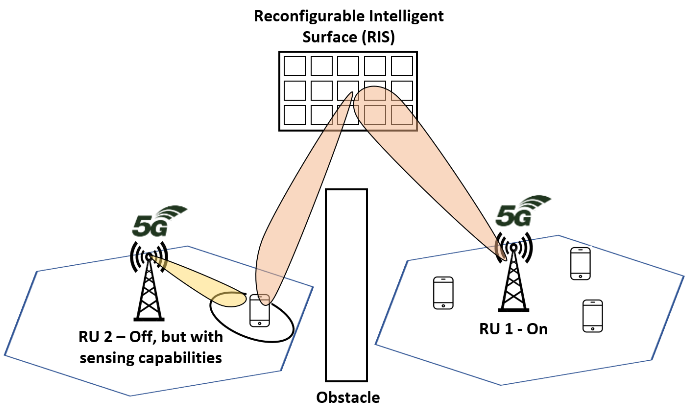

Figure 4-6 Sensing assisted communications demo - use case 2

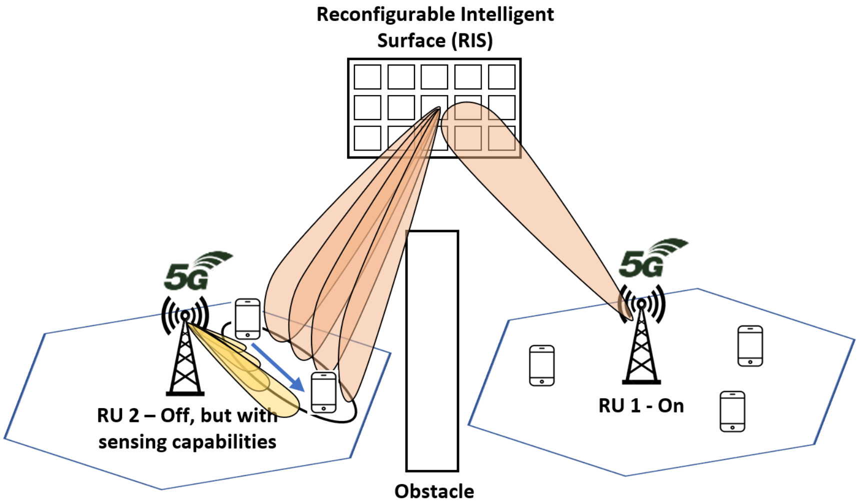

Figure 4-7 Sensing assisted communications demo - use case 3

The use cases previously described, will be evaluated by means of different KPIs. Meanwhile, use case 2 will be evaluated measuring the energy consumption of the setup scenario as the main contribution is to seek opportunities to turn on and off one of the deployed RUs. On the other hand, use case 3 will be evaluated from the point of view of the sensing accuracy showing how accurately we can localize the moving UE and reconfigure the RIS so that the user is still covered.

4.2.2 PoC2 initial planning

Table 4-2 shows the planning of the sensing assisted communications PoC.

Table 4-2 Initial Planning of PoC2

PoC2 Version	Description	Functionalities	Required Developments	Month
V1	Sensing demo with sub-6 GHz and mmWave system. To be used as a prerequisite for the sensing assisted communications demo with RIS	Detection of the objects in the anechoic chamber and their positions. For the tests to be performed out of the anechoic chamber, detection of objects in the surrounding environment.	Development of the necessary hardware and software for sensing	M23
V2	Sensing-assisted demo can successfully demonstrate use case 2	1) The sensing metrics from RU 1 and 2 are pushed to a database which is possible to access. It is possible to retrieve the position of active UEs from the sensing metrics. 2) An rApp reconfigures the reflection properties of a RIS to offload a UE in the coverage area of RU 2 and puts it to sleep mode to save energy.	1) Setting up the necessary demo equipment for RAN functionality and sensing. 2) Verifying that sensing metrics can be retrieved. 3) Verifying that a RIS can be reconfigured, and it is properly calibrated with in the area of the scenario 4) Implementation of the rApp	M24
V3	Sensing-assisted demo can successfully demonstrate use case 3	1) The rApp periodically retrieves the sensed position of the UE and reconfigures the RIS beam according to its position	1) Extending the rApp to periodically reconfigure the RIS reflection properties.	M26

4.3 PoC3: Energy efficient CU and O-RAN RIC

One of the major objectives in improving the performance of the O-RAN components is to reduce energy consumption. Since the CU and RIC are exposed to users, an intelligent control supported by AI/ML services would be desirable in order to materialise the reduction of the energy consumed by the RAN. In particular, using energy optimized hardware such ARM servers can optimize the energy of the overall system. PoC3 is on implementation and demonstration of RAN components (RIC and CU) acceleration to increase performance and reduce energy consumption.

4.3.1 PoC3 description

Figure 4-8 shows the RAN HW setup of PoC3, which follows a baseline O-RAN deployment. In particular, there is a x86 server hosting either the CU and/or the RIC components. The main goal of PoC3 is to compare the performance of the O-RAN CU and RIC components implemented on a x86 platform and on an ARM server. Thus, PoC3 will port the x86 CU and RIC implementations into ARM servers and test whether the ARM implementation improves the energy consumption.

The first step in PoC3 is porting the CU and O-RAN RIC into the ARM server and compare the performance in terms of energy, power consumption and throughput with the original implementation on x86. Figure 4-9 describes the three major scenarios where i) both CU and RIC are deployed on independent x86 servers, ii) CU is deployed on an ARM machine while the RIC is deployed on a x86 server, and iii) where the CU is deployed on a x86 server while the RIC is deployed on an ARM server. These three scenarios will determine the baseline of the system. On top of these scenarios, some mechanisms to enhance the system are foreseen, for instance, the implementation of the eXpress Data Path (XDP) mechanism aiming to improve the data traffic on the CU. The XDP is a mechanism where a XDP is hooked into the NIC and uses the Berkeley Packet Filter (BPF) to intersect GPRS Tunnelling Protocol (GTP) messages and deliver them directly to the CU framework, overlooking its transit over the Linux Network Stack (LNS). The LNS will manage GTP traffic, control traffic and first GTP messages. In high demanding traffic applications, such as those involving vast telemetry or in applications with a data hungry environment, as it is the case for the BeGREEN architecture, it is necessary to support fast and reliable xApps/rApps based on AI/ML learning algorithms implemented in the RIC providing energy efficient RRM strategies. The demo use cases for PoC3 are as following:

- **PoC3 – use case 1: private network scenario:** The first use case is centred around a small-scale deployment scenario of a private network with a small number of UEs.
- **PoC3 – use case 2: (public) cellular scenario:** The second use case is centred around large-scale deployment scenario of a mobile network with many UEs.

The KPIs for RIC hardware acceleration with an Intelligent Plane are outlined in PoC3. The first KPI targets a power consumption reduction exceeding 20% when operating the CU on ARM architecture while utilizing hardware acceleration for the Packet Data Convergence Protocol (PDCP) of CU-UP. The second KPI considered aims for a power consumption reduction surpassing 20% when running the Near-RT RIC on ARM, combined with hardware acceleration for xApp operations.

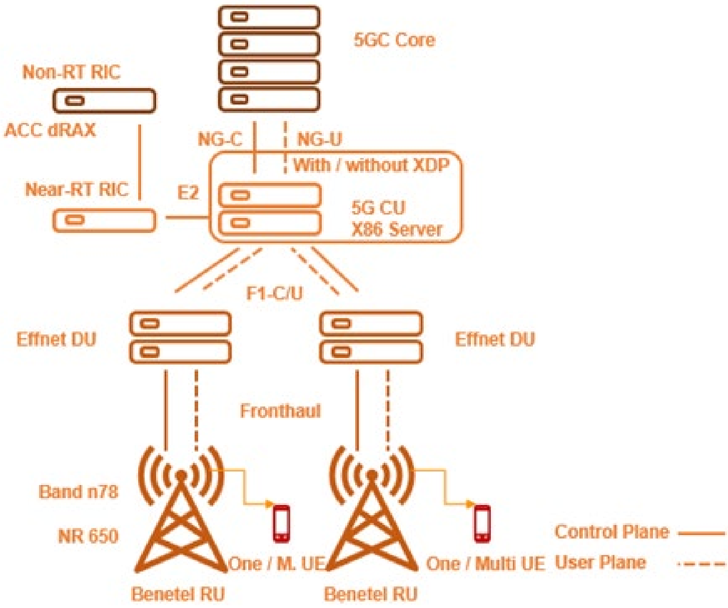

Figure 4-8 RAN CU hardware accelerator demo

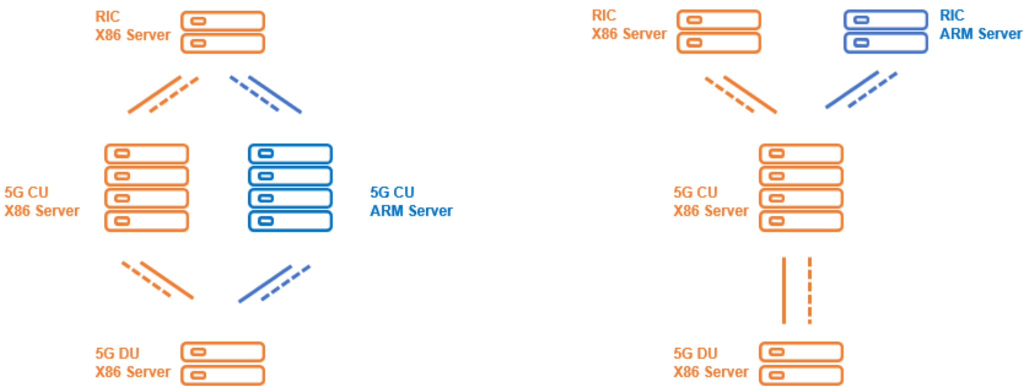

Figure 4-9 CU-RIC HW acceleration proposal. Left) CU energy evaluation. Right) RIC energy evaluation

4.3.2 PoC3 initial planning

PoC3 requires the porting of both CU and RIC into the ARM architecture, and then, the creation of the XDP functionality for both architectures and finally the testing is arranged. The overall performance tests will be organised in single or multiple UEs scenarios. Detailed planning of PoC3 is found in Table 4-3.

Table 4-3 Initial Planning of PoC3

PoC3 Version	Description	Functionalities	Required Developments	Month
V1	CU porting into ARM	dRAX CU implementation working on ARM servers.	Extend Accelleran CU into ARM.	M12
V2	RIC porting into ARM	dRAX CU implementation working on ARM servers.	Extend Accelleran RIC into ARM.	M14
V3	XDP implementation	XDP function ability either for x86 and ARM servers.	Develop the XDP CU enhancement for the performance of the system.	M13
V4	CU in ARM Testing	Single test to probe the capabilities of the CU - RIC on ARM servers	V1 ready	M16
V5	RIC in ARM Testing		V2 ready	M18
V6	Full demo testing	Full capabilities for HW acceleration	Integration of the enhanced capabilities with single UE and multiple UE	M22

4.4 PoC4: Energy-efficient DU implementation using hardware acceleration

This PoC is on implementation of DU high-PHY algorithms, using hadware acceleration techniques, to reduce power consumption compared to legacy implementation. Specific targets are Low Density Parity Check codes (LDPC) decoder and Sphere Decoder as described in [2]. These methods will be demonstrated running over operational platforms.

4.4.1 PoC4 description

PoC4 will be implemented in the laboratories of PW with actual HW using standard test equipment in conducted mode. The setup is shown in Figure 4-10.
PoC4 will show how well different DUs under test do on the power consumption KPIs and receiver performance KPIs. Specifically, PoC4 aims to assess how much power reduction can be achieved using any of the suggested BeGREEN DU architectures (i.e., ARM and ARM+GPU) and compare them to legacy architectures (i.e., x86). PoC4 use cases are as following:

- **PoC4 - use case 1 – testing of DU algorithms in simulation mode:** As a preliminary stage of any DU development and related PoC, before running the DU under test in real-time, a stage of running the platform in simulation mode is performed. This enables better debuggability and helps with algorithms development and tuning. A Matlab simulation is used for generating the UE signals, and for simulating the fading and channel impairments. The UE signal after undergoing the fading channel and impairments, is then injected as a baseband IQ file into the DUT memory, and from there on the DUT will run the code as it would run in a real deployment. Another advantage of this stage is that we can simulate any channel and any impairment, which is sometimes very hard to achieve in a real-time PoC. In this use case, the power consumption of the DU under test can be derived by tools such as the "Jetson Stats Package" [5] and utilisation counters. Receiver performance can be measured using the receiver outputs, for example, by counting how many blocks were decoded successfully, how many blocks failed, and derive the block error rate (BLER).

- **PoC4 - Use case 2 – testing of DU algorithms in emulation mode:** As a second stage the demo will be performed in real-time, after fronthaul integration with an RU. There, a signal generator (such as Keysight) will be used to generate the UE signal in RF frequency and a channel emulator will be used to generate the channel between the UE transmit antennas and the gNB receive antennas. The advantage of the emulation mode demo is that it mimics the operation of the DU in the real product. The channel emulator is a very powerful tool and is used to generate scenarios which emulate real-world deployments. In this demo the power consumption of the DU under test can be either derived like in use case 1 (by tools such as the "Jetson Stats Package" and utilisation counters) or by power measurement tools. The receiver performance can be measured using the test automation entity, which will collect receiver outputs and statistics and analyse them.

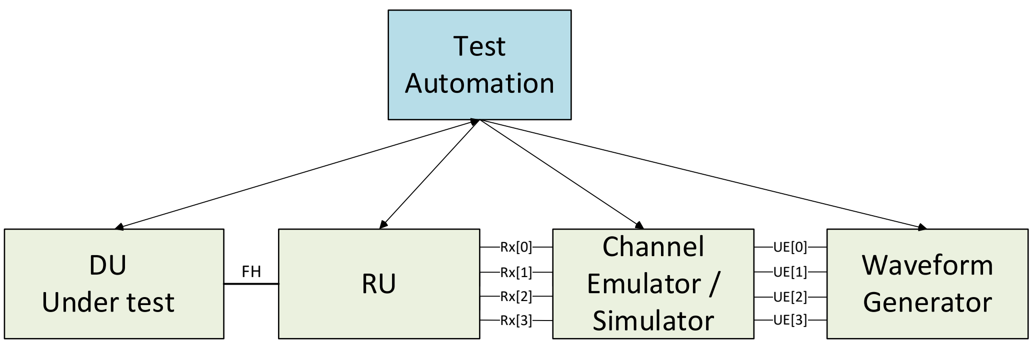

Figure 4-10 PoC4 DU hardware accelerator setup

Table 4-4 PoC4 Scenarios

Scenario	MIMO Mode	# of layers	Target Environment
PoC4 Sen1	SU-MIMO	4	Rural
PoC4 Sen2	SU-MIMO	4	Urban
PoC4 Sen3	MU-MIMO	4	Rural
PoC4 Sen4	MU-MIMO	4	Urban

The simulated scenarios (in use case 1) or emulated scenarios (in use case 2) are listed in Table 4-4.

The intent is that these scenarios will cover different deployment options and different environments. When a UE requires very high throughput and the available allocation is limited, we increase the channel's capacity by introducing SU-MIMO, where multiple layers can be allocated to a single UE. On the other hand, when there are many users requiring UL allocation and the available allocation is not sufficiently large, we use MU-MIMO to be able to multiplex several UEs. The ability to multiplex multiple layers to the same UE or multiple UEs on the same physical resource is possible thanks to the spatial diversity introduced by the channels. When this diversity is large, we can use a simple linear Minimum Mean Square Error (MMSE) equalizer as the demodulator. However, when the spatial separation is not very good (which is a very common case), we would need to use the Sphere Decoder, which performs much better in such scenarios. In the PoC, we will use both linear MMSE and the sphere decoder and compare them while running on the different environments and platforms. In addition to the demodulator, when working with MIMO, as the number of layers increases, the load on the LPDC decoder will increase. Hence, we will also compare the LDPC processing and run it on the different MIMO modes, environments, and platforms.

The PoC will measure various key performance indicators (KPIs) focusing on power consumption. However, accurately measuring and isolating specific algorithm metrics may pose challenges, so the execution times and cycles offer a more accessible means of estimating overall power consumption. The "Jetson Stats Package" is a useful tool for gaining insights into CPU and GPU power consumption. It is anticipated that the BeGREEN implementations on ARM or GPU will exhibit a minimum 15% reduction in power consumption compared to legacy implementations, and a comparison between x86 and ARM architecture factors will be established for benchmarking against the state-of-the-art. In BeGREEN, the modules under development, such as the LDPC decoder and Sphere Decoder, are integral components of the gNB PUSCH receiver, and their performance is evaluated in two key ways, BLER versus Signal-to-Noise Ratio (SNR), with particular interest in the 10% point, and Throughput versus SNR, with a focus on the 90% point. As a fundamental performance benchmark, tests outlined by 3GPP under title "Base Station conformance testing" ([3] Section 8.2.1) will be conducted. Subsequently, BLER and throughput tests representing BeGREEN's use cases, as defined in BeGREEN D2.1 Section 2.3 [1], will be carried out in the next testing phase.

4.4.2 PoC4 initial planning

PoC DU hardware accelerator planning is described in Table 4-5.

Table 4-5 Initial Planning of PoC4

PoC4 Version	Description	Functionalities	Required Developments	Month
V1	Use case 1: Simulation mode demo	DU algorithms running on x86, ARM and GPU	DU algorithms implemented on x86, ARM and GPU	M23

PoC4 Version	Description	Functionalities	Required Developments	Month
		with simulated RF signal	Testing scenario definition and generation	
V2	Platform bring up	Emulated mode setups ready	Fronthaul integration between DU and RU Bring-up of demo setup	M24
V3	Use case 2: Emulation mode demo	DU algorithms running on x86, ARM with emulated RF signal DU algorithms running on GPU with emulated RF signal – nice to have	Testing scenario generation Algorithms modification and improvements	M28

4.5 PoC5: RU power amplifier blanking

Traditionally, the main desire of network operators from the RUs is to obtain high coverage area and high capacity in their networks. These requirements dictate high transmitting power from the RUs power amplifier (PA). Due to the large number of Rus deployed in cellular networks and their low power efficiency (typically 20-30%) they became, by far, the highest power consumer component in the cellular network and one of the major operational expenditures (OPEX) to the operators. PoC5 is on a RU power consumption optimisation technique called "power amplifier (PA) blanking". The PA blanking algorithm reduces the RU power consumption by turning off (blanking) the RF PA, which is the highest power-hungry component in the RU (especially in high power RUs > 1 Watt) at times when there is no data to be transmitted by the RU.

4.5.1 PoC5 description

The diagram in Figure 4-11 shows the main components of the PoC that consist of a 5G end to end system with three main modules:

- REL 5G Sparq-2025-ORU Radio Unit (with PA blanking module).
- Server with (Intel Based) that includes the DU, CU and 5G Core of a 5G SA network as well as a Video streaming application that will create data to be sent over the network to the UEs
- Commercial UE devices.

Here is the relevant use case mapping:

- **PoC5 – use case 1**: PoC5 only use case consists of comparing the power consumption of the RU when the power blanking module is active to the power consumption of the RU when the power blanking module is not active and estimating the power saving that this solution can achieve. The algorithm inspects each orthogonal frequency-division multiple access (OFDMA) symbol that the Scheduler (in the DU) plans to transmit via the RU, and if the next symbol to be transmitted does not include any data, it switches off the RU PA.

To make the PA blanking solution more efficient, it is required to configure the Medium Access Control (MAC) scheduler in the DU to allocate the data to be transmitted mainly in the frequency domain rather than the time domain as depicted in Figure 4-12. This process maximizes the period that PA is turned off for a specific amount of data to be transmitted in a symbol.

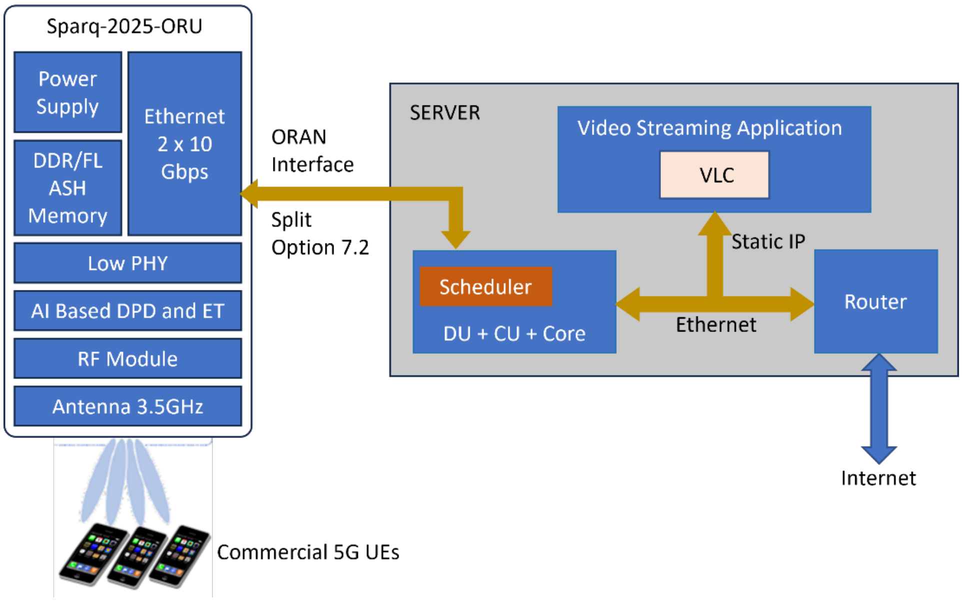

Figure 4-11 RU PA power blanking demo

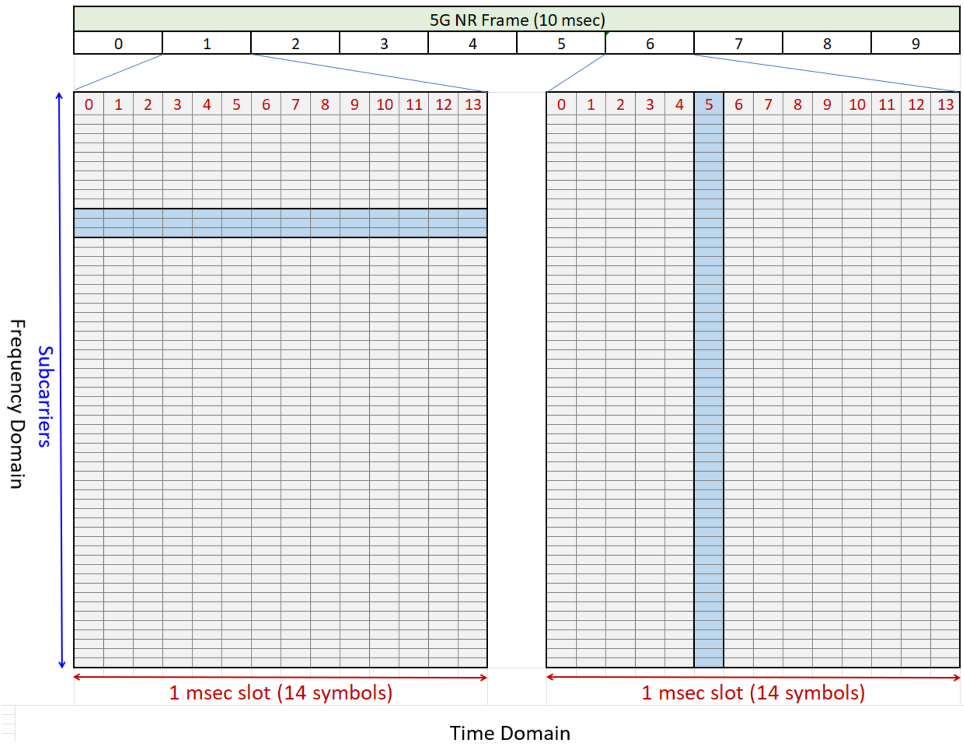

Figure 4-12 5G OFDMA Symbol with data allocated in time domain (right), and in frequency domain (left)

4.5.2 PoC5 initial planning

The initial plan is to deliver the 5G end-to-end system to BT test site in Adastral Park (UK); install the RU in one of the 5G masts in the site and connect it to the Server that will be installed in an indoor cabinet rack. Once that system is installed and powered on, video streams will be delivered from the application server to the commercial 5G UEs connected to the network and the power consumed by the RU will be monitored to measure the power consumption saving achieved when the PA blanking algorithm is activated. Table 4-6 shows the initial planning for PoC5.

Table 4-6 Initial Planning of PoC5

PoC5 Version	Description	Functionalities	Required Developments	Month
V1	RU Power Amplifier Blanking Test	Turn off the Radio Unit Power Amplifier when there is no down link data to be transferred in a specific Symbol	RU internal algorithm that checks the data to be transmitted in the next symbol and turn off the RF module Power Amplifiers when there is no data to be transmitted	M30

5 Final Demonstrations in Adastral Park

The testbed at Adastral Park, described in Section 3.4, will host a number of BeGREEN demonstrations, which will be a subset of the PoCs described in Chapter 4.

The Adastral Park O-RAN testbed will be used to test some of the BeGREEN extensions to its baseline O-RAN testbed. This can provide an additional opportunity for to integrate the BeGREEN PoCs into other O-RAN networks. As we approach the demonstration period, BT will highlight if any other O-RAN systems could be available for testing. For example, the BeGREEN IP demo will use TeraVM RIC tester; but it is also possible to demonstrate functionality over the air if a suitable RAN is available.

Table 5-1 shows the current thinking about which of the BeGREEN PoCs will be demonstrated at Adastral Park.

Table 5-1 Mapping of BeGREEN PoCs to Final Demonstrations at Adastral Park

Demo - Use Cases	Integration in Adastral Park
PoC1 (BeGREEN Intelligent Plane) – use case 1	Yes
PoC1 (BeGREEN Intelligent Plane) – use case 2	Yes
PoC2 (Sensing-assisted communications) – use case 1	No
PoC2 (Sensing-assisted communications) – use case 2	No
PoC2 (Sensing-assisted communications) – use case 3	No
PoC3 (Energy efficient implementation of CU and O-RAN RICs) – use case 1	No
PoC3 (Energy efficient implementation of CU and O-RAN RICs) – use case 2	No
PoC4 (Energy-efficient DU implementation using Hardware accelerators) – use case 1	No
PoC4 (Energy-efficient DU implementation using Hardware accelerators) – use case 2	No
PoC5 (RU Power Amplifier blanking) – use case 1	Yes

6 BeGREEN PoC Mapping to BeGREEN KPIs

The following section will delve into mapping the PoC's use cases to the project KPIs. This mapping ensures we can effectively measure and evaluate BeGREEN's results. BeGREEN's KPIs are our guiding metrics, clearly showing our achievements. We will begin with an overview of the project's KPIs, followed by a detailed mapping of each PoC developed within the project to these KPIs. Finally, we will explore how the measurements gathered from the various PoCs align with the project's KPIs.

6.1 BeGREEN KPI overview

This section will present a comprehensive list of the KPIs that have been meticulously selected to evaluate the BeGREEN technical contributions. These KPIs have been tailored to reflect the project's objectives and commitment to enhancing energy efficiency in mobile networks. Table 6-1 outlines all BeGREEN project's KPIs.

Table 6-1 BeGREEN Project KPIs

KPI ID	KPI Technical Description
KPI-01.1	Development of an energy consumption model for 5G/B5G base-stations, and proposing enhancements reducing energy consumption, while maintaining performance. Applying the output of the work to 3GPP to specific scenarios of practical interest and contributing to the network energy study item in 3GPP Rel. 18.
KPI-01.2	Development of a system level simulator to enable area wide assessment of energy consumption over time, accounting for both traffic growth, service evolution, spectrum availability, user distributions, radio propagation environment parameters, and different proposed architectural solutions.
KPI-01.3	Definition of the appropriate balance between different network evolution strategies to optimize energy efficiency in different target service areas.
KPI-01.4	Creation and demonstration of system level control algorithms, potentially through AI/ML, to achieve the target goal of area-wide network efficiency.
KPI-02.1	GPU optimized mMIMO processing acceleration, including, i) inverse discrete Fourier transform (IDFT)-based Sounding Reference Signal (SRS) channel estimation to provide real-time and accurate channel state estimation required for the beamforming weights generation of up to 64 TRX Massive MIMO array with same level of performance of 4x4, with significant power reduction; ii) zero forcing beamforming weights calculation algorithm to support the real-time generation of up to 16 layers weights with same level of performance with significant power reduction; iii) LDPC FEC decoder to allow the decoding of up to 8 code words with same level of performance with significant power reduction.
KPI-02.2	Accelerated mMIMO processing performance benchmarking against simulation: the goal is to ensure the accelerated design is not including 'short-cuts' to benefit acceleration on expense of algo performance, to ensure proper, close to ideal performance level.
KPI-02.3	GPU based mMIMO processing acceleration utilization and processing time benchmarking compared to accelerated CPU code and accelerated FPGA code where applicable: to inspect compute performance/efficiency of the GPU accelerated code with reference to available accelerated code on other platforms, e.g., FPGA. Concrete example is LPDC compute performance/efficiency can be quantified in terms of the compute time against FPGA setups.
KPI-03.1	RU energy efficiency schemes controllable from both O-RAN rApps and xApps
KPI-03.2	Demonstrated energy reduction of RU of >=40% for emulated day/night (busy/idle) scenario
KPI-03.3	Demonstrated energy reduction of RU of >=40% by use of dynamic AGC, and combined AI-based DPD and Envelope Tracking
KPI-03.4	Total (100%) replacement of RU power source from fossil-based energy to alternative green power energy.
KPI-04.1	Precision of the developed sensing algorithm for detecting potential users

KPI ID	KPI Technical Description
KPI-04.2	Sensing assisted beam search – 20% performance improvement with respect to extensive search and hierarchical search.
KPI-04.3	Detection of users in order to estimate the presumed network load – at least 50% accuracy of estimation of potential mobile users
KPI-05.1	BeGREEN intelligent plane is expected to include the following components an AI-engine, a date lake, and the SMO using R1 interface offering AI/ML based reduction in the energy consumption implemented using RICs and custom rApps/xApps.
KPI-05.2	Overall energy rating is provided at use case and deployment level. With rating A when VNFs are using less energy during less traffic and rating E if VNFs are using higher energy at less traffic conditions.
KPI-05.3	A clear orchestration strategy to achieve improved energy rating from current energy score calculated based on energy influencers, e.g., a strategy to lower the energy score, that indicates key energy influencers are at optimal levels.
KPI-06.1	>20% bare metal server energy consumption reduction at low load with respect to bare metal server energy consumption at peak load, without noticeable impact on user plane traffic performance.
KPI-06.2	At least 2 AI/ML adaptation strategies for energy efficiency demonstrated in laboratory environment with TRL4.
KPI-06.3	An energy rating that effectively rates the VNFs (e.g., in the scale of A to E where A is highly efficient, and E is least efficient). A scoring method (value between 0 to 100) measure the energy efficiency of VNFs with its influencing factors including traffic.
KPI-07.1	>20% power consumption reduction on the server that runs the edge AI service AI service power consumption
KPI-07.2	>20% power consumption reduction on running CU on ARM and HW accelerating PDCP of CU-UP
KPI-07.3	>20% power consumption reduction on running Near-RT RIC on ARM and HW accelerating xApp

6.2 BeGREEN KPIs to PoC mapping

In this section, the analysis on the relation of each how each PoC within the BeGREEN project relates to the project's KPIs. This mapping exercise will provide transparency regarding which aspects of the project each PoC contributes to and how they collectively drive progress toward our overarching goals. Table 6-2 shows the mapping between project KPIs and PoCs.

Table 6-2 BeGREEN KPI to PoC mapping

KPI source	PoC
KPI-02.1	PoC 4 (Energy-efficient DU implementation using Hardware accelerators)
KPI-02.2	PoC 4 (Energy-efficient DU implementation using Hardware accelerators)
KPI-02.3	PoC 4 (Energy-efficient DU implementation using Hardware accelerators)
KPI-03.1	PoC 1 (BeGREEN Intelligent Plane), PoC 2 (Sensing-assisted communications), PoC 3 (Energy efficient implementation of CU and O-RAN RICs)
KPI-03.2	PoC 1 (BeGREEN Intelligent Plane), PoC 2 (Sensing-assisted communications), PoC 3 (Energy efficient implementation of CU and O-RAN RICs), PoC 5 (RU power amplifier blanking)
KPI-04.2	PoC 2 (Sensing-assisted communications)
KPI-04.3	PoC 2 (Sensing-assisted communications)
KPI-05.1	PoC 1 (BeGREEN Intelligent Plane)
KPI-05.2	PoC 1 (BeGREEN Intelligent Plane)

KPI source	PoC
KPI-07.2	PoC 3 (Energy efficient implementation of CU and O-RAN RICs)
KPI-07.3	PoC 3 (Energy efficient implementation of CU and O-RAN RICs)

6.3 Measurements to KPI mapping

The final section will showcase how the measurements from the different PoCs align with our project's KPIs. This mapping is crucial to demonstrate the real-world impact of the project and how it contributes to the attainment of our defined KPIs. These subsequent sections provide an in-depth analysis of the metrics and measurements that underpin our BeGREEN project, helping us to navigate toward a more energy-efficient future in mobile networks. Finally, Table 6-3 shows the mapping between measurements and project KPIs as a means to have a reference on how the measures will be used to compute the KPIs.

Table 6-3 Measurements to KPI mapping

Measurement	Measurement description	BeGREEN KPI	Unit
Energy consumption	Energy consumption refers to the total amount of electrical energy used by a device or system over a specified period. It is typically measured in units such as kilowatt-hours (kWh) or joules (J). Monitoring energy consumption is vital for understanding and managing the power requirements of a device, system, or process. It is essential for optimizing energy efficiency and reducing operational costs.	KPI-03.2, KPI-05.1, KPI-05.2	KWh J
Power usage	Power consumption is the rate at which electrical energy is used or dissipated over time. It is measured in watts (W) and represents how quickly a device or system consumes energy. Power consumption is a critical parameter for assessing the real-time energy requirements of devices. It is crucial in designing and operating energy-efficient systems and prolonging battery life in portable devices.	KPI-02.1, KPI-02.3, KPI-07.2, KPI-07.3	W
Throughput	Throughput is a measure of the amount of data or units of work that can be processed or transmitted within a given time frame. It is typically expressed in bits per second (bps), transactions per second, or other relevant units. Throughput is a key metric in networking, communication systems, and data processing. It helps assess the efficiency and performance of these systems, indicating their capacity to handle data or tasks.	KPI-02.2	bps
Sensing accuracy	Sensing accuracy refers to the precision and reliability with which a sensor or measurement system can detect and report the actual value of the parameter it is designed to measure. Sensing accuracy is crucial in various applications such as scientific research, industrial monitoring, and healthcare. It ensures that the data collected is trustworthy and can be used for decision-making and control.	KPI-04.2, KPI-04.3	m
Block Error Rate (BLER)	BLER is a measurement used in digital communication systems to assess the error rate in the reception of data. It represents the ratio between the incorrect and the correct decoded blocks. BLER is important in evaluating the performance of communication systems, especially in scenarios where data integrity is critical, such as wireless networks and satellite communications. Lower BLER values indicate better data transmission reliability.	KPI-02.2	%
DU processing time	Processing Time measures the time taken by the Distributed Unit (DU) to process a transport block in microseconds, which is a fundamental data unit containing user information, control information, and error correction data. This KPI quantifies the latency associated with the DU's processing of incoming transport blocks, reflecting the responsiveness and performance of the network infrastructure.	KPI-02.3	ms

7 Summary and Conclusions

BeGREEN D5.1 presents a description of BeGREEN PoCs. The document elaborates a list with the PoC descriptions, and the main technologies required for all use cases, and an initial plan for each PoC implementation. Furthermore, the document identifies PoCs that will be integrated into Adastral Park for final demonstration. The outline of the KPIs that each PoC will cover, and the associated measurements are also provided.

The implementation and integration of PoC components are already underway. BeGREEN T5.2 takes care of precise project plans for PoC implementations, the demonstraitons of the results and the integration of the selected PoCs into the final demonstration testbed.

8 References

[1] BeGREEN, D2.1, "BeGREEN Reference Architecture", July 2023.

[2] O-RAN Alliance "O-RAN Working Group 1 (Use Cases and Overall Architecture); O-RAN Architecture Description".

[3] 3GPP TS 38.141-1, "NR; Base Station (BS) conformance testing Part 1: Conducted conformance testing (Release 16)", September 2021.

[4] BeGREEN, D4.1, "State-of-the-art review and initial definition of BeGREEN O-RAN intelligent plane, and AI/ML algorithms for NFV user-plane and edge service control energy efficiency optimization", November 2023.

[5] NVIDIA Jetson Stats package for monitoring and control, https://developer.nvidia.com/embedded/community/jetson-projects/jetson_stats#:~:text=jetson%2Dstats%20is%20a%20package,import%20in%20your%20python%20script.