



***SEamless integrationN of efficient 6G WirelesS
tEchnologies for Communication and Sensing***

D5.3 Lab demos and evaluation

July 2026

6G-SENSES project has received funding from the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme under Grant Agreement 101139282

Project Start Date: 2024-01-01

Duration: 30 months

Call: HORIZON-JU-SNS-2023

Date of delivery: 2026-07-24

Topic: HORIZON-JU-SNS-2023-STREAM-B-01-02

Version 1.0

Co-Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission (granting authority). Neither the European Union nor the granting authority can be held responsible for them.

Dissemination Level: Public (PU)

Grant Agreement Number:	101139282
Project Name:	SEamless integratioN of efficient 6G Wireless tEchnologies for Communication and Sensing
Project Acronym:	6G-SENSES
Document Number:	D5.3
Document Title:	Lab demos and evaluation
Version:	1.0
Delivery Date:	2026-06-30 (2026-07-24)
Responsible:	BubbleRAN (BR)
Editor(s):	Navid Nikaein (BR)
Authors:	Navid Nikaein (BR), Federico Trombetti, Salvatore Pontarelli, Novella Bartolini (UNIROMA1), Khai Nguyen, Mikel Irazabal (BR), Vladica Sark, Mert Özates, Batuhan Sütbas, Jesús Gutiérrez (IHP), Muhammad Umar, Sajjad Ahmad, Padmanava Sen (BI), Luis Díez, Neco Villegas, Ramón Agüero, Alicia Barón, Ignacio Santamaría, Jesús Ibáñez (UC), Xi Ding, Nils Hildebrand, Eduard Jorswieck (TUBS), Ioanna Mesogiti (OTE), Markos Anastasopoulos, Ilias Floudas, Pavlos Doanis, Anna Tzanakaki (IASA), Yazhou Zhu, Jessica Sanson, Valerio Frascolla (INT), Muhammad Saadi, Shahid Mumtaz (NTU).
Keywords:	6G, Testing Methodologies, Testbed, Setup, Proof of Concept.
Status:	Final
Dissemination Level	Public (PU)
Project URL:	https://www.6g-senses.eu

Revision History

Rev. N	Description	Author	Date
0.1	Draft Table of Contents (ToC)	Navid Nikaein (BR), Jesús Gutiérrez (IHP)	2026-04-15
0.2	Chapter 2 contributions	Navid Nikaein (BR), Jesús Gutiérrez (IHP)	2026-07-15
0.3	Chapter 4 contributions	Xi Ding, Eduard Jorswieck (TUBS), Alicia Barón, Jesús Ibáñez, Ignacio Santamaría (UC)	2026-07-17
0.4	Chapter 5 contributions	Germán Castellanos (ACC/FC), Luis Díez (UC), Jesús Gutiérrez (IHP)	2026-07-20
0.5	Chapter 3 contributions	Neco Villegas, Luis Díez, Ramón Agüero (UC), Muhammad Umar (BI), Batuhan Sütbas, Jesús Gutiérrez (IHP)	2026-07-20
0.6	Final Revision of the Document	Anna Tzanakaki (IASA), Ignacio Santamaría (UC), Valerio Frascolla (INT)	2026-07-23
1.0	Last revision and submission of the document	Jesús Gutiérrez (IHP)	2026-07-24

Table of Contents

LIST OF FIGURES	6
LIST OF TABLES.....	9
EXECUTIVE SUMMARY	10
1 INTRODUCTION.....	11
1.1 Last Phase Experimentation Testing.....	11
1.2 Organisation of the Document	12
2 OVERALL 6G-SENSES TESTBED IMPLEMENTATION	13
3 POC#1 – MULTI-TECHNOLOGY ISAC PLATFORM	15
3.1 Multi-technology ISAC RAN implementation.....	15
3.1.1 Final deployment and configuration.....	15
3.1.2 Results and KPIs	19
3.2 Multi-technology transport optimisation	26
3.2.1 Deployment and configuration.....	26
3.2.2 Experimental Validation.....	28
3.3 Other small-scale implementations.....	43
3.3.1 Evaluation of 60 GHz AFE blocks.....	43
3.3.2 Demonstration: LiveSense - Real-Time Wi-Fi Sensing Platform for Range-Doppler on COTS Laptops.....	45
4 POC#2 – CF-MIMO PROTOTYPE	47
4.1 Implementation at TUBS	47
4.1.1 CF-MIMO Prototype.....	49
4.1.2 RIS-based PoC Demonstrator.....	65
4.2 Implementation at UC.....	69
4.2.1 Final deployment and configuration.....	69
4.2.2 Results and KPIs	72
5 POC#3 – NETWORK DIGITAL TWIN (NDT) VIA RAN SENSING.....	81
5.1 Intelligent Plane implementation.....	81
5.1.1 Lab setup under the UC server	82
5.2 Experimental validation of the PoC	85
5.2.1 Final deployment and configuration.....	85
5.2.1.1 Setup 1 Data generation	85
5.2.1.2 Setup 2 Automatic Data Acquisition	86
5.2.1.3 Inference.....	90
5.2.1.4 NDT	92

6	SUMMARY AND CONCLUSIONS	95
7	REFERENCES.....	97
8	ACRONYMS.....	99

List of Figures

Figure 1-1 6G-SENSES Testing and evaluation cycles	12
Figure 2-1 Multi-layer 6G-SENSES Architecture and mapping of PoCs	13
Figure 3-1 PoC#1 implementation at UC.....	16
Figure 3-2 Sensing sources (sub-6 SDR based gNB and non-3GPP mmWave radar) integration within O-RAN instance	17
Figure 3-3 Execution example of the Sub-6 gNB with sensing capabilities	19
Figure 3-4 Execution validation of the sensing supported MAC scheduler	20
Figure 3-5 Execution validation of the sensing SM and end-to-end (E2 agent-xAPP) communication	21
Figure 3-6 Anechoic chamber setup for testing the RIS-assisted sensing.....	23
Figure 3-7 Heat maps for a) 80, b) 85, c) 90, d) 95 degrees reflection angles with respect to the RIS surface	24
Figure 3-8 Stronger reflection received back with use of PA and high gain antenna	24
Figure 3-9 WEC Demo, deploying a MEC infrastructure	26
Figure 3-10 Experimental testbed used in final demonstration for PoC#1 at IASA	27
Figure 3-11 Transport network supporting ISAC. Configuration with emulated air interfaces. Snapshot of 3 moving targets, blue dots are static reflectors	28
Figure 3-12 a) Volume of sensing traffic vs resolution velocity b) Average Optical transport network capacity allocated for sensing with and without optical network switching	30
Figure 3-13 Moving targets as detected by (a) RU4 and (b) RU5, Nframe = 12, W= 100MHz, RAN frequency =3.5GHz	31
Figure 3-14 a) Lab testbed connectivity with electronic and optical data paths, b) physical nodes	32
Figure 3-15 a) Generated Communication and background at Server b) Switch from Centralized to FL learning, c) Total network traffic.....	34
Figure 3-16 RAN reconfiguration time through the SMO/OAM using O1/NETCONF	37
Figure 3-17 CDF Optical Datapath reconfiguration time.....	37
Figure 3-18 Evaluation of external sensing data ingestion through the Non-RT RIC.....	39
Figure 3-19 Experimental multiband sensing and communication testbed. Hardware setup integrating the FR1 OFDM communication/sensing branch and the X-band phased-array radar	40
Figure 3-20 Beam-delay response maps obtained for four reflector placements corresponding to nominal azimuth directions of approximately (a) -60° , (b) -30° , (c) $+20^\circ$, and (d) $+70^\circ$. The dominant reflection shifts consistently across steering directions, demonstrating the angular discrimination capability of the X-band sensing branch.	41
Figure 3-21 FR1 delay-Doppler sensing results for alternating approaching and receding target motion. Positive and negative Doppler signatures are observed on opposite sides of the zero-Doppler axis, demonstrating motion direction discrimination.....	42
Figure 3-22 (a) zoom in for RF-probing, (b) characterization standards for the PCB stackup and transmission results of the Thru standard.....	44

Figure 3-23 (a) manufactured antenna sample on the probe station chuck, (b) simulation vs measured results	44
Figure 3-24 (a) Manufactured balun as a back-to-back connected sample, (b) simulation vs. measurement results	45
Figure 3-25 Microscopic view of the AFE board with four transceiver chipsets integrated in to the RF chain	45
Figure 4-1 Schematic overview of the TUBS CF-MIMO laboratory infrastructure, including time server, Master OctoClock, OctoClock distribution, Ethernet networking, compute nodes, and USRPs	49
Figure 4-2 Joint beamforming concept, redrawn/reused from D5.1	51
Figure 4-3 Distributed/group-wise beamforming concept, redrawn/reused from D5.1	52
Figure 4-4 Physical deployment of the co-located massive-array topology	53
Figure 4-5 Physical deployment of the CF topology	53
Figure 4-6 Measurement and processing workflow. Important: calibration, H-estimation, H-averaging, and beamformer computation must be inside the iteration loop	55
Figure 4-7 Mean received-power matrices for the CF topology, averaged over 31 independent sessions. Rows denote the intended target RX channel, columns denote the observed RX channel, and red boxes mark the intended target position. Left: joint MRT. Right: distributed MRT	57
Figure 4-8 Mean received-power matrices for the co-located massive-array topology, averaged over 31 independent sessions. Rows denote the intended target RX channel, columns denote the observed RX channel, and red boxes mark the intended target position. Left: joint MRT. Right: distributed MRT	58
Figure 4-9 Estimated spectral efficiency for the selected RX pair ch3/ch9, derived from empirical pairwise SINR values. The two bars per case correspond to the two ordered interference directions: target ch3 with ch9 as interferer, and target ch9 with ch3 as interferer	59
Figure 4-10 Pairwise SLNR profile for target RX channel ch9. For each observed non-target RX channel, the plotted value is the received power at the intended target channel minus the received power at the observed RX channel during the same beamforming transmission. Higher values indicate better spatial isolation of the beam. The diagonal entry ch9/ch9 is omitted	60
Figure 4-11 Beamforming margin per target RX channel for both physical topologies and both beamforming methods. Positive values indicate that the intended RX channel received more power than any other RX channel. The cell-free topology maintains positive margins for all targets, whereas the co-located massive topology exhibits systematic degradation at selected RX positions	61
Figure 4-12 Comparison of mean received-power matrices across physical topologies and beamforming methods. Rows correspond to physical topologies and columns correspond to beamforming methods. This overview highlights the stronger spatial selectivity of the CF topology and the broader leakage structure of the co-located massive-array topology	62
Figure 4-13 Numerical simulation of the spatial beampattern based on point-sampled Green model for distributed setup and MRT to one antenna located in the center	63
Figure 4-14 Numerical simulation of the spatial beampattern based on point-sampled Green model for the massive MIMO setup and MRT to one antenna located in the center	64
Figure 4-15 Final deployment of the RIS-based demonstrator. The RF path follows a cascaded TX → RIS1 → RIS2 → RX geometry. A metallic wall blocks or strongly attenuates the direct TX-to-RX path, so that the received signal is dominated by the two-RIS reflection path. RIS1 is configured for the fixed TX-to-RIS2 path, while RIS2 sweeps its horizontal outgoing reflection direction from -90° to $+90^\circ$	67

Figure 4-16 Received signal strength of the cascaded TX → RIS1 → RIS2 → RX link as a function of the configured horizontal output angle of RIS2. RIS1 is kept fixed for the TX-to-RIS2 path, while RIS2 sweeps from -90° to $+90^\circ$. A moving average is applied to the received measurements because the SDR measurement rate is higher than the RIS reconfiguration rate	68
Figure 4-17 Block diagram of UC setup for reciprocity calibration integration in OAI RU gNB.....	70
Figure 4-18 UC CF-MIMO small-scale demonstrator diagram	71
Figure 4-19 UC CF-MIMO small-scale demonstrator setup.....	72
Figure 4-20 OAI gNB executing the reciprocity calibration procedure at UC lab.....	73
Figure 4-21 Estimated reciprocity phase drift versus OFDM subcarrier for each gNB antenna	73
Figure 4-22 Single user downlink performance.....	75
Figure 4-23 Spatial distribution of the average SNDR in the single-user scenario. Left: Cell-Free MRT. Right: Small Cell MRT.....	76
Figure 4-24 Histogram of location-dependent SNDR for Cell-Free MRT and Small Cell MRT	76
Figure 4-25 Reciprocity calibration impact on SINR in multiuser downlink. Left: User 1. Right: User 2	78
Figure 4-26 Power allocation impact on SINR in multiuser downlink: global normalization versus per AP normalization. Left: User 1. Right: User 2	79
Figure 4-27 Power allocation impact on multiuser downlink constellations. Up: User 1. Down: User 2. Left: global normalization. Right: per AP normalization	79
Figure 5-1 PoC#3 implementation Architecture.....	81
Figure 5-2 Schematic organization of the PoC#3 implementation	82
Figure 5-3 Location and physical installation of the Physical Installation of ACC O-RAN BS @ UC testbed ...	84
Figure 5-4 Demo control website	84
Figure 5-5 Date set generation interface.	85
Figure 5-6 Snippet of dataset (JSON format) generated during the validation measurement campaign	86
Figure 5-7 ACC O-RAN solution in OTE testbed.....	87
Figure 5-8 Physical Installation of ACC O-RAN BS @ OTE testbed	87
Figure 5-9 5G COSMOTOOLS GUI used for capturing datasets for NDT training.....	89
Figure 5-10 Control application metrics from RU and UEs.....	90
Figure 5-11 Control application metrics from RU and UEs.....	91
Figure 5-12 Inference box for the control application	91
Figure 5-13 Inference Heat map for inferred location.	92
Figure 5-14 Inference logs for PoC#3	92
Figure 5-15 NDT box for the control application.....	93
Figure 5-16 System metrics for the control NDT and RAN behaviour.....	93
Figure 5-17 Functional architecture of the NDT, based hybrid simulation, to feed the Intelligent Plane	94

List of Tables

Table 3-1 Extensions to the O-RU.....	20
Table 3-2 Extensions to the O-DU	21
Table 3-3 Extensions to the RIC.....	22
Table 3-4 WEC in Sensing application	25
Table 3-5 WEC in MEC application	25
Table 3-6 Experiment Description	31
Table 3-7 Key Performance Indicators	31
Table 3-8 Experiment Description	35
Table 3-9 Key Performance Indicators	35
Table 3-10 Experiment Description	38
Table 3-11 Key Performance Indicators	38
Table 3-12 Experiment Description	42
Table 3-13 Key Performance Indicators – Integration of External Radar Sensing	42
Table 4-1 Measurement parameters for the CF-MIMO beamforming experiments	56
Table 4-2 Pairwise empirical SINR and estimated spectral efficiency for RX pair ch3/ch9.....	59
Table 4-3 Summary of key beamforming KPIs.....	61
Table 4-4 PoC#2 OAI integration of reciprocity calibration.....	74
Table 4-5 PoC#2, OTA reciprocity calibration in UC CF-MIMO platform.....	77
Table 4-6 PoC#2 -Experimental validation of multiuser in UC CF-MIMO platform.....	80
Table 5-1 Configuration @ OTE testbed.....	87
Table 5-2 COSMOTOOLS – Parameters / Data Received	88

Executive Summary

This deliverable presents the final demonstration activities of the 6G-SENSES project. This last phase of the project completes the integration activities, and extends the testing of the 6G-SENSES solution to realistic and dynamic scenarios in laboratory environments, improving system robustness and scalability, enhancing AI/ML-based optimisation mechanisms, and validating end-to-end (e2e) performance against the defined KPIs. These activities will be essential for demonstrating the practical viability, performance, and impact of the 6G-SENSES solution and for advancing the development of intelligent, adaptive, and energy-efficient 6G networks.

PoC#1 presents and validates an e2e multi-technology access and transport network architecture capable of jointly supporting communication and sensing services in Open RAN (O-RAN)-compliant Integrated Sensing and Communication (ISAC) deployments. Current 3GPP positioning functionality does not feature positioning/sensing capabilities good enough to compete with other technologies (LiDAR, Wi-Fi). 6G-SENSES has acknowledged that the most pragmatic solution in 6G is to integrate enabling technologies that allow for improved capabilities in the 3GPP/O-RAN ecosystem and make room for them in enhanced architectural frameworks. The identified technologies in 6G-SENSES are improved 3GPP-based solution involving technology add-ons and to leverage N3GPP technologies for improving the sensing resolution (both in range and angular). The latter systems need to be complementary and not adding up much energy consumption, therefore the seek for improving this parameter has been taken in mind when proposing a novel scalable MIMO front-end for ISAC with spectrum coexistence control.

In **PoC#2**, Cell-Free Multiple-Input Multiple-Output (CF-MIMO) prototype described in this section was deployed as a controlled over-the-air (OTA) measurement platform for evaluating distributed antenna deployments, reciprocity calibration algorithms, and cooperative beamforming strategies under realistic laboratory conditions. The small-scale CF-MIMO platform deployed at UC has implemented and evaluated an OTA multi-antenna and multi-carrier reciprocity calibration algorithms, which have been further integrated into OpenAirInterface (OAI) as an extension of its functionalities. The platform allowed assessing the gain of CF-MIMO over small-cell schemes in single-user and multi-user scenarios based of 5G-NR.

PoC#3 demonstrates an AI-driven Network Digital Twin (NDT) via RAN sensing, where sensing data, telemetry, and AI models are combined to enable closed-loop optimisation of the RAN. The Intelligent Plane is implemented on top of the ACC Open 5G SA NPN and dRAX Radio Intelligent Controller (RIC)/Service Management and Orchestration (SMO) stack, which provides telemetry collection, model execution, xApp/rApp hosting, and the control interfaces needed to actuate decisions back into the network. The PoC follows the model creation, model actualisation/validation, and model deployment lifecycle, ensuring that AI models are validated before being applied to the live network. The DT environment relies on ns-3/5G-LENA integrated with Sionna RT, allowing realistic ray-tracing-based simulations and scenario modelling that reproduce the physical deployment. The hybrid NDT is then used to generate accurate predictions and configuration recommendations that feed the Intelligent Plane and support network optimisation.

Overall, the integration results provide encouraging evidence that the 6G-SENSES vision—combining sensing, communications, and AI-driven intelligence within a unified architecture—is technically feasible. At the same time, they underline the gap that remains between conceptual designs and fully operational, scalable, and robust systems.

1 Introduction

The **6G-SENSES** project aims to advance beyond current 5G capabilities by introducing a unified and intelligent architecture that seamlessly integrates communication and sensing functionalities. Building upon the architectural foundations defined in Work Package 2 (**WP2**) and the technological developments of **WP3** and **WP4**, **WP5** focuses on the integration, validation, and demonstration of these innovations through a set of Proofs of Concept (PoCs).

Deliverable **D5.1** [1] established the design, specification, and planning of the PoCs, including their functional architecture, key components, and evaluation methodologies. In contrast, this deliverable (**D5.2**) focuses on the integration phase, reporting on how the different technological components are being interconnected, deployed, and prepared for experimental validation across the project testbeds. The integration activities in **WP5** are centred around three complementary PoCs, each addressing a distinct aspect of the **6G-SENSES** vision. These PoCs collectively demonstrate the feasibility of combining advanced Radio Access Network (RAN) technologies, sensing capabilities, and AI-driven intelligence within a unified framework.

The primary objective of the integration activities is to ensure that the individual components developed across the project are combined into cohesive and interoperable e2e systems. This involves not only the physical and logical interconnection of hardware and software elements, but also the alignment of interfaces, data models, and control mechanisms across multiple architectural layers, including the RAN, CN, and Intelligent Plane. A key aspect of this integration is the adoption of Open RAN (O-RAN)-based principles, enabling modularity, openness, and programmability using xApps/rApps and standardised interfaces. Additionally, the integration process incorporates AI/ML-driven functionalities, particularly within the Intelligent Plane, to support data-driven optimisation and decision-making. The integration is being carried out in an iterative and incremental manner, following the methodology defined in deliverable **D5.1** [1]. Initial efforts have focused on component-level validation and interface definition, followed by progressive system-level integration across different testbeds (**D5.2** [2]). This approach ensures that interoperability issues are identified at early stages, and are resolved efficiently, while enabling continuous refinement of the system.

Across the three PoCs described in the follow subsections, the project is transitioning from component-level development to system-level integration and validation. While **PoC#1** and **PoC#2** have reached a higher level of integration maturity, **PoC#3** represents a more dynamic and evolving system, driven by data-centric and AI-based processes.

1.1 Last Phase Experimentation Testing

This final stage covers the real-time implementation and operation of the experiments (see Figure 1-1). It includes:

- Executing preparatory activities, such as SUT configuration, validation, and environment setup.
- Conducting testing procedures based on defined scenarios.
- Capturing and storing measurements and telemetry from various components and layers.
- Performing post-processing, correlation, and analytics on the data collected.
- Evaluating results against defined KPIs.
- Documenting and sharing results for cross-site and cross-use case evaluation.

This structured methodology ensures that experimentation in **6G-SENSES** not only validates its cutting-edge innovations but also provides a coherent and repeatable path toward building trust in 6G technologies at scale.




Figure 1-1 6G-SENSES Testing and evaluation cycles

1.2 Organisation of the Document

This deliverable is structured as follows:

- Chapter 2 provides a cohesive vision of the 6G-SENSES solution.
- Chapter 3 presents the final deployment of PoC#1 and lists the final results and KPIs obtained.
- Chapter 4 presents the final deployment of PoC#2 and lists the final results and KPIs obtained.
- Chapter 5 presents the final deployment of PoC#3 and lists the final results and KPIs obtained.
- Chapter 6 provides the conclusions of the deliverable and provides ideas of the future work.
- Finally, at the end of the document the Consent Forms for participation in the Experiments are provided.

2 Overall 6G-SENSES testbed implementation

We presented in deliverable **D5.2** [2] the **6G-SENSES** multi-layer, multi-technology architecture, revised in deliverable **D2.3** [5] (Figure 2-1). It is aligned with O-RAN and 3GPP standards to deliver ISAC capabilities and to expose sensing information both within the system and to third-party applications.

The design follows the O-RAN model, which disaggregates traditional RAN functions and incorporates software-defined intelligence and advanced control mechanisms. A disaggregated gNodeB (gNB) comprises the Radio Unit (O-RU), O-DU, and Central Unit (O-CU), interconnected through standardised fronthaul (FH) and F1 interfaces. The figure incorporates the mapping of the different project PoCs to the architecture.

The **6G-SENSES** architecture is organised as a multi-layer O-RAN/3GPP-aligned architecture spanning the RAN, Core Network (CN), transport layer, and Intelligent Plane. The three PoCs are mapped onto different layers and capabilities (as shown in Figure 2-1):

- **PoC#1** – Multi-Technology ISAC Platform: focuses mainly on the RAN, transport, and O-RAN intelligence layers by integrating communication and sensing sources (3GPP and non-3GPP), exposing sensing information to the RIC, and enabling sensing-assisted communication and transport optimisation.
- **PoC#2** – CF-MIMO Prototype: concentrates on advanced RAN technologies, validating distributed antenna systems, reciprocity calibration, beamforming, and CF-MIMO operation within realistic OAI-based deployments.
- **PoC#3** – NDT via RAN Sensing: primarily targets the Intelligent Plane, integrating sensing, AI/ML, DT, and closed-loop optimisation through the dRAN RIC/SMO platform and NDT environment.




Figure 2-1 Multi-layer 6G-SENSES Architecture and mapping of PoCs

The architecture intentionally did not integrate every component into a single e2e demonstrator because of technology maturity differences, integration complexity, and the need to validate the different **6G-SENSES** innovations at the appropriate architectural layer. Moreover, the complexity of the testbeds, which have

been mostly built within the project duration, makes extremely difficult bringing a single platform that serves all PoCs. PoC#1 and PoC#2 reached a higher level of integration maturity, whereas PoC#3 remained a more dynamic, data-centric and AI-driven system, making synchronised e2e demonstrations difficult.

BubbleRAN (BR) served as the common integration platform for PoC#1, while the ACC Open 5G/dRAX platform formed the common baseline for PoC#3. Two CF-MIMO testbeds from TUBS and UC served the purpose of evaluating different capabilities of such type of networks.

All testbeds built up in the project will be featured with additional capabilities in the context of both national and EU-funded projects.

3 PoC#1 – Multi-Technology ISAC Platform

PoC#1 builds an e2e demonstration showing the potential of a multi-WAT ISAC platform, with the possibility to leverage sensing data stemming from different WATs acting as “sensors” (non-3GPP non-Wi-Fi and Wi-Fi) complementing O-RAN Radio Units (O-RUs), Distributed Unit (DU), and Central Unit (CU)/User Plane Function (UPF) pools with a real-time control fabric providing sub-millisecond control loop over each 3GPP network component. The architecture of the **6G-SENSES** RAN relies on a multi-layer structure, inspired by the 3GPP and O-RAN standards, which has been recently revised in deliverable **D2.3** [5]. Extensions are required to support ISAC services both at the CN and RAN – Near-Real-Time (Near-RT) and Non-Real-Time (Non-RT) RAN Intelligent Controllers (RICs) – to dynamically adjust network behaviour.

3.1 Multi-technology ISAC RAN implementation

The objective of this PoC is to demonstrate the practical integration of heterogeneous sensing and communication technologies within a O-RAN RAN. Rather than developing isolated sensing demonstrators, the implementation focuses on incorporating sensing capabilities into the operational RAN architecture so that sensing information can be collected, transported, and exploited by standard O-RAN control functions to support communication services.

To this end, the implementation combines two complementary sensing sources: a Sub-6 GHz software-defined radio (SDR)-based gNB providing monostatic sensing capabilities and a 60 GHz mmWave sensing platform. Both sensing systems are integrated through a common O-RAN framework using dedicated E2 Service Models (SMs), allowing sensing information generated by heterogeneous technologies to be consumed by xApps running on the Near-RT RIC.

A key implementation aspect is the realization of the sensing-enabled gNB using an SDR platform based on OAI. Instead of introducing dedicated radar transmissions, the solution exploits the flexibility of the 5G NR TDD frame by reusing the flexible symbols available within the standard frame structure to transmit sensing waveforms. This approach preserves compatibility with the communication stack while enabling monostatic sensing using the same radio hardware, illustrating a practical and standards-aligned path towards Integrated Sensing and Communication (ISAC).

Beyond integrating sensing sources, the PoC validates the complete sensing information chain, from signal acquisition at the RU to its exploitation by O-RAN intelligence applications. This includes the implementation of sensing-aware E2 SMs, the extension of O-RAN components to support sensing information exchange, and the validation of sensing-assisted communication functions such as proactive MAC scheduling and mobility support. The resulting platform demonstrates how heterogeneous sensing technologies can be transformed into interoperable O-RAN services capable of supporting closed-loop network optimisation.

3.1.1 Final deployment and configuration

The overall deployment is depicted in Figure 3-1 and comprises the integration of the elements and functionalities described in deliverable **D5.2**, Section 3.1.1. It includes sensing capable devices, as well as a set of services to allow the integration of sensing information to support sensing assisted communications. The description herewith focuses on the integration and validation of the different functional elements described in **D5.2**, and references will be done to the corresponding sections of that deliverable when necessary.




Figure 3-1 PoC#1 implementation at UC

As can be observed, the setup comprises two commercial O-RUs (LiteOn), one SDR based gNB with radar-based sensing capabilities (highlighted in red), a mmWave radar and an additional SDR based gNB. These elements are integrated into the BubbleRAN MX-PDK platform leveraging the O-RAN SMs described in D5.2.

In the following the description of the functionalities is articulated into four different system level functionalities, that combine part of the elements described in section 3.1.1 of deliverable D5.2. The first two ones refer to integration of sensing elements into a wider O-RAN platform, while the last two ones describe the exploitation of the sensing information to support sensing-assisted communication solutions.

3.1.1.1 Sub-6 GHz gNB with radar capability integrated with O-RAN

This section describes the final integration of a Sub-6 SDR-based gNB with the capability to perform radar-based sensing. The gNB is based on OAI version v2.4.0, whose operation in the interaction with the SDR elements has been extended. As described in section 3.1.1.3.3 of D5.2, the radar functionality exploits the flexibility of the TDD NR frame structure to send sensing signals in flexible symbols, which are not used for either downlink or uplink. In particular the integrated solution is configured with numerology 1 (SCS of 30 kHz) and TDD configuration #16, which defines one flexible slot per sub-frame, where the first symbol is used for downlink and the rest 13 symbols are idle. Altogether, the configuration provides a time period of 464.28 μ s every 5 ms for sensing. The integration approach provides flexibility to integrate custom sensing signals within the available idle periods, provided they use the same bandwidth and sampling rate used for the communication.

The integrated system uses USRP N321 SDR devices which are able to operate with bandwidths of 200MHz. However, in order to be compliant with the 3GPP specifications, TS38.101, the gNB has been configured with a bandwidth of 100 MHz, yielding a sampling rate of 122.88 Msps, permitting sending a maximum sensing burst of around 57000 samples.




Figure 3-2 Sensing sources (sub-6 SDR based gNB and non-3GPP mmWave radar) integration within O-RAN instance

Figure 3-2 shows the interaction between the involved elements. As can be seen, the final setup comprises five N321 devices. The top device is used by the OAI legacy operation for communication (transmission and reception), and it is silently used by the sensing extension to transmit the sensing signal during the available sensing periods. The bottom SDR devices are used only by the sensing functionality to capture the reflected signals via the Uniform Linear Array (ULA) of 8 antennas. As commented in deliverable [D5.2](#), all the antennas could have been used for sensing and communication using the same scheme. However, it would have required high computation capacity to process the communication flows. Opposed to that, the adopted approach alleviates the computation load while keeping the same functionality in terms of sensing integration.

Next, the reflected signals are gathered by the E2 agent using Inter Process Communication (IPC) and sent to the near-RT RIC, as depicted in Figure 3-2 (right). The processing of sensing information can be done locally at the gNB or fully offloaded to the xApp, and the adopted O-RAN SM needs be aligned with the sent data, as defined in section 3.1.1.1 of [D5.2](#) and section 2.2.1 of [D4.3](#). The final implementation splits the processing between the gNB and xApp to simplify the integration with existing E2 SMs. In this sense, the eight reflected signals (one from each antenna of the ULA) are correlated with the transmitted sensing signal and only the result of the correlation is sent through the [E2](#) interface. This implementation decision was taken due to the limited capacity of the buffers of the current implementation of the near-RT RIC, which prevented from sending the raw data. Eventually, the sensing data is gathered by an xApp that finishes the processing and generated the information plotting.

Lessons learnt:

- The 3GPP standards limits the bandwidth of a sub-6 GHz carrier to 100MHz, which limits the sensing capacity at FR1. To overcome this limitation Carrier Aggregation (CA) can be consider to increase the overall bandwidth.
- Experimental setups when using multiple antennas, as it is the case, imposes demanding requirements in terms of both computation and specially communication. Modularity and separation of concepts allows keeping functionality while alleviating the required capabilities.

- Internal implementation of current services, such as the near-RT RIC, may not be readily usable to transport sensing data due to its high volume and potentially bursty nature. Both services and network management techniques need to be revisited to ensure sensing support.

3.1.1.2 mmWave analogue frontend integrated within O-RAN

The second sensing source integrated is related to a mmWave ISAC system operating at 60 GHz. Most of the integration details was already presented in section 3.1.1.2.2 of [D5.2](#). Overall the integration follows the same approach as the one detailed in the previous section, as represented in Figure 3-2. Differently to the previous section, in this case all the sensing processing is locally performed, and only the heatmaps are delivered to the near-RT RIC using the E2 SM detailed in section 3.1.1.1 of [D5.2](#).

3.1.1.3 Sensing supported MAC scheduling

The first sensing assisted communication functionality refers to the MAC scheduler whose overall design was introduced in section 3.1.1.3.4 of [D5.2](#). It has been implemented as an extension of OAI (version 2.4) and equipped with a customized E2 SM that extends the standard O-RAN one. The integrated SM allows both gathering MAC layer information, as well as controlling the operation of the scheduler overtime. An xApp has been implemented able to simultaneously hand the MAC SM and the sensing exploited by the mmWave radar. This way, the xApp is able to detect the motion of the UE and proactively adapt the scheduler configuration. As will be shown in the following section, the xApp detects anticipately that the UE will be hidden and reconfigures the scheduler to avoid abrupt QoS degradation.

3.1.1.4 Handover-based sensing integration

This demonstration extends the handover-based sensing integration reported in [6G-SENSES D3.2](#) (Sec. 2.1.4.2, “Hardware abstractions – BubbleRAN MX-RIC Solution”) and [D5.2](#) (Sec. 3.1.1.3.2, “Handover-based sensing integration”), and reports the corresponding measurements and KPIs obtained at the [UC](#) lab. Note that this [UC](#)-lab test is the main integration result for [D5.3](#), and is distinct from the earlier Globecom 2025 demo video referenced in deliverable [D3.2](#) (Scenario 1), which used a 2×USRP B210 deployment at [BR](#)’s own testbed as an initial proof of concept for the same sensing-triggered handover workflow.

Setup: the scenario uses the split-RAN deployment described in Section 3.1.1, with one CU and two gNB-DUs (DU1, DU2). Unlike the Globecom 2025 proof of concept (2×USRP B210, [D3.2](#) Scenario 1), the [UC](#) lab test reported here uses the COTS O-RAN Radio Units already deployed at the [UC](#) premises (Figure 3-1): one gNB-DU is connected to a LiteON FlexFi Sub-6 O-RU (5G SA, O-RAN option 7.2, n78 band) and the other to a Benetel O-RU. A commercial Google Pixel 7 acts as the UE. The BubbleRAN Near-RT RIC (MX-RIC) hosts two monitoring xApps – one per DU, subscribed via the O-RAN Low Link Control (LLC) Service Model – and one handover-control xApp (O-RAN RC Service Model, F1AP-based).

Workflow: the UE first attaches to DU1. The monitoring xApp continuously extracts Sounding Reference Signal (SRS) I/Q samples, noise samples and channel-frequency-domain estimates from DU1 and streams them to a real-time processing pipeline that computes the SNR and the Channel Impulse Response (CIR) peak, with an end-to-end measurement latency of approximately 1 ms. When the observed SNR fluctuation exceeds a configurable threshold, the handover-control xApp triggers an F1AP-based handover, moving the UE from DU1 to DU2. The second monitoring xApp then confirms the resulting channel quality on DU2.

Measured KPIs:

- End-to-end measurement latency (I/Q capture → xApp): ≈ 1 ms.
- Pre-handover channel quality (DU1): SNR fluctuation $\approx \pm 20$ dB (unstable).
- Post-handover channel quality (DU2): SNR fluctuation $\approx \pm 5$ dB (stable).
- Handover outcome: completed via F1AP with no service interruption.

3.1.1.5 SRS demo

This section documents the SRS-based low-layer sensing capability of the BubbleRAN MX-RIC/MX-PDK platform, introduced in **D3.2** (Sec. 2.1.4.2, Scenario 2 – “I/Q Collection with O-RAN 7.2 Deployment Under UE Mobility”) and demonstrated more broadly in BubbleRAN's public MX-PDK ISAC technology demonstration (bubbleran.com/news/isac-with-mx-pdk).

Setup: same O-RAN 7.2 split deployment as the handover demo (single gNB-DU, LiteON O-RU). A monitoring xApp subscribes to the DU via the LLC Service Model and continuously extracts, per SRS occasion, the I/Q samples, noise samples and channel-frequency-domain estimate; a Python-based pipeline post-processes these into SNR and CIR time series.

Mobility characterisation: the UE was physically relocated to different positions within the lab at 50-second intervals while the RIC continuously logged the low-layer SRS measurements, in order to characterize how channel quality (SNR/CIR) varies with UE position and mobility.

Measured KPIs:

- SRS-derived measurement latency: ≈ 1 ms e2e.
- Measurement granularity: per-SRS-occasion I/Q, noise floor and channel-frequency-domain estimate, from which SNR and CIR peak are derived.
- Distinguishable SNR/CIR signatures observed across UE relocation steps (50 s dwell per position).

Broader capability: the same SRS I/Q $\rightarrow$ feature-extraction pipeline underpins **BR's** published ISAC applications — indoor positioning, human/object detection and multi-UE density/occupancy mapping (see D4.4 and the MX-PDK public demonstration) — and provides the architecture needed to trigger the handover/decision xApp above from a mmWave-radar event instead of (or in addition to) the SRS-SNR threshold, as discussed for the UC lab integration.

3.1.2 Results and KPIs

This section presents validation results of the aforementioned integrated functionalities. Furthermore, the tables of the sections summarize the capabilities and functionalities tested.



Figure 3-3 Execution example of the Sub-6 gNB with sensing capabilities

Figure 3-3 depicts an execution of the sub-6G gNB equipped with mono-static radar functionality. The plot shows (left) the messages generated when started the gNB. As can be observed, information related to

sensing configuration appears among the legacy communication-related messages. At right, a capture of the execution is shown where we indicate the number of samples, whether they correspond to communication (HW) or sensing (SENS) and the time instant (in microseconds). As can be observed, there is a transmission pattern corresponding with the sub-frame duration of 5 ms, corresponding to numerology 1. First, there are seven sendings of 61440 (0.5 ms for a sampling rate of 122.88 Msps) which correspond to the downlink slots shown in the diagram. Afterwards, a short burst of 4448 samples is delivered, which corresponds to the first symbol of the flexible slot. Then, the sensing samples are delivered to the SDR device.

Table 3-1 Extensions to the O-RU

Test 1.X	Extensions to the O-RU
Phase	Validation and performance evaluation
Description	Provision of monostatic radar functionalities to a gNB based on OAI implementation using an SDR platform (N321). Use of digital beamforming as in section 3.2.2.2 (D3.1) to perform AoA estimation. The sensing system features a 100 MHz BW for sensing
KPIs	Measurable requirements: - Angular resolution: 20° with 8 antennas. - Range resolution: 1,5 m Non-measurable requirements: - Added E2 functionality for direct gNB – RIC interfacing to transport sensing information (see Table 3-3)
PoC	PoC#1, UC testbed




Figure 3-4 Execution validation of the sensing supported MAC scheduler

The integration of the sensing-assisted MAC scheduling is validated with a small experiment, depicted in Figure 3-4. In this experiment a UE moves behind an obstacle (the column in the plot), while its motion is captured with the mmWave radar. The sensing information is consumed by an xApp that jointly manages the MAC scheduler. In particular, we have used the scheduler design described in section 4.1 of deliverable D4.3. In this case, the scheduler operation is not adapted using AI/ML but a simple rule. The scheduler is designed to jointly consider the occupancy of the RLC queues, the Guaranteed Flow Bit Rate (GFBT) of the services and the amount of PRBs used. In this regard, the scheduler monitors the GFBT deficit over time, implemented as a virtual queue, and aims to stabilize all queues. In top side on Figure 3-4 the temporal evolution of the scheduler is shown, with (right) and without (left) sensing information. In the configuration without sensing information, the same relative importance is given to all considered parameters. On the other hand, when using sensing information, the usage of PRBs is not considered in the scheduler when the radar detects that the UE is hidden behind the obstacle. In the top plots, period of time when the UE hides behind the column is clearly indicated by a sudden upturn of the PRBs used. When no sensing information is exploited, the MAC scheduler reacts to the queue variation, leading to an initial increase of the GFBT deficit queue just when the UE moves behind the column. Opposed to that, the proactive reaction when using sensing information avoids such variability, thus improving the service QoS.

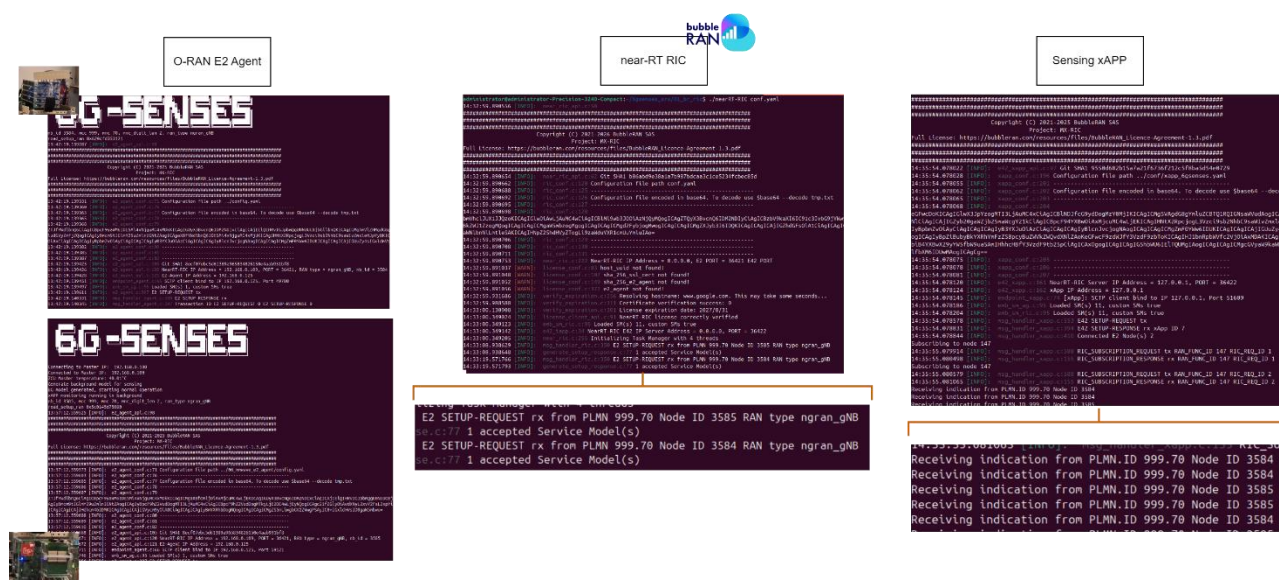


Figure 3-5 Execution validation of the sensing SM and end-to-end (E2 agent-xAPP) communication

As an example of the integration of the MAC and sensing SMs, Figure 3-5 shows the execution console of two E2 agents (left), one corresponding to each sensing source, the near-RT RIC (center), and a sensing xApp (right) able to gather information from both E2 agents. As can be observed, the RIC processes the E2 SETUP requests from both nodes, which belong to the same network (same PLMN ID) but with different node IDs. It is worth noting, although one node is a 3GPP gNB and the other a non-3GPP device, the E2 agents of both devices are configured in such a way that both are seen as RAN nodes by the RIC. Finally, it can be seen that the xAPP is able to receive E2 INDICATION messages from both devices, once it subscribes to such information.

Table 3-2 Extensions to the O-DU

Test 1.X	Extensions to the O-DU
Phase	Validation and performance evaluation

Description	he O-DU is provided with techniques to better manage the radio resources based on sensing information. Particularly, it incorporates a multi-objective MAC scheduler able to manage traffic flows with different characteristics and requirements, whose operation can be tweaked from a sensing xApp through the near RT-RIC using a dedicated E2 Service Model (E2SM).
KPIs	Measurable requirements - User performance improvement with sensing information: avoidance of sudden GFBR degradation upon communication conditions variability Non-measurable requirements: - E2 support at the DU for xApp assisted MAC scheduling configuration
PoC	PoC#1, UC testbed

Table 3-3 Extensions to the RIC

Test 1.X	Extensions to the RIC
Phase	Validation and performance evaluation
Description	Surrounding sensing SM: this SM provides a heatmap of the surroundings provided by a modified O-RU with radar-based sensing capabilities. It is assumed that the I/Q processing has been already performed. As a result, the sensing information consists of range/angle coordinates around the sensing device. MAC scheduler SM: this SM provides MAC scheduler configuration information along with metrics used for MAC scheduling, such as buffer occupancy or level of compliance with 5QI parameters. Besides, the SM allows modifying configurable parameters of the scheduler (i.e. weights applied to different parameters).
KPIs	Non-measurable requirements: - Capability to ingest sensing data from multiple sources in O-RAN RICs. - Definition of E2SMs for general sensing information collection.
PoC	PoC#1, UC testbed

3.1.2.1 RIS-assisted sensing

To assess the feasibility of RIS-assisted sensing, we have carried out tests with the Sub-6 sRU to perform sensing behind obstacles (shared with the BeGREEN project). The target setup was such that an sRU is configured to perform sensing over the RIS to detect an object (corner reflector) that stands behind an obstacle (wall with absorbers). To that end, we transmit a signal through the RIS with the intention of receiving a reflection from the object, back to the sRU through the RIS. The system was integrated and tested initially in a controlled environment, such an anechoic chamber. As the developed system works in the 5 GHz Industrial Scientific and Medical (ISM) band, strong multipath is expected in an indoor environment. Having a strong multipath would make hard drawing any meaningful conclusions about the behaviour of the system.

The setup for testing the integrated RIS/ISAC system in the anechoic chamber is shown in Figure 3-6. The Sub-6 sRU marked in the figure, is the developed ISAC system using SDRs. The RIS is placed in front of the ISAC sensing system at a distance of 3.4 meters. Additionally, a corner reflector is placed a bit off side, hidden behind a wall made of absorbing material and completely undetectable by the sensing system.




Figure 3-6 Anechoic chamber setup for testing the RIS-assisted sensing

Few different issues were discovered when testing. Given that the signal needs to travel two times over the RIS, it was concluded that the transmit power of the radio was not sufficient. A power amplifier and a phased array antenna were used to achieve high equivalent isotropic radiated power (EIRP). The use of additional PAs did increase the detection range and the returned signal strength (see Figure 3-8).




Figure 3-7 Heat maps for a) 80, b) 85, c) 90, d) 95 degrees reflection angles with respect to the RIS surface




Figure 3-8 Stronger reflection received back with use of PA and high gain antenna

Given the difficulties in making such a system working in a controlled environment, it remains unclear given the experiments with the existing hardware is how the system would work in an indoor environment with strong multipath.

3.1.2.2 WEC demo

In this small-scale demo we implement novel caching techniques in the context of Wireless Edge Caching (WEC). The work carried out for the purpose of this demo has been documented through the 6G-SENSES project across multiple deliverables.

The whole process of the evolution of the caching strategy and its implementation in the project has been introduced in deliverable D5.1 [1] and its work carried out and documented mainly through two WP4 deliverables.

The work was first introduced in deliverable D4.1 [8], where the first iteration of the WEC strategy was proposed in a sensing use case, and evaluated in a simplified ad-hoc sensing environment. This is summarized in Table 3-4.

The work has expanded then in the later stages of the project, and the proposed strategy has evolved to a more coherent approach with regard to the whole 6G-SENSES project, and a more standardized environment


for what concerns the simulation aspect of the studied use cases. This final stage has been documented in Deliverable [D4.2](#) [9] and summarized in section 3.4 of deliverable [D2.3](#) [5]. This is reported in Table 3-4.

Table 3-4 WEC in Sensing application

Test 1.1	Sensing Application
Phase	Methodology and Implementation
Description	In this initial phase, we implement the first version of a WEC strategy within a wireless sensing scenario. This first sensing application serves to define and validate the initial WEC framework while also establishing the ROS simulation environment that will later support the small-scale demonstration. The sensing infrastructure simulated in this first iteration consists of a set of sensors distributed throughout the area of interest. Each sensor continuously monitors its surroundings and periodically broadcasts the collected information within a predefined communication range. Alongside the sensing infrastructure, multiple edge devices are deployed to interact with the sensors and process incoming sensing requests issued by a central base station. Sensing is carried out continuously, whereas data requests are processed at fixed discrete time intervals. At every interval, a RAN entity generates a sensing query associated with a specific target within the monitored area. The goal is to obtain the latest valid sensing information produced by the sensor located closest to the requested target. To improve response time and reduce communication costs, sensing data may be cached at edge devices. A query from the base station can be satisfied by any edge device that already stores the required information, as long as the cached data remain valid. Each cached item is associated with a freshness metric that determines its validity over time; once this metric drops below a predefined threshold, the data are considered stale and are no longer eligible to satisfy sensing requests.
KPIs	• Energy Consumption: we measure the amount of energy consumed in terms of sensor activations, i.e., how many times sensors are activated by a given caching policy. • End-to-end Delay: we measure the average time it takes to satisfy requests, given standardized times on tasks such as data retrieval and cache interrogation.
PoC	PoC#1, small-scale demo, ____ testbed

Table 3-5 WEC in MEC application

Test 1.2	MEC Application
Phase	Validation and performance evaluation
Description	In this iteration, we aim to enhance the performance experienced by end-user devices operating on a field, by leveraging edge computing caching mechanisms. With this objective, we not only aim to reduce latency through edge computing, but to improve the cache eviction strategies at the MEC layer under dynamic wireless and mobility conditions. While conventional MEC solutions rely mostly on content popularity, they often overlook the real-time state of wireless channels and the mobility of user devices. We demonstrate how location-aware caching, enabled by sensing information, can be exploited to adapt cache contents to the current communication context. Specifically, the proposed approach leverages channel-quality indicators and proximity information to prioritize the storage of data that is more likely to be accessed by nearby users, thereby reducing cache MISS events.

	
	Figure 3-9 WEC Demo, deploying a MEC infrastructure Figure 3-9 shows an example instance of Gazebo running the use case. In this particular instance, entities are spawned close together for a better visualization. The figure shows edge devices, antenna relays, and users taking part in the simulation.
KPIs	The main performance metric considered in this use case is the cache MISS RATE, defined as the proportion of requests that cannot be served using data already stored at the edge. The evaluation assesses the proposed CIFO caching strategy under different system configurations and operating conditions by varying several key parameters. A lower cache miss rate indicates a more effective caching policy, resulting in improved edge resource utilization and better overall application performance.
PoC	PoC#1, small-scale demo, ____ testbed

3.2 Multi-technology transport optimisation

The PoC#1 implementation at IASA demonstrates two complementary capabilities of the 6G-SENSES architecture. First, it validates a multi-technology transport network able to jointly support and dynamically orchestrate communication and sensing traffic across O-RAN infrastructures. Second, it demonstrates the integration of sensing information generated by non-3GPP technologies into the O-RAN ecosystem, enabling its delivery to Near-RT RIC xApps through the Service Management Orchestrator (SMO) and O-RAN orchestration framework.

3.2.1 Deployment and configuration

The objective of this PoC is to experimentally validate a programmable multi-technology transport infrastructure capable of jointly supporting communication and sensing services in O-RAN compliant ISAC deployments. The demonstrator validates how heterogeneous communication and sensing traffic generated by multiple O-RUs can be efficiently transported, orchestrated, and dynamically managed over a hybrid packet-optical transport network while satisfying the bandwidth, latency, and service differentiation requirements of both traffic classes.

The main features validated under this PoC include:

- E2e support of ISAC services over O-RAN compliant infrastructures.
- Joint transport of communication and sensing traffic over shared transport resources.
- Integration of packet switching and all-optical switching under a unified Software Define Networking (SDN) control framework.
- Dynamic allocation of transport resources according to service requirements.

- Logical isolation and differentiated treatment of communication and sensing traffic.
- Programmable transport control through SDN.




Figure 3-10 Experimental testbed used in final demonstration for PoC#1 at IASA

Figure 3-10 illustrates the final deployment of the **IASA** transport testbed.

The **IASA** demonstrator implements a multi-layer, multi-technology architecture compatible with O-RAN and 3GPP specifications. The transport infrastructure interconnects the disaggregated RAN components with the core network while simultaneously supporting communication and sensing traffic.

Each O-RU generates baseband I/Q samples transported using enhanced Common Public Radio Interface (eCPRI) over Ethernet. The generated traffic is encapsulated into Virtual Local Area Network (VLAN)-tagged Ethernet flows, allowing logical separation according to radio unit, sensing stream, antenna, beam, or service type. At the O-DU, the incoming VLAN-tagged flows are demultiplexed and forwarded to the appropriate communication or sensing processing pipelines according to the O-RAN scheduling policies.

The demonstrator consists of:

- USRP N310 O-RUs equipped with 10G SFP+ interfaces;
- OpenFlow-enabled Aruba 3810M packet switches;
- Polatis 6000 all-optical switching fabric;
- OpenDaylight SDN controller;
- Centralized transport orchestrator.

The packet network is controlled through OpenFlow and RESTCONF interfaces, while the optical switching fabric is configured using Standard Commands for Programmable Instruments (SCPI)/Telnet commands. The transport orchestrator maintains a global view of packet forwarding paths, optical cross-connections, VLAN mappings, wavelength assignments and resource utilization, enabling coordinated management across the heterogeneous transport domains.

3.2.2 Experimental Validation

3.2.2.1 Experiment 1 – Joint Communication and Sensing Transport

The objective of this experiment is to validate the proposed sensing-aware transport optimization framework over a multi-technology packet-optical transport network supporting ISAC services in O-RAN compliant networks. The experiment demonstrates the capability of the transport infrastructure to simultaneously transport communication and sensing traffic, while dynamically allocating network resources according to communication requirements, sensing demands, and optical transport constraints.

The optimization framework jointly considers the radio access network, the optical transport network, and the sensing application to determine the most efficient allocation of transport resources while maintaining communication service continuity and sensing performance. In particular, the framework optimizes the association between O-RUs and O-DUs, routing and optical resource assignment of communication and sensing flows, and the activation of sensing streams according to their sensing quality. Only sensing streams satisfying the required signal quality threshold are transported towards the sensing application, thereby reducing unnecessary transport overhead and improving network resource utilization.

3.2.2.1.1 Experiment

The overall layout of the network architecture used in the experiments is shown in Figure 3-11. This exploits an optoelectronic transport network interconnecting the RAN with the core functions located at the Multi access Edge Compute (MEC) servers. The hierarchical structure of the proposed architecture offers RAN connectivity and collects and aggregates transport network and sensing traffic streams from the RUs, while it transports these to the servers. Both the RAN and aggregation transport network segments are implemented through multi-vendor SDN-controlled optoelectronic switches with different capabilities (port no, capacity per port & latency). The access network is equipped with low energy consuming switching nodes with limited number of input ports and relatively small capacity (1/10GbE/port).




Figure 3-11 Transport network supporting ISAC. Configuration with emulated air interfaces. Snapshot of 3 moving targets, blue dots are static reflectors

The aggregation transport network comprises high-end optoelectronic and all-optical fast switches with higher capacity and density (10G/40G/100GbE/ports), offering much higher energy efficiency. All-optical fast switches are a key enabler that facilitates transparent routing of IQ sensing flows to the MEC servers needed to achieve acceptable transport delays. For the RAN segment, a Software Defined Radio (SDR) platform (based on USRP B210 and N310) transmitting 5G NR Orthogonal Frequency-Division Multiplexing (OFDM) modulated signals is used to support the functionality of the RUs. For the compute domain, edge servers attached to the access switches and central cloud servers connected to the aggregation/core switches are adopted. These servers are used to host containerized RAN elements (i.e. DU/CU) and sensing functions (storage of I/Q echoes and processing).

The sensing app implements a passive OFDM-based radar that uses as input, signals reflected on objects in the surrounding area, when connectivity between BSs and UEs is established. The sensing app evaluates the quality of the reflected signals captured per BS and, if these are above a specific SNR threshold, they are used to estimate the position of the corresponding object. These results are then communicated to the end-to-end orchestrator that calculates the optimal resource allocation policy needed to accommodate both communication and sensing services. This is achieved by identifying the appropriate capacity for the paths implementing “RU-DU-CU” and “RU-Sensing” function connectivity, executed by the SDN transport network controller. Finally, for the wireless channel emulation, a dedicated virtual machine has been deployed hosting the open-source GNU Radio platform, emulating channel impairments and objects in the surrounding area using the “Noise Source” and “Static Target Simulator” blocks, respectively.

Transport Optimisation Methodology

The experiment validates the proposed sensing-aware transport optimization framework for networked ISAC services. The optimization periodically collects communication traffic demands, sensing traffic requirements, optical network resource availability, and sensing quality measurements reported by the individual O-RUs.

Based on this information, the optimization jointly determines:

- the association between O-RUs and O-DUs;
- the routing and wavelength assignment of communication and sensing traffic;
- the activation of sensing streams according to their received SNR;
- the allocation of optical transport capacity.

Details of the optimisation framework are provided in [19].

The optimization operates under radio access, transport capacity, wavelength continuity, propagation delay, and sensing quality constraints. Sensing streams with received Signal to Noise Ratio (SNR) below a predefined threshold are discarded, preventing the transport of sensing information that cannot contribute to reliable target detection while simultaneously reducing unnecessary network resource consumption. The optimization objective is to minimize the total transport resources consumed by both communication and sensing traffic while preserving communication connectivity and sensing functionality.

Experimental Procedure

The experiment is executed through the following sequence:

1. Communication services are established between the O-RUs and the virtualized O-DU through the packet-optical transport network.
2. The passive OFDM sensing application is activated using the received communication waveforms.
3. The transport orchestrator continuously collects communication traffic demands, sensing traffic requirements, and sensing quality measurements.

4. The sensing-aware optimization framework computes the optimal allocation of transport resources and selects the sensing streams satisfying the required SNR threshold.
5. The SDN controller provisions the corresponding packet forwarding rules and optical network connectivity.
6. Communication performance, sensing performance, and transport resource utilization are continuously monitored throughout the experiment.

3.2.2.1.2 Experimental Results and KPIs

The experimental evaluation validates the effectiveness of the proposed sensing-aware transport optimization framework. The first set of experiments investigates the impact of sensing resolution on transport resource utilization. Figure 3-12 (a) shows that increasing the sensing resolution by reducing the velocity resolution from 8.4 m/s to 1.7 m/s increases the transported sensing traffic from approximately 0.6 GB to 3.0 GB, corresponding to an almost five-fold increase in transport demand. The results demonstrate an approximately linear relationship between sensing resolution and transport load, highlighting the need for dynamic transport resource allocation capable of adapting to varying sensing requirements.




Figure 3-12 a) Volume of sensing traffic vs resolution velocity b) Average Optical transport network capacity allocated for sensing with and without optical network switching

The second set of experimental results evaluate the effectiveness of optical switching aggregation in reducing transport resource utilization. As illustrated in Figure 3-12 (b), Optical Circuit Switching (OCS) significantly reduces the required optical network capacity for scenarios involving a limited number of moving targets. For a single target, the required transport capacity decreases from approximately 25 Gbps to 5 Gbps, corresponding to an 80% reduction in required transport resources. As the number of simultaneously observed targets increases, the required sensing traffic also increases, reducing the benefit of dynamic optical aggregation. For five simultaneous targets, continuous connectivity between the radio units and the sensing application becomes necessary, resulting in similar transport capacity requirements with and without optical switching aggregation.

The sensing performance is evaluated through both monostatic and multistatic sensing configurations. The experimental results demonstrate that individual monostatic radio units detect only subsets of the available targets due to their individual observation geometry. In the evaluated scenario, one radio unit successfully detects two targets while another detects three targets. By jointly processing sensing information collected from multiple radio units, the multistatic sensing configuration successfully detects all targets within the monitored area, demonstrating the benefit of cooperative sensing and sensing diversity.

Finally, the detection probability is evaluated for different sensing scenarios. A snapshot of the sensing scenario is shown in Figure 3-13. In this set-up we consider the case of an OFDM-based monostatic radar with all RUs forwarding I/Q samples to the “Low-Phy” sensing function hosted in server 1. Each set of I/Q samples

is processed independently by the sensing function. From Figure 3-13(a) we observe that when applying the monostatic radar algorithm the received sensing signals captured from RU4 can be used to detect only 2 targets. In contrast, for this layout sensing information from RU5 can detect only 3 targets. Therefore, utilizing the transport network to combine sensing from multiple RUs all moving targets can be detected.




Figure 3-13 Moving targets as detected by (a) RU4 and (b) RU5, Nframe = 12, W= 100MHz, RAN frequency =3.5GHz

The results show that the detection probability remains close to 100% for a single target and gradually decreases to approximately 81–82% when simultaneously tracking four moving targets. Increasing the number of cooperating radio units significantly improves sensing performance, with the detection probability increasing from approximately 86% for groups of two radio units to nearly 98% when five radio units jointly participate in the sensing process. These results demonstrate that cooperative sensing substantially improves target detection while maintaining efficient utilization of transport resources

A summary of the experimental setup along with the main results is provided in Table 3-6 and Table 3-7, respectively.

Table 3-6 Experiment Description

Item	Description
Test ID	Test 3.1
Test Name	Joint Communication and Sensing Transport
Phase	Validation and Performance Evaluation
Objective	Validate the sensing-aware transport optimization framework supporting simultaneous communication and sensing services over a multi-technology packet-optical transport network
Experimental Setup	USRP B210/N310 SDRs, Aruba OpenFlow switches, Polatis OCS, OpenDaylight SDN Controller, GNU Radio channel emulator, MEC servers hosting O-DU/O-CU and sensing applications
Traffic	5G NR communication traffic and OFDM sensing I/Q streams
Expected Outcome	Efficient coexistence of communication and sensing services while minimizing optical transport resource utilization
Related PoC	PoC#1 – Multi-Technology Transport Network

Table 3-7 Key Performance Indicators

KPI	Measured Value	Unit
Sensing traffic volume	0.6–3.0	GB
Velocity resolution	8.4–1.7	m/s
Maximum sensing traffic increase	≈5x	—

Optical network capacity with OCS	5–98	Gbps
Optical network capacity without OCS	25–100	Gbps
Maximum optical capacity reduction	≈80%	%
Targets detected (individual RU)	2–3	Targets
Target coverage (multistatic sensing)	100%	%
Detection probability (1 target)	≈100%	%
Detection probability (4 targets)	≈81–82%	%
Detection probability (RU group size = 5)	≈98%	%
Detection probability (RU group size = 2)	≈86%	%

3.2.2.2 Experiment 2 – Multi-O-RU Sensing Aggregation

The objective of this experiment is to validate adaptive cooperative sensing using multiple sensing-enabled O-RUs operating under the hierarchical AI orchestration framework. The experiment evaluates two complementary sensing modes: centralized sensing, where range-Doppler maps generated at multiple sensing nodes are transported to a centralized sensing aggregation function, and distributed sensing based on Federated Learning (FL), where AI models are trained locally and only model updates are exchanged. The hierarchical AI framework dynamically selects the sensing mode according to communication traffic, transport resource utilization and compute availability, enabling an adaptive trade-off between sensing accuracy and transport resource consumption.

3.2.2.2.1 Experimental Setup

The experiment is performed on the experimental platform presented in Figure 3-14 , extending the transport infrastructure described in Experiment 1. Multiple sensing-enabled O-RUs capture RF communication signals and generate range-Doppler radar images through the Low-PHY sensing functions implemented at the O-DUs. The sensing information is processed either centrally at the sensing aggregation function hosted at the MEC or locally at the distributed sensing nodes using Federated Learning (FL).




Figure 3-14 a) Lab testbed connectivity with electronic and optical data paths, b) physical nodes

In the centralized mode, range-Doppler maps extracted by the Low-PHY sensing functions at different O-RUs are forwarded through the optical transport network to a centralized sensing aggregation function, where they are jointly processed. The fusion of these multi-view sensing observations enables multi-static sensing and significantly improves trajectory estimation accuracy by exploiting spatial diversity and complementary Doppler and range perspectives from multiple sensing nodes. The sensing aggregation function performs trajectory prediction using the hybrid Convolutional Neural Network - Long Short-Term Memory (CNN-LSTM) architecture described in deliverable D4.4 [11], where convolutional neural networks extract spatial features from the range-Doppler maps and a (Long Short-Term Memory) LSTM model captures their temporal evolution to estimate the trajectory of moving targets.

When transport resources become constrained, sensing processing switches to distributed operation. AI models are trained locally at distributed sensing nodes using locally generated range-Doppler radar images. Instead of transmitting high-volume sensing data over the transport network, the nodes periodically exchange only model updates with a federated aggregation function located at the MEC. The aggregation function combines the received updates using the Federated Averaging (FedAvg) algorithm to construct a global sensing model, which is subsequently redistributed to all participating sensing nodes for the next training round. By exchanging model parameters instead of raw sensing data, the proposed FL framework significantly reduces fronthaul and backhaul bandwidth requirements while preserving data privacy and improving scalability.

3.2.2.2.2 Cooperative Sensing Optimization

The experiment consists of three consecutive phases.

Initially, communication traffic remains low and sensing operates in the centralized mode. Range-Doppler maps generated by multiple sensing-enabled O-RUs are transported to the centralized sensing aggregation function, where multi-view sensing observations are fused to estimate the trajectory of moving targets.

During the second phase, background communication traffic is gradually increased, increasing transport resource utilization. The hierarchical Artificial Intelligence (AI) framework continuously monitors communication traffic, sensing traffic, transport resource utilization and sensing accuracy. When transport resources become constrained, the AI framework dynamically switches from centralized sensing to FL. Each sensing node independently trains its local sensing model and periodically exchanges only model parameters with the MEC aggregation function.

Finally, when communication traffic decreases and sufficient transport resources become available again, the AI framework automatically restores centralized sensing, allowing complete range-Doppler maps to be transported and jointly processed.

The experimental set up is shown in Figure 3-14 .The developed framework was deployed across three servers. The Polatis switch provided the optical datapath to the servers which are interconnected with 10G SFP+ transceivers. A typical 1GbE switch has been used to support the optoelectronic switch functionality. For server 1 a Network Interface Card (NIC) card with 2xSFP+ and 2xGbE transceivers has been used whereas for Server 1xSFP+ and 2x1GbE ports have been activated. In Server 3, only 1GbE ports have been activated which are used for the control of the transport network. Server1 to Server 2 connections are established supporting backhaul and sensing traffic. In addition to this, dynamic background traffic is created between two VMs hosted in server1 and server 2.




Figure 3-15 a) Generated Communication and background at Server b) Switch from Centralized to FL learning, c) Total network traffic

The ability of the proposed AI framework to reconfigure RAN, transport network and compute resources is evaluated for the testbed shown in Figure 3-14 for the scenario where the background traffic generated by a Virtual Machine (VM) in server 1 destined to server 2 was varying over time. In addition to background traffic, the transport network should also provide resources for user-plane (communications traffic) and sensing traffic connectivity. The servers were interconnected using an optical datapath with 10Gbps capacity and an electronic path with 1Gbps capacity. Figure 3-15 a) shows the egress traffic from server 1. During the initial period (0–5 s), the network experiences a total load of approximately 2–3 Gbps, while the backhaul communication traffic generated by the RAN fluctuates between 50–300 Mbps. At $t=5$ s, the background traffic increases significantly, reaching values close to 9–10 Gbps, which brings the total network load close to the transport capacity. After $t \approx 10$ s, the background traffic decreases again to a moderate level. In Figure 3-15 (b) we observe that the proposed framework switches from centralized learning to FL significantly reducing the transport load by transmitting only model updates instead of high-volume sensing data Figure 3-15 (c). The use of FL prevents the sensing traffic from exceeding the available transport capacity. Instead of transmitting full RD maps, only model updates are exchanged, keeping the aggregate traffic below the network capacity limit. When the background traffic decreases again, the Reinforcement Learning (RL) policy switches to centralized sensing. These results demonstrate the adaptive decision-making capability of the proposed hierarchical AI framework, which continuously balances transport resource utilization and sensing

performance according to current network conditions. Consequently, the framework maintains communication service continuity while avoiding transport congestion through intelligent orchestration of the sensing processing mode.

3.2.2.2.3 Results and KPIs

Table 3-8 Experiment Description

Item	Description
Test ID	Test 3.2
Test Name	Adaptive Cooperative Sensing
Objective	Validate adaptive switching between centralized cooperative sensing and Federated Learning
Experimental Platform	Multiple sensing-enabled O-RUs, O-DUs, optical transport network, MEC sensing aggregation function
AI Models	CNN-LSTM trajectory estimation, Federated Averaging (FedAvg)
Traffic	Communication traffic, sensing traffic, FL model updates
Expected Outcome	Maintain sensing performance while reducing transport resource utilization
Related PoC	Transport Interconnect PoC

Table 3-9 Key Performance Indicators

KPI	Value	Comments
Number of sensing nodes	6 O-RUs	Distributed sensing deployment
Number of sensing processing nodes	2 O-DUs	Local sensing processing
MEC aggregation server	1	Federated aggregation (FedAvg)
Sensing processing modes	Centralized / Federated Learning	Dynamically selected by the AI orchestrator
Federated learning algorithm	FedAvg	Global model aggregation
Local training epochs	5 epochs	Per FL round
FL aggregation interval	10 scheduling intervals	Adaptive orchestration
Local batch size	16	Local SGD training
Learning rate	0.05	Local optimizer
Model update size	≈4 MB per DU per FL round	Transported during FL
Range-Doppler map size	≈20 MB per DU	Centralized sensing mode
Transport traffic reduction	≈80%	Compared with centralized sensing
Communication traffic	50–300 Mbps	RAN backhaul traffic
Background traffic	2–3 Gbps → 9–10 Gbps	Dynamic congestion scenario
Optical transport capacity	10 Gbps	Optical datapath
Electronic control path	1 Gbps	SDN control network
Optical path reconfiguration	10.32–11.23 ms	Experimental measurement
AI sensing model	CNN-LSTM	Trajectory prediction
Cooperative learning	FedAvg	Global model aggregation

3.2.2.3 Experiment 3 – Transport Reconfiguration and Resilience

This scenario evaluates the resilience and dynamic reconfiguration capabilities of the proposed multi-technology transport infrastructure under changing traffic conditions and link failures. The orchestration framework continuously monitors the state of the packet and optical transport domains and dynamically updates policies to maintain end-to-end service continuity for both communication and sensing traffic flows. Upon detection of transport link failures, the framework automatically re-establishes connectivity by selecting alternative packet and/or optical paths and reconfiguring the underlying transport resources in real time.

3.2.2.3.1 Experimental Setup

The experiment is performed using the 6G-SENSES packet-optical transport testbed comprising O-RAN radio units, SDN-controlled packet and optical switches, an optical circuit switch, and the hierarchical AI orchestration framework. Radio resources are configured through the SMO/ Operation and Maintenance (OAM) using the O1/NETCONF interface, while optical transport resources are managed by the SDN controller through a dedicated 1 Gbps Ethernet control network. Communication and sensing traffic are transported over the optical infrastructure, allowing the orchestration framework to dynamically reconfigure transport paths according to network conditions.

3.2.2.3.2 Methodology

The hierarchical orchestration framework continuously monitors the operational state of the radio, packet and optical transport domains. Based on network conditions and service requirements, the AI agents trigger appropriate control actions to maintain end-to-end service continuity.

The experimental evaluation consists of two complementary parts. First, representative radio reconfiguration actions are executed through the SMO/OAM, including sensing O-RU creation and activation, radio parameter modification, and VLAN configuration for communication and sensing traffic. In addition, optical datapath reconfiguration is evaluated by dynamically establishing and modifying optical circuits under SDN control. The corresponding reconfiguration times are measured to assess the responsiveness of the proposed transport infrastructure

3.2.2.3.3 Experimental Procedure

The experiment is executed in the following steps:

1. The transport infrastructure is initialized with active communication and sensing services.
2. The RAN AI agent performs representative radio reconfiguration actions, including:
 - creation of sensing-enabled O-RUs,
 - activation of sensing functions,
 - modification of radio parameters,
 - VLAN configuration.
3. The SDN controller dynamically reconfigures optical paths through the optical circuit switch while communication and sensing traffic remain active.
4. The reconfiguration latency of both radio and optical domains is measured and analysed.

3.2.2.3.4 Experimental Results

The experimental evaluation confirms the capability of the proposed framework to rapidly adapt both radio and transport resources.




Figure 3-16RAN reconfiguration time through the SMO/OAM using O1/NETCONF

Figure 3-16 presents the measured reconfiguration latency for representative RAN management operations performed through the SMO/OAM using the O1/NETCONF interface. Among the evaluated operations, the Activate Sensing RU exhibits the highest reconfiguration latency as it requires transitioning the radio unit from a configured state into operational mode. This process involves parameter validation, resource allocation, synchronization and activation of the user-plane processing chain, making it inherently more complex than simple configuration updates.

Figure 3-17 presents the Cumulative Distribution Function (CDF) of the optical datapath reconfiguration time. The SDN controller manages the optical switch through a dedicated 1 Gbps Ethernet control interface, while optical traffic remains on the data plane. The measured optical reconfiguration delay ranges from 10.32 ms to 11.23 ms, with the observed variation mainly attributed to queuing delays within the electronic control network and processing delays at the SDN controller. The results demonstrate deterministic optical path establishment within a few milliseconds, confirming the suitability of the proposed SDN-controlled packet-optical transport network for latency-sensitive ISAC services.




Figure 3-17CDF Optical Datapath reconfiguration time.

Table 3-10 Experiment Description

Parameter	Value
Experiment ID	Test 3.3
Objective	Validation of dynamic radio and optical transport reconfiguration
Transport technologies	Packet switching and SDN-controlled optical switching
Radio management	SMO/OAM (O1/NETCONF)
Optical management	SDN Controller
Control network	1 Gbps Ethernet
Data network	Optical transport
Evaluated functions	RAN reconfiguration and optical path reconfiguration
Measured metrics	RAN configuration latency, optical reconfiguration latency
Target application	Resilient ISAC transport

Table 3-11 Key Performance Indicators

Parameter	Value
Experiment ID	Test 3
Objective	Validation of dynamic radio and optical transport reconfiguration
Transport technologies	Packet switching and SDN-controlled optical switching
Radio management	SMO/OAM (O1/NETCONF)
Optical management	SDN Controller
Control network	1 Gbps Ethernet
Data network	Optical transport
Evaluated functions	RAN reconfiguration and optical path reconfiguration
Measured metrics	RAN configuration latency, optical reconfiguration latency
Target application	Resilient ISAC transport

3.2.2.4 Integration of non-3GPP technologies

The goal of this PoC is to demonstrate integration and delivery of non-3GPP sensing information toward O-RAN xApps through the O-RAN orchestration framework. In particular, the PoC validates how sensing data generated by external sensing platforms and (such as a radar sensor) can be ingested to the network in a secure way through the SMO using the A1-S and delivered toward Near-RT RIC xApps for real-time sensing-aware optimization and control. More specifically, this PoC demonstrates:

- Ingestion of sensing-related information into the SMO through broker-based interfaces and A1-S mechanisms.
- Delivery of sensing information toward sensing-aware xApps operating in the Near-RT RIC.
- Real-time exploitation of external sensing information for communication and sensing optimization functions.
- Fusion of non-3GPP sensing information with 3GPP sensing data.

The overall data ingestion pipeline is shown in Figure 3-18.




Figure 3-18 Evaluation of external sensing data ingestion through the Non-RT RIC

3.2.2.4.1 Experimental Setup

The experimental platform consists of four functional domains:

- an external radar sensor where for this set up in an X-band phased array system;
- the SMO and Non-RT RIC, responsible for sensing registration, filtering, aggregation, and context generation;
- the Near-RT RIC, hosting sensing-aware xApps;
- the O-RAN radio access network, providing communication services and native RAN telemetry.

The radar sensor continuously monitors the surrounding environment and generates sensing observations including target detections, range estimates, velocity measurements, target tracks, confidence indicators, timestamps, and sensor metadata.

The radar observations are transmitted through the external sensing gateway, which converts the native radar output into the common sensing information model adopted by the 6G-SENSES architecture. The sensing information is then ingested into the SMO, where it is validated, filtered, aggregated, and transformed into technology-independent sensing context.

The processed sensing context is exposed by the Non-RT RIC through the A1 interface and delivered to the Near-RT RIC. The A1 mediation functions inject the information into the Fast Message Router, allowing subscribed xApps to combine radar-derived context with real-time RAN measurements collected through the E2 interface.

The equipment used in the experiment is shown in Figure 3-19.




Figure 3-19 Experimental multiband sensing and communication testbed. Hardware setup integrating the FR1 OFDM communication/sensing branch and the X-band phased-array radar

The integration process consists of two sequential stages.

3.2.2.4.2 Stage 1 – Radar Sensor to SMO

The external radar sensor first registers as a sensing producer together with its sensing information schema and sensing processing jobs.

The radar periodically generates sensing observations that are transmitted to the Sensing Control Function (SeCF). The SeCF validates each sensing message against the registered sensing schema before forwarding the measurements to the sensing processing pipeline.

Within the SMO, sensing services perform:

- sensing registration;
- sensing validation;
- sensing filtering;
- sensing aggregation;
- sensing context generation.

The resulting sensing context represents the environment independently of the underlying sensing technology.

3.2.2.4.3 Stage 2 – SMO to Near-RT RIC

The Non-RT RIC exposes the processed sensing information through the A1 interface.

The A1 mediation layer validates the sensing messages, converts them into the internal messaging format of the Near-RT RIC, and injects them into the Fast Message Router.

Subscribed xApps receive the sensing-derived context and combine it with native O-RAN telemetry obtained through the E2 interface. The fused information is then used to support sensing-aware RAN optimization functions such as beam management, mobility optimization, resource allocation, interference mitigation, and service orchestration

3.2.2.4.4 Experimental Procedure

The experimental procedure comprises the following steps:

1. The external radar sensor is activated.
2. The radar sensor registers as a sensing producer within the SMO.
3. The sensing schema and sensing processing jobs are instantiated.
4. The radar continuously generates sensing observations.
5. The observations are transmitted to the SeCF.
6. The sensing information is validated, filtered, and aggregated.
7. The Non-RT RIC generates technology-independent sensing context.
8. The sensing context is delivered through the A1 interface.
9. The Near-RT RIC distributes the sensing information to subscribed xApps.
10. The xApps combine external radar information with E2 telemetry and perform sensing-aware network optimisation.

3.2.2.4.5 Experimental Results

The experiment successfully demonstrates the integration of heterogeneous sensing technologies within the 6G-SENSES O-RAN control architecture. The external X-band phased-array radar operating at 10.2 GHz is successfully registered as an external sensing producer, while its sensing observations are represented through the common sensing information model implemented within the SMO. The sensing processing pipeline validates incoming radar measurements and transforms them into technology-independent sensing context without requiring modifications to either the radar platform or the O-RAN infrastructure.

The Non-RT RIC successfully exposes the generated sensing context through the A1 interface, enabling the Near-RT RIC to distribute sensing information to subscribed xApps. The xApps combine external radar observations with native NR sensing measurements obtained through the E2 interface, demonstrating seamless interoperability between heterogeneous sensing technologies and O-RAN control applications.

The X-band sensing branch validates directional sensing using a shared OFDM waveform and phased-array beamforming. Experimental measurements were performed for four reflector placements corresponding to nominal azimuth directions of -60° , -30° , $+20^\circ$ and $+70^\circ$. The resulting beam-delay response maps exhibit a clear shift of the dominant reflected component according to the reflector position, demonstrating consistent angular discrimination and beam-domain spatial selectivity across the measured field of view. The clutter suppression stage effectively attenuates static reflections, allowing the dominant target response to be clearly identified. The experimental results shown in Figure 3-20 confirm that the X-band sensing subsystem provides reliable directional discrimination rather than high-resolution ranging, which is consistent with its intended beam-domain sensing functionality.




Figure 3-20 Beam-delay response maps obtained for four reflector placements corresponding to nominal azimuth directions of approximately (a) -60° , (b) -30° , (c) $+20^\circ$, and (d) $+70^\circ$. The dominant reflection shifts consistently across steering directions, demonstrating the angular discrimination capability of the X-band sensing branch.

These measurements are then ingested in the 5G-NR system, operating at 2.3 GHz which provides motion-sensitive sensing. As shown in Figure 3-21, Delay-Doppler processing successfully distinguishes approaching and receding targets through positive and negative Doppler signatures appearing on opposite sides of the zero-Doppler axis. The derived motion-direction metric remains consistent with the ground-truth target movements throughout the experiments. Although weak opposite-Doppler components are observed due to indoor multipath propagation, the dominant Doppler response remains clearly distinguishable, demonstrating robust motion-direction discrimination under realistic indoor conditions.




Figure 3-21 FR1 delay-Doppler sensing results for alternating approaching and receding target motion. Positive and negative Doppler signatures are observed on opposite sides of the zero-Doppler axis, demonstrating motion direction discrimination.

Table 3-12 Experiment Description

Parameter	Value
Experiment ID	Test 3.4
Experiment Name	Integration of External Radar Sensing
External sensing technology	X-band phased-array SDR radar
Radar operating frequency	10.2 GHz
Complementary sensing technology	FR1 OFDM sensing
FR1 operating frequency	2.3 GHz
Communication waveform	Shared OFDM waveform
External sensing interface	External sensing producer
Sensing representation	Common sensing information model
O-RAN components	SMO, Non-RT RIC, Near-RT RIC, A1, E2, xApps
External sensing processing	Beam-delay response generation
Native RAN sensing	Delay-Doppler processing
Data fusion	External radar + NR sensing measurements
Demonstrated capability	Multi-technology sensing integration

Table 3-13 Key Performance Indicators – Integration of External Radar Sensing

KPI	Value / Result
-----	----------------

External radar frequency	10.2 GHz
NR sensing frequency	2.3 GHz (FR1)
Communication waveform	Shared OFDM
Number of sensing technologies integrated	2
Number of sensing modalities	2 (Beam-domain + Delay-Doppler)
Radar registration	Successful
Common sensing model generation	Successful
Radar data ingestion into SMO	100% successful
Non-RT RIC context generation	Successful
A1 dissemination	Successful
Near-RT RIC subscription	Successful
xApp sensing data reception	Successful
Fusion of external and NR sensing data	Successful
Directional discrimination	4 angular sectors validated (-60°, -30°, +20°, +70°)
Motion direction discrimination	Approaching and receding targets successfully distinguished
Clutter suppression	Validated
Beam-domain sensing	Validated
Delay-Doppler sensing	Validated
End-to-end interoperability	Successfully demonstrated

3.3 Other small-scale implementations

In this section we report small-scale implementations that could not yet been implemented in the integrated solution

3.3.1 Evaluation of 60 GHz AFE blocks

In **6G-SENSES** we support the development of low-power technology components and foster the sovereignty of EU microelectronics. The **6G-SENSES** millimeter wave (mmWave) ISAC demonstrator features a COTS analogue front-end (AFE) from Sivers Semiconductors, whose purpose is for communications solely. We proposed a modular and low-power AFE of to support MIMO and digital beamforming applications in ISAC.

The design of the AFE has been reported in **D5.2**, while here we report on the latest AFE board measurements. A full integration with RFSoc board and adapter board is in progress and it is expected to be ready over the summer, ready to be shown and presented at the Final Periodic Review. The delays associated with this development stem from delays in the delivery of Surface Mount Device (SMD) components.

The individual RF blocks of the Analog front-end (AFE), namely 60 GHz serial feed antenna and 60 GHz balun, have been manufactured as a standalone structures on the same PCB stack up. They provide individual characterisation of their performance in the RF chain. These individual blocks have been manufactured with specially designed Ground-Signal-Ground (GSG) pads and tested with 325 µm pitch RF probes.




Figure 3-22 (a) zoom in for RF-probing, (b) characterization standards for the PCB stackup and transmission results of the Thru standard

Two different samples of the **60 GHz antenna** are measured, i.e. with decoupling structure (wDS) and without decoupling structure (woDS). The measurement results of the 60 GHz antenna show a resonance frequency shift of 3%, which is attributed to the extension of the feed line due to measurement purposes. The results are given in Figure 3-23.




Figure 3-23 (a) manufactured antenna sample on the probe station chuck, (b) simulation vs measured results

The **balun** samples are prepared in a back-to-back connected fashion (see Figure 3-24) to avoid differential measurement setup. The balun performance was verified with high accuracy as given in the following figure. The wideband operation of the balun serves as the cushion for any response shifts occurring due to the PCB etching inaccuracies in the production of the AFE board.

The individual RF blocks have been integrated on the PCB with the selected PCB stackup. 60 GHz transceiver chipsets have been integrated into the PCB cavities such that to flush the chip pad height to the PCB traces. That results in the short possible wirebond length and hence the associated parasitics. Figure 3-25 shows a microscopic view of four transceiver chipsets integrated into the PCB with wirebonds and associated balun structures for each chipset.




Figure 3-24 (a) Manufactured balun as a back-to-back connected sample, (b) simulation vs. measurement results




Figure 3-25 Microscopic view of the AFE board with four transceiver chipsets integrated in to the RF chain

An adapter board was developed to connect the ultra-low-power MIMO front-end to an AMD/Xilinx UltraScale+ RFSoc board, which handles the radar signal processing. This work was necessary to adapt the design to the interface, the number of I/O channels, and the connector count on the MIMO front-end. In addition, an SDR design was revised to support the eight DAC/ADC channels required by the MIMO front-end.

The adapter board has been designed and manufactured. Integration and assembly of the MIMO RADAR hardware platform are currently underway. Due to the delay in SMD component delivery, the final hardware solution and its performance will be reported on at the Final Review audit.

3.3.2 Demonstration: LiveSense - Real-Time Wi-Fi Sensing Platform for Range-Doppler on COTS Laptops

As part of the experimental proof of concept and technology validation for the 6G-SENSES network solution, Intel (INT) successfully evaluated and demonstrated **LiveSense**, a real-time cross-platform Wi-Fi sensing system. This demonstration directly addressed the project's key performance indicators (KPIs) for ISAC by

transforming commercial off-the-shelf (COTS) laptops into high-precision, centimetre-level Range-Doppler sensors without compromising simultaneous data communication.

1. Public Demonstration and Real-World Validation

To validate the technology beyond controlled laboratory conditions, the LiveSense platform was exhibited as a live, real-time public demonstration at the prestigious **IEEE PerCom conference in March 2026**.

The public exhibition served as a rigorous real-world testbed. During the event, the system was subjected to dynamic, unpredictable interference and continuous testing by multiple conference attendees. Despite the crowded and noisy wireless environment, the demo performed flawlessly. Attendees were able to interact with the system in real-time, successfully validating its robustness, high accuracy, and immediate responsiveness.

2. Technical Capabilities Showcased

The demonstration highlighted the system's ability to operate in a fully monostatic configuration (using a single device's transceiver) without requiring any external infrastructure, access points, or radar hardware. Key technical achievements showcased during the live demo included:

- **Cross-Platform Robustness:** The system ran seamlessly on standard laptops equipped with COTS Intel AX211 (Wi-Fi 6E) and Intel BE201 (Wi-Fi 7) Network Interface Cards (NICs).
- **Real-Time Human Presence and Motion Tracking:** The GUI continuously streamed and visualized the distance and radial velocity of the interacting attendees within a 5-meter radius.
- **Micro-Motion and Gesture Ranging:** The demo proved the system's high-fidelity capabilities by accurately detecting minute movements, including user respiration and hand-gesture ranging.
- **Low Latency and High Accuracy:** real time visualization of the user position and velocity

The successful deployment at PerCom 2026 proved that the **6G-SENSES** monostatic Wi-Fi sensing framework is highly mature and scalable. The ability of the LiveSense platform to extract fully synchronized Channel State Information (CSI) at 100 Hz, perform on-device time-phase alignment, and cancel self-interference in real-time, all while operating perfectly under the stress of continuous public interaction, confirms the viability of utilizing existing Wi-Fi infrastructure as an advanced 6G sensing ecosystem.

4 PoC#2 – CF-MIMO Prototype

PoC#2 delivers a CF-MIMO PHY prototype with closed-loop reconfiguration capabilities. The algorithms related to this implementation are deployed on top of representative real-time SDR hardware components to demonstrate the feasibility of the concept as well as achieving high network performance. All relevant aspects of an operational CF-MIMO network are demonstrated, including tight timing synchronization, OTA reciprocity calibration for TDD operation, distributed signal processing for channel estimation and beamforming, and fronthaul/backhaul data plane traffic exchange. In addition to the main platform at **TUBS**, a small-scale platform has been developed at **UC**, whose results and associated KPIs complement those obtained by the TUBS platform. Both have focused on PHY-layer implementation aspects. In the following, we present the hardware setup for both CF-MIMO prototypes, as well as the results and lessons learned—particularly its integration into OAI.

4.1 Implementation at TUBS

The objective of **PoC#2** was to establish a flexible software platform for experimental research on large-scale digital beamforming and massive MIMO systems. At the beginning of the project, the implementation was planned to build upon the National Instruments (NI) Massive MIMO Prototyping System and its accompanying LabVIEW-based software framework. This appeared to be an attractive solution as it provided an integrated hardware/software environment specifically designed for rapid prototyping of large antenna arrays. Over the course of the project, however, it became evident that the corresponding software framework is no longer actively maintained or supported, making it unsuitable as a sustainable basis for future research developments and long-term operation.

To establish a future-proof experimental platform, the implementation was therefore migrated to the open-source OAI 5G software stack. This migration required a comprehensive investigation of the OAI physical layer implementation and its suitability for large-scale antenna systems. During this work, several important technical limitations were identified.

First, the current OAI physical layer is primarily designed for conventional base stations with a relatively small number of antennas. The Layer-1 processing architecture is optimised for standard cellular deployments and does not scale efficiently when the number of transmit or receive antennas is significantly increased. Extending the implementation towards massive MIMO, therefore, requires substantial redesign of the PHY processing architecture, memory management, and thread organization rather than isolated code modifications. This observation is consistent with ongoing discussions within the OAI developer community, where a redesign of the Layer-1 beamforming framework is currently under active development.

Second, increasing the antenna count beyond the standard configurations revealed several implementation issues related to memory allocation, array indexing, and buffer management. These resulted in memory access violations and buffer overflows throughout different parts of the PHY layer. The identified issues were systematically analysed and corrected by adapting numerous source files within the OAI implementation. Although these modifications enabled operation with larger antenna configurations, they also demonstrated that current OAI releases have not yet been fully optimised for large-scale massive MIMO research platforms.

A further important finding concerned the implementation of digital beamforming algorithms. While OAI provides the general infrastructure for MIMO transmission, the current software release does not yet contain complete implementations of scalable linear beamforming algorithms required for experimental massive MIMO research. Widely used precoding techniques such as Maximum Ratio Transmission (MRT), Zero Forcing (ZF), and Minimum Mean Square Error (MMSE) beamforming are only partially supported or require substantial additional implementation effort. Consequently, advanced beamforming functionality had to be developed beyond the standard OAI software framework. This finding is also reflected in historical OAI

documentation and in the current roadmap for the ongoing redesign of the Layer-1 beamforming architecture¹.

In parallel, receiver-side improvements, including automatic gain control (AGC) optimisation, were investigated to improve measurement stability and overall system robustness.

Based on the experience gained during these investigations, and considering the remaining project duration, the implementation strategy was adapted during the final phase of the project. Instead of continuing with a full 5G software stack, the experimental platform was migrated to a lightweight implementation directly based on the USRP Hardware Driver (UHD) framework. This approach provides complete control over the digital signal processing chain, enables rapid implementation of custom beamforming algorithms, and avoids the architectural constraints imposed by a complete cellular protocol stack. As a result, the developed platform represents a considerably more flexible basis for future research on experimental massive MIMO beamforming and distributed antenna systems.

Lesson learned. The project demonstrated that existing open-source cellular software platforms provide an excellent basis for standards-compliant 5G experimentation but still require substantial architectural extensions before they can efficiently support research on large-scale massive MIMO systems. In particular, scalable Layer-1 processing, efficient memory management for large antenna arrays, and comprehensive implementations of advanced digital beamforming algorithms remain important research and software engineering challenges. These findings significantly influenced the software architecture adopted for the experimental platform developed within this project.

Future research with 6G-SENSES testbed

The hardware platforms established and further developed within this project will continue to serve as key experimental infrastructures in several follow-up research activities.

The Massive MIMO Laboratory will form one of the central experimental platforms within the German research initiative xG-RIC², where it will support the activities of Work Group TD2 – Next Generation MIMO. The platform will be used for the development and experimental validation of advanced distributed MIMO concepts, scalable beamforming algorithms, AI-assisted radio resource management, and O-RAN functionalities. The experience gained during this project regarding scalable software architectures and flexible beamforming implementations provides an important foundation for these future developments.

In addition, both the Massive MIMO Laboratory and the Mobile MIMO Laboratory will be utilized within the 6G-MIMOLink project³. Although the primary focus of 6G-MIMOLink is on operation in the emerging FR3 frequency range, the existing platforms provide an essential basis for the early development, implementation, and validation of communication algorithms, distributed signal processing concepts, synchronization methods, and experimental procedures. The platforms will therefore be employed for algorithm development and proof-of-concept demonstrations before migrating the developed techniques to dedicated FR3 hardware as it becomes available during the project.

Beyond these specific projects, both laboratories will remain part of the institute's long-term experimental infrastructure. They will continue to support research on massive MIMO, distributed antenna systems, O-

¹ Eurecom Gitlab, "Suggestions and discussion solicitation around L1 implementation of precoding/beamforming", https://gitlab.eurecom.fr/oai/openairinterface5g/-/work_items/1055?utm_source=chatgpt.com

² xG-RIC, „Technology Transfer Hub for Medicine and Mobility of the Future“, <https://xg-ric.de/>

³ 6G-MIMOLink, "Skalierbare leistungsstarke Antenneneinheiten für effiziente und zuverlässige 6G-Netze", <https://www.forschung-it-sicherheit-kommunikationssysteme.de/projekte/6g-mimolink>

RAN, ISAC, and AI-native wireless communication systems, while also serving as an important platform for graduate education and the training of doctoral researchers.

4.1.1 CF-MIMO Prototype

The CF-MIMO prototype at **TUBS** is implemented on top of the existing laboratory infrastructure for large-scale SDR experiments. The platform combines a synchronized multi-USRP radio infrastructure, a dedicated Ethernet and compute backbone, and a custom UHD/RFNoC-based software stack for channel estimation and beamforming. While the conceptual CF-MIMO transmission schemes and their mathematical background are described in earlier project deliverables (cf. [D2.3](#), [D3.1](#), [D3.2](#)), this section focuses on the practical laboratory implementation used for the hardware-based validation presented in this deliverable.




Figure 4-1 Schematic overview of the TUBS CF-MIMO laboratory infrastructure, including time server, Master OctoClock, OctoClock distribution, Ethernet networking, compute nodes, and USRPs

A schematic overview of the laboratory-side infrastructure is shown in Figure 4-1. The setup is built around a shared timing and synchronization architecture (green). A GNSS-disciplined time reference provides the global timing source, which is distributed through a central Master OctoClock and several downstream OctoClock units. These distribute 10 MHz reference clock and PPS signals to the USRP devices, enabling coherent operation across a large number of RF frontends. In parallel, the SDR devices and compute nodes are connected through the laboratory Ethernet infrastructure (red), which provides the required control and sample-streaming connectivity between the radio hardware and the processing hosts.

The main radio platform is the existing **TUBS** MIMO rack, which comprises **24 Ettus USRP X310 devices**, three OctoClock synchronization units, and the associated network switching infrastructure. Each USRP X310 provides two independently tunable RF chains covering 10 MHz to 6 GHz, with a maximum instantaneous real-time bandwidth of 160 MHz and I/Q sample rates of up to 200 MS/s. This makes the rack suitable for scalable multi-antenna experiments with a large number of synchronized RF chains. For the CF-MIMO measurements reported in Section 4.1.1.2, eight X310 devices are used as the transmit-side array, corresponding to 16 active TX channels. The receive side is implemented using a further set of X310 devices

from the same synchronized infrastructure, with 15 active RX channels used during the reported measurement campaign.

The same X310-based radio hardware is used for both physical antenna deployments investigated in this deliverable. In the **co-located massive-array configuration**, all transmit antennas are arranged as one compact array. In the **CF configuration**, the same transmit antennas are split into two spatially separated antenna groups. Thus, the comparison between both deployments is based on identical RF hardware and identical measurement parameters; only the physical antenna geometry is changed. The exact antenna arrangements, distances, and measurement parameters are described in Section 4.1.1.1.

In addition to the existing X310-based infrastructure, **4 Ettus USRP X410** devices were procured within the project and integrated into a portable mobile rack. Each X410 provides four independently tunable transmit and receive channels, covering 1 MHz to 7.2 GHz, tunable up to 8 GHz, with up to 400 MHz instantaneous bandwidth per channel and sample rates of approximately 490 MS/s. Compared to the X310, the X410 offers higher RF-chain density, higher instantaneous bandwidth, and higher sample-rate capability, making it particularly suitable for future receive-side or UE-side extensions of the CF-MIMO platform. Although the X410 devices include onboard GPSDO/OCXO reference capabilities, the intended laboratory operation follows the same external synchronization concept as the rest of the infrastructure, using the OctoClock-distributed reference and PPS signals.

The X410-based mobile rack was originally introduced in the context of the earlier OAI-based CF-MIMO implementation activities. Following the transition towards the native UHD/RFNoC implementation described in this deliverable, the hardware remains relevant as a higher-bandwidth extension of the platform, in particular for future receive-side and mobile demonstration scenarios. For the final measurement campaign, the X410 devices were used as the receive-side platform, replacing the X310-based RX chain used during earlier validation runs.

On the software side, the prototype uses a custom Python implementation built directly on UHD and RFNoC. This native UHD approach provides direct control over RFNoC graph construction, device synchronization, timed tuning, TX/RX streaming, pilot generation, channel estimation, and beamforming-weight computation. The implementation is structured into hardware-facing TX/RX classes, a transmission orchestration layer, a hardware-independent DSP layer, and offline analysis scripts. This modular structure allows the same measurement procedure to be repeated across different antenna geometries and beamforming strategies while keeping the RF configuration and signal-processing pipeline consistent.

The following subsection describes the final deployment and measurement configuration, including the two antenna geometries, the downlink channel-estimation procedure, the implemented joint and distributed MRT beamforming modes, and the measurement parameters used for the reported KPI evaluation.

4.1.1.1 Final deployment and configuration

The CF-MIMO prototype described in this section was deployed as a controlled over-the-air (OTA) measurement platform for evaluating distributed antenna deployments and cooperative beamforming strategies under realistic laboratory conditions. Building on the laboratory infrastructure introduced in Section , the experiment uses the same synchronised USRP-based radio platform for all investigated scenarios. The different test cases are, therefore, not defined by different hardware, but by the physical placement of the antenna arrays and by the beamforming algorithm applied to the estimated downlink channel.

The underlying mathematical concepts for the CF-MIMO transmission configurations were introduced in deliverable **D3.2**, Section 3.1.1 [7]. In particular, **D3.2** distinguishes between non-cooperative transmission, cooperative joint transmission on shared resources, and phase-aligned multi-beam transmission with per-RU constraints. The measurements reported here focus on the experimental validation of the corresponding

MRT-based beamforming primitives on real hardware. More specifically, the implementation evaluates how accurately a desired RX position can be spatially addressed when the transmit antennas are either arranged as one co-located massive array or as two spatially separated CF sub-arrays.

To keep the experimental comparison controlled and reproducible, the same TX and RX hardware, frequency, gain, waveform parameters, calibration procedure, and analysis pipeline are used for both physical topologies. Only the antenna geometry is changed between the two deployments.

Beamforming configurations

For each physical antenna topology, two beamforming strategies are evaluated. These strategies correspond to the MRT-based beamforming principles discussed in **D3.2**, Section 3.1.1, but are implemented here in a controlled measurement setup using a larger number of RX observation points. The purpose of this subsection is, therefore, to explain the beamforming concepts, while the concrete physical antenna deployments used for the measurements are described afterwards.

The first method is referred to as joint MRT beamforming. In this case, all active TX antennas are treated as one coherent antenna array, and a single MRT weight vector is computed over the full 16-element TX aperture for each target RX channel. This corresponds to the fully coordinated case, where the complete TX array contributes jointly to focusing power at the selected receive position.

Figure 4-2 illustrates this joint transmission concept using the simplified two-UE representation introduced in deliverable **D5.1** [1] and illustrates joint coherent transmission over the complete TX aperture. The figure is used to explain the beamforming principle and does not depict the exact measurement geometry used in this deliverable. In the actual experiments, the two UE positions are replaced by a linear RX measurement array with 15 receive antennas, allowing the spatial power distribution and leakage towards neighbouring RX positions to be evaluated.




Figure 4-2 Joint beamforming concept, redrawn/reused from D5.1

The second method is referred to as distributed MRT beamforming. Here, the 16 TX antennas are partitioned into two groups of eight antennas, matching the physical grouping of the CF deployment. Each group computes its own MRT weight vector based only on the channel coefficients of that group, and the two group-wise weight vectors are then combined for transmission. This realises a per-group power-normalised beamforming strategy, which is intentionally not identical to joint MRT. It approximates a distributed antenna deployment in which each RU or antenna group has its own local power budget.

Figure 4-3 illustrates the distributed/group-wise beamforming concept, again using the two-UE representation from D5.1. The figure shows two antenna groups serving UEs and is used to explain the distributed MRT principle. As for the joint case, the figure serves as a conceptual illustration of the beamforming method rather than as a drawing of the exact 15-RX measurement setup. This distinction is important: in the following experiments, distributed beamforming refers to the signal-processing method, whereas the CF topology refers to the physical placement of the antenna groups.




Figure 4-3 Distributed/group-wise beamforming concept, redrawn/reused from D5.1

The mathematical formulation of the corresponding CF-MIMO transmission schemes, including MRT-based phase compensation and cooperative/distributed multi-beam concepts, is provided in D3.2, Section 3.1.1. The implementation in this deliverable uses these principles as beamforming primitives, but evaluates them experimentally through sequential single-target transmissions and a full spatial RX power measurement. Pairwise SINR estimates reported in Section 4.1.1.2 are then derived from the measured target-by-receiver power matrices.

Physical antenna deployments

The two beamforming strategies described above are evaluated under two different physical TX antenna topologies. Both topologies use the same 16 TX antenna ports from the X310-based MIMO rack, only the physical antenna placement is changed. This makes the comparison between the two topologies controlled and reproducible, since the hardware, carrier frequency, gains, waveform parameters, calibration procedure, and analysis pipeline remain identical.

The RX side is kept fixed for both deployments. It consists of 15 active receive antennas arranged as a uniform linear array with 8 cm element spacing. During each beamforming transmission, the received power is measured simultaneously at all RX channels. Thus, each transmission provides not only the received power at the intended target position, but also the leakage towards all other RX positions. This 15-point RX array is the experimental counterpart to the simplified UE positions shown in the conceptual beamforming figures above.

In the co-located massive-array topology (see Figure 4-4), all 16 TX antennas are arranged as one uniform linear array with 4 cm element spacing. The TX array is placed parallel to the RX array and centered with respect to the middle RX antenna. The distance between the TX and RX array centers is 1.5 m. This topology represents a conventional centralized massive-MIMO deployment, where all transmit antennas are located in one compact aperture and can be used jointly for beamforming.




Figure 4-4 Physical deployment of the co-located massive-array topology

In the CF topology (see Figure 4-5), the same 16 TX antennas are split into two spatially separated sub-arrays of 8 antennas each (left part of the figure). The centers of the two TX sub-arrays are separated by 2 m. Both sub-arrays are rotated such that they point towards the center of the RX array, and the distance from each sub-array center to the central RX antenna is again 1.5 m. This topology emulates a simplified CF or distributed antenna deployment, where the effective transmit aperture is enlarged by separating the antenna groups while still operating them from the same synchronized laboratory infrastructure.




Figure 4-5 Physical deployment of the CF topology

The physical two-group structure of the CF topology matches the algorithmic grouping used for distributed MRT beamforming. In other words, the two physical TX sub-arrays correspond to the two beamforming groups used by the distributed MRT implementation. However, both joint and distributed MRT are evaluated for both physical topologies, so that the influence of the antenna geometry and the influence of the beamforming strategy can be compared separately.

The laboratory environment is intentionally not an anechoic or reflection-controlled setup. Reflections from walls, racks, equipment, and surrounding objects are present and contribute to the measured channel. The results therefore reflect a realistic indoor laboratory propagation environment rather than an ideal free-space array response.

Downlink channel estimation and feedback

The channel state information used for beamforming is obtained directly from a downlink calibration transmission. During calibration, all TX antennas transmit known Zadoff-Chu pilot sequences simultaneously. Each TX antenna uses a different pilot sequence, so that the individual TX-to-RX channel contributions can be separated at the receiver by cross-correlation.

From the received pilot signals, the full complex downlink channel matrix is estimated. Its rows correspond to TX antennas and its columns correspond to RX channels. This matrix is then used by the central processing software to compute the beamforming weights for every target RX channel and for every beamforming method. Conceptually, this corresponds to downlink channel estimation followed by feedback of the measured channel state to the transmitting access points or antenna groups.

The mathematical basis for the MRT weight computation is given in deliverable **D3.2**, Section 3.1.1. The present implementation does not repeat the full derivation, but uses the same principle: the channel phase is compensated such that the contributions from the selected transmit antennas combine coherently at the intended RX position.

Because the channel is estimated directly in the downlink direction used for the subsequent beamforming transmission, no reciprocity-based channel conversion is required for the measurements reported here.

Measurement procedure

The complete processing and measurement flow is shown in Figure 4-6. The experiment is organised hierarchically into sessions, iterations, calibration rounds, and individual beamforming transmissions.

At the beginning of each session, the USRP hardware is initialized, synchronized, and tuned. The RFNoC graphs for the TX and RX sides are created once and then kept alive for the complete session. Although TX and RX share the same external reference clock and PPS distribution, they are operated as two independent RFNoC graphs with independent timekeepers. This reflects the intended base-station-to-UE measurement model, where the transmitter and receiver are not controlled by one shared RFNoC time domain.

Each iteration starts with a fresh channel calibration. The calibration consists of several calibration rounds. In each round, the TX side sends the code-multiplexed Zadoff-Chu pilot burst, while the RX side records all RX channels. The received data is correlated with the known pilot sequences to estimate one channel matrix per calibration round. Since TX and RX operate on independent timekeepers, repeated calibration rounds may contain a common phase and scale offset. Therefore, the per-round channel matrices are phase-aligned before averaging. This produces one aggregated channel matrix for the current iteration.

After calibration, the beamforming weights are computed from the aggregated channel matrix. For every target RX channel, the software computes one joint MRT weight vector and one distributed MRT weight vector. The two methods are then tested back-to-back for the same target before moving to the next RX channel. This ordering minimizes the impact of slow environmental drift when comparing joint and distributed beamforming.

For each target and each beamforming method, the measurement procedure consists of three steps. First, the RX noise floor is measured with the TX side idle. Second, a common random test signal is transmitted using the corresponding beamforming weights. Third, the received signal power is measured simultaneously at all RX channels. This produces a complete target-by-receiver power matrix: each row corresponds to one intended target RX channel, while each column corresponds to one actually observed RX channel.

The resulting power matrices form the basis for all KPIs reported in Section 4.1.1.2, including received signal strength, beamforming margin, Signal-to-Leakage-plus-Noise Ratio (SLNR), and pairwise SINR estimates.




Figure 4-6 Measurement and processing workflow. Important: calibration, H-estimation, H-averaging, and beamformer computation must be inside the iteration loop

Measurement parameters

The measurement configuration is kept identical for the co-located massive-array topology and the cell-free topology. The only intended difference between the two data sets is the physical TX antenna placement described above. All waveform parameters, calibration settings, beamforming methods, gain settings, and analysis settings are kept unchanged to enable a direct comparison between the two topologies.

For each topology, the experiment is repeated over 31 independent sessions with 30 iterations per session. As described in the previous subsection, these sessions correspond to independent UHD/RFNoC re-initializations of the measurement stack, while the iterations provide repeated calibration and beamforming measurements within one session. The resulting configuration is summarised in Table 4-1.

Table 4-1 Measurement parameters for the CF-MIMO beamforming experiments

Parameter	Value
TX hardware	8 × USRP X310
Active TX channels	16
RX hardware	8 × USRP X310
Active RX channels	15 (channels 0–14)
Center frequency	3.6 GHz
Sample rate	1 MHz
TX gain	20 dB
RX gain	25 dB
Beamforming methods	joint MRT, distributed MRT
Distributed beamforming groups	2 groups × 8 TX antennas
Pilot sequence type	Zadoff-Chu
Zadoff-Chu sequence length	1021 samples
Calibration rounds per iteration	5
Calibration duration per round	3 s
Beamforming data duration	0.5 s
Iterations per session	30
Independent sessions per topology	31
Noise measurement	before every beamforming transmission

The measurements stored with this configuration form the basis for the received-power heatmaps, beamforming margins, SLNR values, pairwise SINR estimates, and cross-session stability analysis reported in Section 4.1.1.2.

4.1.1.2 Results and KPIs

This section summarises the experimental results obtained with the CF-MIMO prototype described in Section 4.1.1.1. The evaluation compares two physical antenna topologies — the co-located massive-array topology and the spatially separated cell-free topology — and two MRT-based beamforming strategies — joint MRT and distributed MRT. For every combination of topology and beamforming method, the received signal power is evaluated across all RX positions, and several KPIs are derived from the resulting target-by-receiver power matrices.

The main KPIs considered are:

- **Received signal strength**, represented as a target-by-receiver power matrix.
- **Beamforming margin**, defined as the power difference between the intended target RX channel and the strongest non-target RX channel.
- **Dominance rate**, i.e. the fraction of sessions in which the intended RX channel is the strongest receiver.

- **Empirical SLNR**, computed from the measured received power at the intended RX channel relative to leakage towards all other RX channels.
- **Pairwise SINR estimates**, computed from the measured power matrices to estimate the interference between selected RX positions.
- **Cross-session stability**, quantified through the variation of the estimated channel matrix and KPI values over independent hardware/software re-initializations.

The pairwise SINR values reported in this section are derived from sequential single-target measurements under the assumption of linear power superposition. They should therefore be interpreted as an estimate of the potential mutual interference between two simultaneously served RX positions, rather than as a direct measurement of a simultaneous multi-user transmission. A direct simultaneous multi-user validation with code-separated data streams is left as a possible extension.

All reported values are aggregated over 31 independent sessions per topology, with 30 iterations per session. No session failed in either measurement campaign.

CF topology

The CF topology uses two spatially separated 8-element TX sub-arrays, both oriented towards the center of the RX array. Figure 4-7 shows the mean received-power matrix for the CF topology. Each row corresponds to one intended target RX channel, while each column corresponds to the actually observed RX channel. The diagonal therefore represents successful spatial focusing towards the intended target, while off-diagonal entries correspond to leakage or sidelobes.




Figure 4-7 Mean received-power matrices for the CF topology, averaged over 31 independent sessions. Rows denote the intended target RX channel, columns denote the observed RX channel, and red boxes mark the intended target position. Left: joint MRT. Right: distributed MRT

For the CF topology, both beamforming methods achieve reliable spatial focusing over all evaluated targets. The intended target channel is the strongest received channel in every evaluated session and for every RX target, resulting in a dominance rate of 100 % for both joint MRT and distributed MRT.

The distributed MRT implementation achieves a mean beamforming margin of approximately **2.47 dB**, with per-target mean margins ranging from **0.85 dB** to **4.93 dB**. Joint MRT improves the average margin to approximately **2.98 dB**, with per-target mean margins between **0.73 dB** and **5.72 dB**. This confirms the

expected behaviour: joint MRT benefits from optimizing over the complete 16-element TX aperture, while distributed MRT applies an independent normalization per TX group and is therefore more constrained.

The channel estimation is also stable across independent sessions. The phase-corrected channel-matrix deviation across the 31 independent sessions shows a maximum corrected standard deviation of **-36.68 dB** and a mean corrected standard deviation of **-46.22 dB**, indicating that the calibration and synchronization procedure is highly repeatable across independent UHD/RFNoC re-initializations.

Co-located massive-array topology

The co-located massive-array topology uses the same 16 TX antennas, but arranged as one compact linear array parallel to the RX array. Figure 4-8 shows the corresponding received-power matrix.




Figure 4-8 Mean received-power matrices for the co-located massive-array topology, averaged over 31 independent sessions. Rows denote the intended target RX channel, columns denote the observed RX channel, and red boxes mark the intended target position. Left: joint MRT. Right: distributed MRT

Compared with the CF topology, the massive-array topology shows a broader spatial response. For most RX targets, the intended channel is still the strongest receiver; however, two target positions (**ch4** and **ch11**) are consistently outperformed by neighbouring or nearby RX channels. This results in a dominance rate of **13/15 targets**, or approximately **86.7 %**, for both beamforming methods.

For distributed MRT, the mean beamforming margin over all targets is approximately **0.93 dB**, with a minimum per-target mean margin of **-0.54 dB**. For joint MRT, the mean margin increases to approximately **1.26 dB**, but the minimum per-target margin remains negative at **-0.79 dB**. These negative margins are not interpreted as a software failure. Rather, they reflect the physical limitations of the co-located array geometry in the given laboratory environment: the RX antennas are spaced only 8 cm apart, the room is not reflection-controlled, and the compact TX aperture can produce broader beams and stronger sidelobes or multipath-induced ambiguities than the spatially separated cell-free deployment.

The channel-matrix stability remains good, although slightly worse than for the cell-free topology. Across the 31 independent sessions, the phase-corrected channel-matrix deviation reaches a maximum of **-30.84 dB** and a mean value of **-42.14 dB**.

Pairwise SINR and achievable-rate estimates

To relate the single-target beamforming measurements to a potential multi-user interpretation, pairwise SINR estimates are derived from the measured power matrices. For a given target RX channel, the pairwise

SINR indicates how strongly another RX channel would interfere if its beam were active at the same time. This provides an experimentally grounded estimate of spatial separability between pairs of RX positions.

As an example, the RX pair **ch3/ch9** is selected. This pair is spatially separated within the RX array and avoids the extreme edge positions, while still showing clear separation in the measured CF results. The empirical pairwise SINR values and the corresponding estimated spectral efficiencies are shown in Table 4-2 and Figure 4-9.

Table 4-2 Pairwise empirical SINR and estimated spectral efficiency for RX pair ch3/ch9

Topology	Method	SINR ch3 with ch9 as interferer (dB)	Rate (bit/s/Hz)	SINR ch9 with ch3 as interferer (dB)	Rate (bit/s/Hz)
CF	Distributed MRT	26.96	8.96	28.68	9.53
CF	Joint MRT	32.24	10.71	33.43	11.11
Massive	Distributed MRT	9.27	3.25	12.80	4.33
Massive	Joint MRT	14.79	4.96	13.56	4.57




Figure 4-9 Estimated spectral efficiency for the selected RX pair ch3/ch9, derived from empirical pairwise SINR values. The two bars per case correspond to the two ordered interference directions: target ch3 with ch9 as interferer, and target ch9 with ch3 as interferer

The pairwise comparison highlights the improved spatial separability of the cell-free topology for this RX pair. While the co-located massive topology still achieves usable SINR values, the cell-free deployment provides substantially higher pairwise SINR due to its larger effective aperture and more favourable angular diversity.

Pairwise SLNR profile over the RX array

In addition to the pairwise SINR estimate for a selected RX pair, we evaluate how strongly a beam targeted at one RX channel leaks towards the remaining RX positions. For a target RX channel i , the pairwise SLNR profile compares the received power at the intended RX channel with the received power at each individual non-target RX channel during the same beamforming transmission.

For each observed RX channel j , this pairwise SLNR value is computed as the difference between the target-channel power and the leakage power towards channel j , both taken from the same row of the received-power matrix. Higher values therefore indicate better spatial isolation of the beam towards the intended target. The diagonal entry is omitted, since it would compare the target channel with itself.

This metric differs from the pairwise SINR estimate discussed above. The pairwise SINR evaluates how another user's beam would interfere at the target receiver. In contrast, the pairwise SLNR profile evaluates where the target beam itself leaks across the RX array.




Figure 4-10 Pairwise SLNR profile for target RX channel ch9. For each observed non-target RX channel, the plotted value is the received power at the intended target channel minus the received power at the observed RX channel during the same beamforming transmission. Higher values indicate better spatial isolation of the beam. The diagonal entry ch9/ch9 is omitted

For the cell-free topology, target **ch9** shows a pronounced spatial isolation profile: leakage is strongly suppressed for many RX positions, while the lowest pairwise SLNR values occur at nearby or geometrically coupled RX channels such as **ch8**, **ch10**, and **ch12**. This behaviour is consistent with the finite spatial resolution of the array.

For the co-located massive-array topology, the same target exhibits a flatter and generally lower pairwise SLNR profile. The reduced separation between the target channel and neighbouring RX channels reflects the broader spatial response of the compact TX aperture in the considered indoor environment.

Overall, the pairwise SLNR profiles provide a direct view of the spatial leakage pattern of one beam across the RX array. They complement the received-power heatmaps and support the same conclusion: the cell-free topology provides improved spatial discrimination over a larger part of the RX array, while the co-located massive topology is more sensitive to nearby RX positions and multipath-induced ambiguities.

Overall comparison of beamforming methods and topologies

The combined KPI overview in Table 4-3 and the margin comparison in Figure 4-11 summarise the main trends observed across all four evaluated cases. The table condenses the average behaviour over all RX targets, while Figure 4-11 shows how the beamforming margin varies from target to target. This per-target view is important, because the difference between the two physical topologies is not only visible in the average margin, but also in the systematic behaviour at specific RX positions.

Table 4-3 Summary of key beamforming KPIs

Topology	Method	Mean margin (dB)	Min. margin (dB)	Dominant targets	Mean empirical SLNR (dB)	H stability mean / max corrected std (dB)
Cell-free	Distributed MRT	2.47	0.85	15/15	-3.34	-46.22 / -36.68
Cell-free	Joint MRT	2.98	0.73	15/15	-2.72	-46.22 / -36.68
Massive	Distributed MRT	0.93	-0.54	13/15	-3.91	-42.14 / -30.84
Massive	Joint MRT	1.26	-0.79	13/15	-3.16	-42.14 / -30.84




Figure 4-11 Beamforming margin per target RX channel for both physical topologies and both beamforming methods. Positive values indicate that the intended RX channel received more power than any other RX channel. The cell-free topology maintains positive margins for all targets, whereas the co-located massive topology exhibits systematic degradation at selected RX positions

As shown in Table 4-3, the cell-free topology achieves positive margins for all RX targets and both beamforming methods. This is also visible in Figure 4-11, where all CF margin values remain above the 0 dB threshold. In contrast, the co-located massive-array topology exhibits negative margins for selected targets, most notably **ch4** and **ch11**. These channels are the same target positions that already appeared as problematic in the received-power matrices, confirming that the effect is not an isolated outlier but a systematic consequence of the compact array geometry and the laboratory propagation environment.

The comparison also shows the expected difference between joint MRT and distributed MRT. For a fixed physical topology, joint MRT generally provides higher average margins and better empirical SLNR values than distributed MRT. This trend is visible in the KPI table and in the per-target margin plot: in the cell-free topology, joint MRT increases the mean margin from **2.47 dB** to **2.98 dB**, while in the co-located topology it increases the mean margin from **0.93 dB** to **1.26 dB**. This is consistent with the fact that joint MRT optimises over the full 16-antenna aperture, whereas distributed MRT applies an independent normalization per antenna group.

At the same time, the physical topology has a stronger impact than the choice between the two MRT variants. Both joint and distributed MRT perform more robustly in the cell-free topology than in the co-located massive topology. The CF deployment maintains **15/15 dominant targets** for both methods, whereas the co-located topology reaches **13/15 dominant targets**. This confirms the main spatial conclusion of the experiment: in the tested indoor geometry, separating the TX aperture into two sub-arrays improves spatial discrimination across the RX measurement line.

For a compact visual overview, Figure 4-12 additionally combines all four received-power matrices in a single comparison plot.




Figure 4-12 Comparison of mean received-power matrices across physical topologies and beamforming methods.
Rows correspond to physical topologies and columns correspond to beamforming methods. This overview highlights the stronger spatial selectivity of the CF topology and the broader leakage structure of the co-located massive-array topology

The 2×2 power-matrix overview in Figure 4-12 ties these KPI observations back to the raw received-power distributions. The cell-free matrices show a consistently strong diagonal structure, while the co-located matrices show stronger off-diagonal components around specific target regions. This visual pattern explains both the reduced margins and the reduced dominance rates observed for the co-located topology. Together, the table and figures demonstrate that the cell-free deployment provides the most robust spatial separation in the evaluated laboratory setup, while joint MRT provides the best performance within each topology.

4.1.1.3 Summary

The measurements demonstrate that the native UHD/RFNoC CF-MIMO prototype can reliably estimate the downlink channel, compute MRT-based beamforming weights, and evaluate the resulting spatial focusing performance over many independent hardware/software re-initializations. Across both evaluated topologies, the intended target channel is correctly addressed for most or all RX positions, and joint MRT consistently provides higher average margins than distributed MRT.

The comparison between the two physical topologies shows a clear trade-off. The cell-free topology provides superior spatial selectivity and maintains positive margins for all RX targets, whereas the co-located massive-array topology exhibits broader beams and systematic leakage at selected RX positions. These observations are consistent with the different effective apertures and angular diversity of the two deployments, as well as with the non-anechoic laboratory environment.

A corresponding Green-function-based simulation of the measured antenna geometry was used as a reference model and closely reproduces the experimentally observed topology-dependent spatial selectivity. As illustrated in Figure 4-13 and Figure 4-14, the distributed deployment comprising two 8-element arrays exploits near-field propagation more effectively due to its larger effective aperture. This enables tighter spatial focusing and improved control of the electromagnetic field at the intended receiver locations. In contrast, the co-located 16-element massive MIMO array provides a smaller effective aperture and is therefore able to concentrate only part of the transmitted energy onto the desired receiver.

The measurement results, shown in the four matrices above (Figure 4-12), exhibit the same qualitative behavior and are in good agreement with the simulation. In the distributed configuration, pronounced off-diagonal components are observed, indicating that the wider aperture enables stronger and more spatially selective focusing across different receiver positions. In contrast, the co-located massive MIMO configuration is dominated by the main diagonal, reflecting a more localized beam pattern with less pronounced spatial selectivity over the considered measurement grid.




Figure 4-13 Numerical simulation of the spatial beampattern based on point-sampled Green model for distributed setup and MRT to one antenna located in the center




Figure 4-14 Numerical simulation of the spatial beampattern based on point-sampled Green model for the massive MIMO setup and MRT to one antenna located in the center

The following tables provide a summary/update of the Tests that have been described in the sections above.

Table 4-1 Software platform implementation and migration

Test 4.1	Implementation of a scalable software platform for Cell-free and Massive MIMO experiments
Phase	Validation and performance evaluation
Description	The software environment for the TUBS Cell-free/Massive MIMO prototype was developed and evaluated in several stages. The initially envisaged LabVIEW-based software of the National Instruments Massive MIMO Prototyping System could not be retained because the corresponding software framework is no longer actively maintained. The implementation was subsequently migrated to OpenAirInterface to investigate the suitability of a standards-compliant open-source 5G stack for large-scale multi-antenna operation. This evaluation revealed limitations concerning PHY-layer scalability, memory management for antenna configurations beyond the commonly supported setups, and the availability of complete digital beamforming implementations. Several memory-access, indexing, and buffer-management issues were identified and corrected, and receiver-side improvements, including AGC optimisation, were investigated. Based on these findings, the final prototype was implemented using a lightweight UHD/RFXNoC-based software architecture. This solution provides direct control of the baseband processing chain and enables the flexible implementation of channel estimation, calibration, and custom beamforming algorithms.
KPIs	• Successful migration from the discontinued NI/LabVIEW framework to a maintainable software architecture • Identification and correction of memory-access and buffer-management issues for enlarged antenna configurations • Evaluation of the scalability of the OAI PHY implementation for large-scale MIMO operation • Implementation of a UHD/RFXNoC-based signal-processing chain • Support for custom channel-estimation and beamforming algorithms • Improved receiver stability through AGC optimisation • Reproducible control of the multi-USRP experimental platform
PoC	PoC#2 – TUBS Cell-free/Massive MIMO prototype

Table 4-2 Deployment and validation of the Cell-free/Massive MIMO experimental infrastructure

Test 4.2	Deployment and validation of the CF-MIMO experimental infrastructure
Phase	Validation and performance evaluation
Description	A scalable SDR-based laboratory infrastructure was deployed to experimentally investigate Cell-free and Massive MIMO transmission. The platform is based on synchronized Ettus USRP X310 software-defined radios connected through a dedicated Ethernet backbone and a common timing and synchronization network. A GNSS-disciplined timing reference distributes a common 10 MHz reference clock and PPS signals via Master and slave OctoClock units, enabling coherent operation of all radio front-ends. The platform supports flexible antenna deployments ranging from conventional co-located massive MIMO arrays to spatially distributed Cell-free configurations using identical RF hardware and measurement procedures. The laboratory infrastructure provides a flexible environment for channel sounding, reciprocity calibration, distributed beamforming, and over-the-air validation of advanced PHY-layer algorithms.
KPIs	• Stable synchronization of all SDR nodes using common clock and PPS reference • Coherent operation of distributed USRP X310 radios • Flexible support for both co-located Massive MIMO and distributed Cell-free deployments • Real-time operation with up to 160 MHz instantaneous bandwidth per RF chain • Reproducible over-the-air measurement environment • Support for distributed channel estimation, reciprocity calibration, and beamforming experiments
PoC	PoC#2 – TUBS Cell-free/Massive MIMO prototype

Table 4-3 Experimental validation of Cell-free and Massive MIMO systems

Test 4.3	Experimental comparison of Cell-free and co-located Massive MIMO beamforming
Phase	Validation and performance evaluation
Description	The developed SDR platform was used to experimentally compare two antenna deployment strategies using identical RF hardware and signal processing: (i) a conventional co-located 16-element massive MIMO array and (ii) a distributed Cell-free deployment consisting of two spatially separated 8-element arrays. Beamforming performance was evaluated by measuring the spatial receive-power distribution over a predefined measurement grid. The measurements were compared with Green-function-based electromagnetic simulations of the identical antenna geometry to validate the experimental observations. The comparison demonstrates that the distributed deployment benefits from the larger effective aperture, leading to improved near-field spatial selectivity and tighter focusing of the transmitted energy. The measured receive-power distributions closely match the simulated field distributions, confirming the validity of the underlying propagation and beamforming models.
KPIs	• Experimental validation of coherent distributed beamforming • Comparison of co-located Massive MIMO and Cell-free deployments under identical hardware conditions • Improved spatial selectivity through distributed antenna deployment • Validation of near-field beam focusing effects • Close agreement between Green-function simulations and over-the-air measurements • Repeatable and reproducible measurement results
PoC	PoC#2 – TUBS Cell-free/Massive MIMO prototype

4.1.2 RIS-based PoC Demonstrator

In addition to the CF-MIMO prototype described in Section 4.1.1, **TUBS** implemented a RIS-based demonstrator to evaluate passive, software-controlled beam steering in a compact laboratory setup. The RIS hardware was procured within the project and complements the active SDR-based beamforming experiments by addressing a different class of reconfigurable radio environment: instead of shaping the

transmitted signal at an active antenna array, the propagation path itself is modified by programmable reflecting surfaces.

The general RIS integration concept was already introduced in **D5.1**, Section 5.5.2 [REF]. In that deliverable, the RIS-based PoC was described as a platform for studying RIS-assisted transmission, including multi-panel operation, software-controlled configuration, and cascaded RIS links. The present deliverable builds on this concept and reports a concrete laboratory demonstration of the cascaded RIS scenario.

The demonstrator uses NovoFlect RIS-5Mk1-16x16⁴ surfaces. Each RIS consists of **256 unit cells** and operates in the 5 GHz WLAN frequency range. RIS provides **1-bit phase control**, corresponding to an approximate 180° phase step between its two element states, and supports linear polarization. The unit cells have a size of 20 mm × 13 mm, and a full reconfiguration of all 256 elements via USB is specified with an update time of approximately 17 ms.

For the demonstrator, the RIS are not controlled directly via long USB cables from the measurement laptop. Instead, each RIS is connected via USB to a Raspberry Pi mounted close to the surface, following the control architecture introduced in **D5.1**. The Raspberry Pi powers and controls the RIS and is itself supplied via Power over Ethernet (PoE). A lightweight control service developed at **TUBS** runs on the Raspberry Pi, automatically detects connected RIS, and announces them in the local network via Zeroconf. From the perspective of the laptop application, the network transport is transparent: the RIS can be discovered and configured similarly to a locally attached USB device.

The RF measurement is performed with two USRP B210 devices, one acting as transmitter and one as receiver. This makes the demonstrator portable and independent of the large CF-MIMO rack infrastructure. While the CF-MIMO prototype focuses on active multi-antenna beamforming, the RIS demonstrator focuses on programmable passive reflections and cascaded RIS-assisted propagation.

4.1.2.1 Final deployment and configuration

The final RIS demonstrator uses the cascaded RIS scenario already proposed in **D5.1**, but with an updated geometry. The RF path is arranged as a three-hop link: The final RIS demonstrator uses the cascaded RIS scenario already proposed in **D5.1**, but with an updated laboratory geometry. The RF path is arranged as a three-hop link: TX→RIS1→RIS2→RX.

In the current setup, TX, RIS1, RIS2, and RX form an around-the-corner triangular geometry. A metallic wall is placed between TX and RX, suppressing the direct line-of-sight path and forcing the received signal to rely primarily on the RIS-assisted reflection path. This makes the setup well suited for demonstrating whether two RIS can be configured in series to guide a signal around an obstacle.

All relevant propagation directions lie approximately in a common horizontal plane. The steering task is therefore treated as a two-dimensional horizontal-angle problem. In the evaluated configuration, RIS1 is configured for the fixed geometry between the transmitter and RIS2. Its reflection pattern redirects the incident signal from TX towards RIS2. RIS2 is configured with the correct incoming direction from RIS1, while its outgoing reflection direction is swept horizontally from -90° to +90°. For every configured output angle of RIS2, the received signal strength at the receiver is recorded.

The RIS configurations are generated from the geometric incident and outgoing directions of each surface. For a given incident direction and desired reflection direction, the control software computes the required phase profile across the 16×16 RIS aperture. Since the employed RIS provides 1-bit phase control, the continuous phase profile is quantized to the two available element states, corresponding approximately to 0° and 180°. The resulting 256-bit pattern is then transmitted to the RIS controller. In the measurement

⁴ NovoFlect Custom RIS, <https://novoflect.de/en/custom-ris-2/>

campaign, the pattern of RIS1 remains fixed, because its incident and outgoing directions are fixed by the TX → RIS1 → RIS2 geometry. For RIS2, the incident direction from RIS1 is kept fixed, while the desired outgoing direction is varied over the horizontal sweep angle. For every sweep point, a new 1-bit RIS pattern is generated and applied before the received signal strength is evaluated.

The measurement was performed at approximately 5.89 GHz.

This procedure allows the angular response of the cascaded RIS link to be measured. Since RIS1 remains fixed and RIS2 performs the sweep, changes in received signal strength can be attributed primarily to the outgoing steering direction of RIS2. The setup therefore provides a compact validation of serial RIS operation and passive beam steering.




Figure 4-15 Final deployment of the RIS-based demonstrator. The RF path follows a cascaded TX → RIS1 → RIS2 → RX geometry. A metallic wall blocks or strongly attenuates the direct TX-to-RX path, so that the received signal is dominated by the two-RIS reflection path. RIS1 is configured for the fixed TX-to-RIS2 path, while RIS2 sweeps its horizontal outgoing reflection direction from -90° to $+90^\circ$

4.1.2.2 Results and KPIs

The main KPI of the RIS demonstrator is the received signal strength as a function of the configured output angle of RIS2. This KPI differs from the broader SINR/CQI/data-rate targets outlined in D5.1, because the present demonstration focuses specifically on validating the cascaded RIS steering behaviour in a controlled laboratory geometry. Received signal strength is therefore the most direct and robust indicator for the angular response of the two-RIS path.

Figure 4-X shows the received signal strength over the horizontal RIS2 steering sweep from -90° to $+90^\circ$.




Figure 4-16 Received signal strength of the cascaded TX → RIS1 → RIS2 → RX link as a function of the configured horizontal output angle of RIS2. RIS1 is kept fixed for the TX-to-RIS2 path, while RIS2 sweeps from -90° to $+90^\circ$. A moving average is applied to the received measurements because the SDR measurement rate is higher than the RIS reconfiguration rate

During the sweep, the SDR receiver records signal-strength samples faster than the RIS can be reconfigured. For the plot, a moving average is therefore applied to the received measurements. This suppresses fast sample-to-sample variation and makes the angular response of the RIS sweep visible.

The measurement shows a clear angular dependency of the received signal strength. A pronounced maximum occurs when the configured output angle of RIS2 matches the geometric direction towards the receiver. Away from this main lobe, the received power decreases, while secondary maxima appear as sidelobes of the finite RIS aperture. The curve also contains distinct minima, corresponding to angular configurations where destructive interference or poor alignment between the RIS pattern and the receiver direction reduces the received signal strength.

The observed response is consistent with the expected behaviour of a finite, 1-bit phase-quantized reflecting surface. The main lobe demonstrates that the RIS can steer the reflected signal towards a desired direction.

Overall, the demonstrator validates the feasibility of the cascaded RIS-assisted link envisioned in [D5.1](#). By configuring one RIS for a fixed incoming/outgoing geometry and sweeping the second RIS, the experiment shows that cascaded RIS can be used to guide a signal over multiple passive reflections around an obstacle. This provides a compact proof of concept for RIS-assisted coverage extension and controlled smart-radio-environment experiments within [PoC#2](#).

Table 4-4 RIS-based PoC Demonstrator

Test 4.4	Experimental validation of a cascaded RIS-assisted beam steering demonstrator
Phase	Validation and performance evaluation
Description	A portable RIS-based proof-of-concept demonstrator was developed to experimentally validate programmable passive beam steering using cascaded reconfigurable intelligent surfaces (RIS). The demonstrator employs two NovoFlect RIS-5Mk1-16×16 panels, each comprising 256 programmable unit cells with 1-bit phase control, together with two USRP B210 software-defined radios. The RIS are controlled via Raspberry Pi units connected through Power-over-Ethernet, enabling flexible network-based configuration independent of the measurement computer. The experimental setup realizes a cascaded propagation

	path (TX → RIS1 → RIS2 → RX), where a metallic obstacle suppresses the direct line-of-sight link and forces communication via two successive passive reflections. During the measurements, the first RIS remains fixed while the second RIS performs a horizontal beam-steering sweep. The received signal strength is recorded as a function of the configured steering angle and compared with the expected beam pattern, demonstrating programmable passive wavefront manipulation and controlled signal routing around an obstacle.
KPIs	• Successful implementation of a portable RIS-based demonstrator • Reliable software control of two cascaded RIS panels via Ethernet-connected Raspberry Pi controllers • Experimental validation of cascaded RIS-assisted transmission (TX → RIS1 → RIS2 → RX) • Programmable passive beam steering using 1-bit phase-quantized RIS • Measurement of the angular response of the cascaded RIS link • Clear main lobe observed at the intended steering direction with expected sidelobe behaviour • Experimental validation of passive coverage extension and smart radio environment concepts.
PoC	PoC#2 – RIS-based beam steering demonstrator

4.2 Implementation at UC

A second CF-MIMO USRP-based platform was developed by the **UC** team throughout the **6G-SENSES** project to complement the results from the **TUBS** platform. This small-scale platform has up to six AP antennas and 2 UE antennas and considers only the PHY-layer implementation of TDD CF-MIMO scenarios. The main objective of this second platform was to implement and evaluate OTA multi-antenna multi-carrier reciprocity calibration algorithms and integrate them into OpenAirInterface (OAI). The lower hardware complexity of this platform compared to the implementation at TUBS, provides greater flexibility and speed in the development and evaluation of distributed beamforming algorithms, which has made it possible, for example, to evaluate KPIs that complement those considered in Section 4.1. In particular, in addition to reciprocity calibration algorithms, it was used to evaluate the advantages of CF-MIMO versus standard small-cell implementations and multi-user scenarios.

4.2.1 Final deployment and configuration

4.2.1.1 OAI integration of reciprocity calibration

An OAI-based RU/gNB was implemented using 5G NR numerology, while relying on the hardware available at UC (USRP N310, USRP N321, and an OctoClock).

It was verified that OAI manages all RU antennas through a single UHD transmission streamer and a single UHD reception streamer. As a result, the native execution of the reciprocity calibration procedure is not possible, since the calibration process requires one RU antenna to transmit while the remaining antennas receive, and vice versa. To overcome this limitation, the **UC** RU incorporates an additional antenna (implemented using an extra USRP) whose sole function is to serve as a reference antenna for the OTA reciprocity calibration procedure.

The calibration procedure takes as input the OAI structures `openair0_cfg_t` and `usrp_state_t`, together with a pointer to a vector used to return the calibration results. During execution, a `uhd::usrp::multi_usrp` object is instantiated for the USRP dedicated to calibration, along with its corresponding RX and TX streamers. The feasibility of this approach has been validated through experimental testing at **UC**.

The Argos OTA reciprocity calibration algorithm [6] [19], extended to support multicarrier communications, has been successfully integrated into and validated on an OAI RU/gNB deployed at **UC**.

The **UC** experimental setup, shown in Figure 4-17, consists of a USRP N310 operating as a four-antenna gNB and a USRP N321 acting as the external calibration device. The gNB runs OAI version 2025.w32 and is configured with the following parameters:

- Number of antennas: 4.
- Number of resource blocks: 51.
- Subcarrier spacing: 30 kHz.
- Center frequency: 3619.20 MHz.
- Sample rate: 30.72 Msps.




Figure 4-17 Block diagram of UC setup for reciprocity calibration integration in OAI RU gNB

UC has developed dedicated C++ software to control the RF hardware, including both the gNB USRP and the additional calibration USRP, through the Ettus UHD driver. This software implements the reciprocity calibration algorithm and has been fully integrated into the standard OAI gNB C++ codebase. The integrated solution has been successfully tested on the UC OAI implementation.

Lessons Learned. The current OAI gNB-RU implementation, which relies on the `multi_usrp` object, requires an external device to perform OTA reciprocity calibration. This requirement could potentially be removed by adopting an RFNoC-capable USRP for the RU gNB and introducing substantial modifications to the OAI framework, particularly in its management of UHD streamers.

4.2.1.2 CF-MIMO small-scale demonstrator

The CF-MIMO demonstrator developed at **UC** is based on low-cost SDR hardware, specifically Ettus USRP B210 devices for signal transmission and reception, and an Ettus OctoClock-G CDA-2990 for synchronization. The synchronization unit provides both carrier-frequency and baseband clock references, as well as timing synchronization.

The setup is managed by a central control PC running MATLAB under Windows and using the GTEC Testbed Interface Software (GTIS), developed by the University of A Coruña (**UDC**) in collaboration with **UC**, for the remote control of multiple USRP devices. The control PC communicates via SSH/TCP-IP with several Ubuntu-based processing nodes, each connected through USB 3.0 to one or more USRP B210 devices using the Ettus UHD driver.

Figure 4-18 shows a schematic diagram of the UC CF-MIMO small-scale demonstrator, which comprises three APs with two antennas each (two USRP B210 units per AP) and two single-antenna UEs. A photograph of the demonstrator is shown in Figure 4-19. The figure highlights the three APs, each equipped with two antennas, together with the single-antenna UEs. A two-axis positioning table is also visible, enabling controlled UE movement within a 1 m × 1 m area with a positioning resolution of 1 mm. Additionally, to emulate a non-line-of-sight (NLoS) propagation environment with rich multipath conditions, several obstacles and reflective panels were placed in front of the AP antennas during the experimental campaign.

The demonstrator has been used to evaluate both the feasibility and the necessity of OTA reciprocity calibration in a practical multicarrier (5G NR-like waveform) CF-MIMO deployment based on commercial off-the-shelf transceivers and low-cost synchronization hardware.

The physical-layer experiments were conducted using 5G NR-inspired frames with the following parameters:

- Number of resource blocks: 52
- Subcarrier spacing: 15 kHz
- Number of OFDM symbols per slot: 14
- Sample rate: 15.36 Msps
- Centre frequency: 3000.00 MHz
- Modulation schemes: QPSK and 16-QAM




Figure 4-18 UC CF-MIMO small-scale demonstrator diagram




Figure 4-19 UC CF-MIMO small-scale demonstrator setup

4.2.2 Results and KPIs

4.2.2.1 OAI integration of reciprocity calibration

Several experimental test and measurements have been carried out to validate the OAI integration and implementation of the OTA reciprocity calibration procedure. Figure 4-20 presents a screenshot of the operational gNB displaying the reciprocity calibration log messages. These logs include:

- Reciprocity calibration setup: initialization and configuration of the calibration USRP.
- Reciprocity calibration procedure:
 - Stage 1: The calibration USRP transmits five OFDM symbols, which are received by the four RF chains of the gNB USRP.
 - Stage 2: Each gNB antenna transmits five OFDM symbols using frequency-domain multiplexed pilots, enabling the simultaneous reciprocity calibration of four or more antennas.
- Phase estimation: Linear regression-based estimation of the phase drift as a function of frequency (subcarrier index) for each gNB RF chain. The resulting phase characteristics are shown in Figure 4-21.

The reciprocity calibration coefficients computed for each subcarrier and gNB antenna are then provided to OAI, enabling the exploitation of TDD channel reciprocity and, consequently, the use of uplink channel estimates for downlink beamforming.

```

16 de abr 18:04
gtas@corcel: ~/openairinterface5g-2025.w32/cmake_targets/ran_build/build

[HW] Device timestamp: 1.711451...
[HW] [CAL] RECIPROCITY CALIBRATION SETUP
[HW] [CAL] Number of physical antennas: 4
[HW] [CAL] Center Frequency: 3619.200000 MHz
[HW] [CAL] Sample Rate: 30.720000 Msps
[HW] [CAL] Bandwidth: 100.000000 MHz
[HW] [CAL] Calibration USRP: multi_usrp creation...
[INFO] [MPMD] Initializing 1 device(s) in parallel with args: mgmt_addr=192.168.20.3,type=n3xx,product=n320,serial=32C0C0
clock_source=external,time_source=external,master_clock_rate=245.76e6,subdev=A:0
[INFO] [MPM.PeriphManager] init() called with device args 'clock_source=external,fpga=WX,master_clock_rate=245.76e6,mgmt_
time_source=external'.
[INFO] [MPM.Rhodium-0] init() called with args 'clock_source=external,fpga=WX,master_clock_rate=245.76e6,mgmt_addr=192.16
=external'.
[INFO] [MPM.Rhodium-1] init() called with args 'clock_source=external,fpga=WX,master_clock_rate=245.76e6,mgmt_addr=192.16
=external'.
[HW] [CAL] ...multiusrp created
[HW] [CAL] Calibration USRP master clock rate : 245.760000 MHz
[HW] [CAL] Calibration USRP carrier frequency : 3619.200000 MHz
[HW] [CAL] Calibration USRP: streamers creation...
[WARNING] [0/Radio#0] Attempting to set tick rate to 0. Skipping.
[HW] [CAL] ...streams created
[HW] [CAL] CAL USRP TX antenna 1 gain = 60 dB
[HW] [CAL] CAL USRP RX antenna 1 gain = 40 dB
[HW] [CAL] Calibration USRP clock set to 1776355361.000000 sec
[HW] [CAL] OAI USRP and Calibration USRP PPS's synchronized.
[HW] [CAL] END OF RECIPROCITY CALIBRATION SETUP
[HW] [CAL] RECIPROCITY CALIBRATION PROCEDURE
[HW] [CAL] STAGE 1
[HW] [CAL] CAL USRP transmits and OAI USRP receives at t = 1776355362.249694
[HW] [CAL] OAI USRP received 8000 samples for each antenna
[HW] [CAL] STAGE 2
[HW] [CAL] OAI USRP transmits and CAL USRP receives at t = 1776355362.350114
[HW] [CAL] CAL USRP received 8000 samples
[HW] [CAL] ReceivedSignalStage1:
[HW] [CAL] Antenna 0: peak value = 3434 (Range: 0 - 32768)
[HW] [CAL] Antenna 1: peak value = 2386 (Range: 0 - 32768)
[HW] [CAL] Antenna 2: peak value = 6838 (Range: 0 - 32768)
[HW] [CAL] Antenna 3: peak value = 2754 (Range: 0 - 32768)
[HW] [CAL] ReceivedSignalStage2:
[HW] [CAL] peak value = 9433 (Range: 0 - 32768)
[HW] [CAL] DONE: 1024 phases of 4 antennas wrote to memory and file ReciprocityCalTables.csv
[HW] [CAL] END OF RECIPROCITY CALIBRATION PROCEDURE
[HW] [CAL] Calibration USRP Released
gtas@corcel:~/openairinterface5g-2025.w32/cmake_targets/ran_build/build$

```

Figure 4-20 OAI gNB executing the reciprocity calibration procedure at UC lab




Figure 4-21 Estimated reciprocity phase drift versus OFDM subcarrier for each gNB antenna

Table 4-4 PoC#2 OAI integration of reciprocity calibration

Test 4.5	OAI integration of reciprocity calibration
Phase	Validation and performance evaluation
Description	Implementation of OTA reciprocity calibration algorithms in OAI gNB: Real-time implementation of reciprocity calibration on a USRP N310 operating as a four-antenna OAI gNB and a USRP N321 acting as the external calibration device.
KPIs	Success criteria: Feasibility of OTA calibration algorithms integration into OAI.
PoC	PoC#2, OAI integration of reciprocity calibration

4.2.2.2 Single user CF-MIMO

The first experimental campaign focused on a single-user scenario comprising three APs with two antennas each and one single-antenna UE positioned at 441 different locations on a 5 cm × 5 cm grid covering the 1 m × 1 m positioning area.

Both the Argos [19] and BeamSync [21] procedures, extended to multicarrier communications, were executed at the beginning of the experimental campaign to estimate the reciprocity calibration coefficients for each AP antenna and subcarrier, which were subsequently used for downlink channel compensation.

At each location, during the uplink phase, the UE transmits one slot. The APs receive the signal, and the central processing unit estimates the uplink channel. During the downlink phase, the APs transmit ten slots, each corresponding to a different transmission strategy:

- Slots 1–6 (**SISO₁–SISO₆**): Transmission using only one of the six available antennas.
- Slot 7 (**MRT Uncal.**): Coordinated beamforming using all six antennas. Maximum Ratio Transmission (MRT) is applied on a per-subcarrier basis using the uplink channel estimates directly as downlink channel estimates, without reciprocity calibration.
- Slot 8 (**MRT Argos**): Same as MRT Uncal., but incorporating reciprocity calibration obtained through the Argos procedure.
- Slot 9 (**MRT BeamSync**): Same as MRT Uncal., but using reciprocity calibration obtained through the BeamSync procedure.
- Slot 10 (**MRT SmallCell**): Only the AP providing the highest average SNDR across the 624 subcarriers serves the UE. MRT is performed using its two antennas, with Argos-based reciprocity calibration.

Figure 4-22 shows the cumulative distribution function (CDF) of the SNDR for all strategies. Each curve is based on 1,373,715 SNDR measurements, corresponding to five repetitions at 441 UE locations across 623 subcarriers.




Figure 4-22 Single user downlink performance

The results lead to several important observations.

First, without reciprocity calibration (MRT Uncal.), coordinated beamforming provides essentially no benefit and performs similarly to conventional SISO transmission. This demonstrates that reciprocity calibration is indispensable for exploiting channel reciprocity in TDD-based CF-MIMO systems.

Second, the Argos and BeamSync techniques (MRT Argos and MRT BeamSync) provide nearly identical performance. This can be attributed to the relatively benign propagation conditions in the laboratory environment and the close proximity of the APs. Under more challenging propagation conditions, BeamSync is expected to outperform Argos.

In this setup, only 10 subcarriers from two consecutive OFDM symbols (20 pilots in total) were used to estimate the calibration coefficients for all 624 subcarriers by exploiting the approximately linear relationship between calibration phase and frequency. Experimental results show that the resulting CF-MIMO MRT performance is virtually identical to that obtained using up to 6,240 pilots distributed across all subcarriers of ten consecutive OFDM symbols.

Third, calibrated coordinated beamforming (CF-MIMO) significantly outperforms the conventional Small Cell approach as well as the uncalibrated solution. MRT Argos and MRT BeamSync achieve an average SNDR of 19.5 dB, whereas MRT SmallCell reaches 16 dB and MRT Uncal. achieves only 11.4 dB. This corresponds to:

- A 3.5 dB array gain over the Small Cell approach.
- More than 8.5 dB gain over uncalibrated MRT.

Furthermore, the slopes of the curves in Figure 4-22 indicate that reciprocity calibration is essential for exploiting the available spatial diversity gain.

Finally, for an outage BER of 10^{-3} and QPSK signaling (requiring approximately 9.8 dB SNR), the outage probabilities are:

- 4×10^{-4} for MRT with reciprocity calibration (Argos or BeamSync).
- 4×10^{-2} for MRT SmallCell.

- 5×10^{-1} for MRT without reciprocity calibration.

These results indicate that the Small Cell strategy requires approximately 4.4 dB more transmit power than the calibrated CF solution to achieve the same outage target. Moreover, the uncalibrated Cell-Free approach requires approximately 30 dB more power than the calibrated Argos and BeamSync solutions.

When analysing downlink spatial coverage, Figure 4-23 illustrates the averaged UE SNDR across all subcarriers over the measurement grid with a resolution of 5 cm $\times$ 5 cm. MRT-based CF-MIMO provides consistently better coverage, higher SNDR values, and lower spatial variability than the MRT Small Cell approach. This observation is further confirmed in Figure 4-24, which shows the corresponding SNDR histograms.

It is worth noting that constraining service provision to a single AP, as in the Small Cell approach, leads to a performance loss that becomes even more pronounced in multicarrier systems. In highly frequency-selective channels, the two best antennas are unlikely to be the same across all 624 subcarriers, nor are they necessarily associated with a single AP.




Figure 4-23 Spatial distribution of the average SNDR in the single-user scenario. Left: Cell-Free MRT. Right: Small Cell MRT




Figure 4-24 Histogram of location-dependent SNDR for Cell-Free MRT and Small Cell MRT

Lesson Learned. Reciprocity calibration of the APs is essential for exploiting channel reciprocity in TDD-based CF-MIMO systems. Without reciprocity calibration, downlink channel state information must be obtained through feedback mechanisms.

Lesson Learned. Experimental results show that the reciprocity calibration phase exhibits a linear dependence on frequency. This behavior simplifies phase estimation and provides greater flexibility in the allocation of calibration pilots in OFDM deployments. Furthermore, the calibration phases remain stable over extended periods (up to several hours) in laboratory environments with stable temperature conditions after hardware warm-up. Consequently, the calibration procedure can be scheduled at intervals of several minutes or even hours.

Table 4-5 PoC#2, OTA reciprocity calibration in UC CF-MIMO platform

Test 4.5	Single user scenario
Phase	Validation and performance evaluation
Description	Implementation of reciprocity calibration algorithms in a TDD CF-MIMO system: Real-time implementation of reciprocity calibration algorithms on a CF-MIMO platform based on USRP B210 nodes. The platform consists of 3 APs, with 2 antennas each, beamforming 5G-NR frames coherently and one UE. Both intra-AP calibration algorithms (using one of the antennas of each AP as a reference) and inter-AP calibration algorithms (using one of the APs as a reference) are implemented. The OTA calibration algorithms are based on bidirectional pilot transmissions between the antennas of an AP or between the APs.
KPIs	Success criteria: • Feasibility of OTA calibration algorithms demonstrated. • Improved SNR and improved coverage in the UE service area with respect to DL transmissions without calibration. Measurable requirements: • Phase synchronization error: no SNDR degradation even with only 10 pilots. • Average SNDR: ○ Measured improvement of 8.5 dB wrt to uncalibrated systems. ○ Measured improvement of 3.5 dB of CF-MIMO over Small Cell. • Outage probability: ○ Measured 100 times outage probability improvement of CF-MIMO over Small Cell for uncoded BER of $1e-3$. ○ Measured 4.4 dB efficiency increase of CF-MIMO over Small Cell for an outage probability of $4e-4$ (uncoded BER of $1e-3$). ○ Measured 30 dB efficiency increase of CF-MIMO with reciprocity calibration over CF-MIMO without reciprocity calibration for an outage probability of $4e-4$ (uncoded BER $1e-3$). • SNR maps: improved uniform coverage in the measurement area.
PoC	PoC#2, UC CF-MIMO small-scale testbed

4.2.2.3 Multiuser CF-MIMO

The second experimental campaign, devoted to multi-user scenarios, deploys three APs with two antennas each and two single-antenna UEs. One UE is placed at 121 different positions on a $10\text{ cm} \times 10\text{ cm}$ grid covering a $1\text{ m} \times 1\text{ m}$ positioning table, while the other UE remains at a fixed location.

Argos procedure, extended to multicarrier communications, were executed at the beginning of the experiment to estimate the reciprocity calibration coefficients for each AP antenna and subcarrier, which were subsequently used for downlink channel compensation.

At each location, during the uplink phase, the UEs transmit orthogonal pilots for uplink channel estimation and non-orthogonal data symbols (i.e., using the same resource blocks simultaneously). The CPU estimates

the uplink channels and performs coordinated maximum-ratio combining (MRC) to recover the transmitted data.

During the downlink phase, the CPU corrects the previously estimated uplink channels using the reciprocity calibration coefficients obtained through either the Argos or BeamSync procedure and uses the resulting channels as downlink channel estimates. These estimates are then employed for simultaneous coordinated transmission using minimum mean square error (MMSE) precoding on a per-subcarrier basis across the six available antennas of the three APs. Both UEs receive the transmitted signals and perform conventional linear processing of 5G NR-style slots.

Figure 4-25 compares the downlink signal-to-interference-plus-noise ratio (SINR) observed at both UEs with and without reciprocity calibration. Each curve is based on 75,504 SINR measurements, corresponding to 121 UE1 positions across 624 subcarriers. The results show a gain of between 10 dB and 20 dB for calibrated MMSE compared with uncalibrated MMSE, depending on the position of UE1. This suggests that reciprocity calibration has a greater impact in multi-user scenarios than in single-user scenarios.




Figure 4-25 Reciprocity calibration impact on SINR in multiuser downlink. Left: User 1. Right: User 2

Finally, an additional set of measurements was carried out to compare SINR performance under different AP power-allocation strategies. It should be noted that MMSE precoding is applied independently to each active downlink subcarrier and that the resulting OFDM signals transmitted by the six antennas may have different power levels. Therefore, two strategies were evaluated:

- Strategy 1 (**MMSE_global**): all six transmitted signals are scaled by the same factor so that the transmit power of every AP remains below or equal to the maximum power allowed per AP.
- Strategy 2 (**MMSE_perAP**): the pair of signals transmitted by each AP is independently scaled so that the transmit power of every AP is equal to the maximum power allowed per AP

It should be noted that MMSE_perAP always transmits at the maximum power permitted by the hardware or regulatory constraints, whereas MMSE_global may transmit at a lower power level depending on the precoding vectors, i.e., on the channel realization. As shown in Figure 4-26, global power normalization

outperforms per-AP normalization by more than 3 dB despite using less transmit power. This is because applying different scaling factors to the downlink signals transmitted by different APs distorts the intended MMSE precoding and increases inter-user interference, even though it provides higher received power at each UE.




Figure 4-26 Power allocation impact on SINR in multiuser downlink: global normalization versus per AP normalization. Left: User 1. Right: User 2

An example of the self-induced inter-user interference resulting from the application of different scaling factors at each AP is shown in Figure 4-11, which depicts the received constellations of both users for a representative UE location under the two power-allocation strategies. Under MMSE_perAP normalization, the constellation of UE1 exhibits clear interference from UE2, appearing as a smaller, noisy, and rotated 16-QAM constellation superimposed on the desired symbols. The same phenomenon can also be observed in the constellation of UE2, although to a lesser extent.




Figure 4-27 Power allocation impact on multiuser downlink constellations. Up: User 1. Down: User 2. Left: global normalization. Right: per AP normalization

Table 4-6 PoC#2 -Experimental validation of multiuser in UC CF-MIMO platform

Test 4.6	Multiuser scenario
Phase	Validation and performance evaluation
Description	Implementation of reciprocity calibration algorithms in a TDD CF-MIMO multiuser system: Real-time implementation of reciprocity calibration algorithms on a CF-MIMO platform based on USRP B210 nodes. The platform consists of 3 APs, with 2 antennas each, beamforming 5G-NR frames coherently and two UEs. Both intra-AP calibration algorithms (using one of the antennas of each AP as a reference) and inter-AP calibration algorithms (using one of the APs as a reference) are implemented. The OTA calibration algorithms are based on bidirectional pilot transmissions between the antennas of an AP or between the APs.
KPIs	Success criteria: • Feasibility of OTA calibration algorithms demonstrated. • Improved SINR and improved coverage in the UEs service area with respect to DL transmissions without calibration. Measurable requirements: • Phase synchronization error: no SNDR degradation even with only 10 pilots. • Average SINR: ○ Measured improvement of at least 10 dB wrt. to uncalibrated systems.
PoC	PoC#2, UC CF-MIMO small-scale testbed

5 PoC#3 – Network Digital Twin (NDT) via RAN Sensing

PoC#3 leverages the sensing capabilities of the RAN to optimise network performance and energy efficiency with the support of an NDT. The NDT is being developed based on network planning information and network performance/telemetry data, and is implemented as an xApp associated with the O-RAN RICs. The demonstrator is currently integrating three complementary technologies—Sub-6 GHz, mmWave, and 5G NR—as sensing information sources, which are being fed into the O-RAN framework. Within this framework, the processes of network planning enabling dynamic re-configuration decisions that improve network performance and energy efficiency. These assessment and optimisation functions rely on network performance datasets that are continuously being collected and processed by the NDT application. A general implementation architecture is presented in Figure 5-1. This architecture, focused on the integration and implementation, is an extract of the architecture defined for **PoC#3** in deliverable **D5.1** [1].




Figure 5-1 PoC#3 implementation Architecture

5.1 Intelligent Plane implementation

This section reports the Intelligent Plane as realised for **PoC#3** in the final experimental phase. It connects three prior strands of work: the IP design and implementation framework in **D4.4** [11], the **PoC#3** architecture, interfaces and test planning in **D5.1** [1], and the integration status documented in **D5.2** [2]. Where **D5.1** specified what to build and test, and **D5.2** described how the components were integrated into the platforms, this deliverable (**D5.3**) focuses on the implemented IP baseline used in the **UC** Santander closed-loop demonstration, and on how that baseline supports Inference, NDT lookup and RAN actuation in the live trial.

Within the **6G-SENSES** architecture, the Intelligent Plane is not a new standalone block. As defined in **D4.4**, it is the functional layer that links sensing, communication, computing and orchestration with AI/ML-driven decision-making, distributed across management, control and data-processing domains. It is both an

architectural concept and an operational framework: the means by which models are trained, deployed, monitored and connected to real network control actions. In O-RAN terms, those functions map onto the existing SMO / Non-RT RIC (analytics, model lifecycle, policy, longer-horizon optimisation) and Near-RT RIC (xApps for near-real-time inference and control), rather than introducing a parallel management plane. dRAX provides the RIC/SMO host environment—live telemetry buses, historical databases, and the execution context for rApps/xApps—while MLRun (or equivalent MLOps tooling) supplies the lifecycle/CI-CD part of the IP; it is not the Intelligent Plane by itself.

For **PoC#3** — Network Digital Twin (NDT) via RAN Sensing — **D5.1** and **D5.2** organise the system around two application roles on top of this IP: a Modelling path (model management, monitoring, serving and model database) and an Optimisation Control path (user targets → objectives such as energy efficiency, coverage or throughput → select, evaluate and apply configurations). Two environments participate in the loop: the real network; ACC Open 5G SA NPN / dRAX at **UC**, with sensing enrichment, and the NDT / simulator sandbox (ns-3 5G-LENA and Sionna RT in the partner design). Operation follows the three interrelated cycles already introduced in **D4.4** [11] and **D5.2** [2]— Model Creation, Model Actualisation and Validation, and Model Implementation — so that data-driven models are created from live telemetry, refined against NDT and field measurements, and only then applied to the RAN under operator-driven objectives.

Accelleran's contribution in this chain is the Open 5G SNPN and dRAX Intelligent Plane baseline (pre-trialled with **OTE** and finalised at **UC**), including northbound exposure for NDT/SMO integration and the Near-RT path that carries inference and control towards the live cells. The subsections that follow describe how that IP baseline was implemented for the **PoC#3** demonstration; experimental setups, Inference, NDT and KPI results are reported in §5.2 and §5.3.

5.1.1 Lab setup under the UC server

The **PoC#3** demonstration at **UC** is built on Accelleran's Open 5G Standalone Non-Public Network and dRAX RIC/SMO stack, extended with the 6G-SENSES Intelligent Plane baseline and the UC-hosted applications. From an operator's point of view the stack splits into four layers: the radio access path (RU, DU, and CU under the Cell Wrapper / gNB configuration), the RIC data and control plane (dRAX), the three xApp pods that implement Modelling, Intelligent Plane lifecycle, and Optimisation Control, and a dedicated demo-control host that orchestrates everything through a single HTTP service. A schematic representation of the system is presented in Figure 5-2.




Figure 5-2 Schematic organization of the PoC#3 implementation

On the RAN side, the live cell used by the closed loop is identified in configuration as cell-1, with the Benetel radio unit associated as ru-cell-1. Configured transmit power for the Control application path is read and written through the dRAX Configuration API (field `maximum_eirp_dbm` on the gNB RAN cell object), so that actuation goes through the Cell Wrapper rather than bypassing it with direct RU SSH.

On the RIC side, dRAX provides the near-real-time telemetry bus and the configuration API used by the xApps. The Modelling xApp consumes UE measurement traffic from the Kafka topic used for RIC raw UE measurements and writes rolling KPI files (including 30-second SSB averages of RSRP, RSRQ and SINR). Those files are the live input to both the Dataset lab and the Control closed loop. The Optimisation Control xApp hosts the Inference, the NDT, and the Control pipeline and, the `cell_tx_control` client that talks to drax-dashboard on the cluster. The Intelligent Plane xApp is registered for lifecycle (start/stop/status/logs) from the dashboard; its deeper MLOps / training stack is out of scope for the Control application pages in this monorepo, consistent with the **D4.4** split between IP capabilities and PoC application logic.

As can be seen in Figure 5-3, physically the demo-control webservice, as well as all the 3GPP and O-RAN services, run on a **UC** lab control host (bare-metal in the 192.168.112.0/24 lab LAN). located in a server room at the UC. In turn, it is connected to the Benetel O-RU, located in a lab, through a PTP switch to provide synchronization. SSH from that host reaches a bastion (UCantabria_server) and then the dRAX/Kubernetes node (alias cu), where each xApp is exposed on a dedicated NodePort-style SSH port: Modelling, Intelligent Plane, and Optimisation Control. Operators and browsers never connect to the RU or to the xApp pods directly for the demo UX: all orchestration is mediated by demo-control over SSH BatchMode. The demo webservice is a Python HTTP service that binds on port 30880 and serves both static operator pages and a JSON API. It can be started in the foreground or as a background process via project scripts, and optionally as a user-level systemd unit so that the service survives logout. Beside the HTTP process, the same host runs long-lived orchestrator helpers and a roughly 30-second Tx / UE history sampler used by the Control charts. Configuration is file-based. Remote actions always go through an SSH runner with key-based BatchMode authentication; health and lifecycle actions map to a common remote script contract on each xApp PVC (start/stop/restart/status/logs). Deploying updated xApp code is done by syncing the Modelling and Optimisation Control trees onto the corresponding PVCs on the cu node and restarting the services; changing the webservice or static UI requires a demo-control restart and a hard browser refresh. In short: demo-control is the single northbound entry point for the PoC#3 lab demo. It does not replace dRAX or the Cell Wrapper; it orchestrates the three xApps, exposes operator UIs, and keeps browsers away from SSH and RU credentials.

The main **PoC#3** demo-control page, shown in Figure 5-4, is the operator dashboard that hosts the three xApps as lifecycle cards on one screen. From this page, a user can start, stop, restart, and query status and logs for Modelling, Intelligent Plane, and Optimisation Control, and can validate SSH reachability to the UC hosts before a trial. The page does not embed separate xApp web UIs; it is the common control surface for process lifecycle, while specialised workflows live on linked pages served by the same webservice. The Modelling card is the richest of the three: besides lifecycle actions it shows live connected-UE information and the per-UE SSB table (30-second averages of RSRP, RSRQ and SINR) read from the Modelling xApp KPI files over SSH. That live view is what operators use to confirm that Kafka UE measurements are flowing and that the radio path is healthy before starting Dataset lab collection or a Control cycle. The Intelligent Plane and Optimisation Control cards expose the same lifecycle buttons; deeper Intelligent Plane MLOps remains out of scope on this dashboard, and the full Inference / NDT / Tx closed loop is operated from the Control application pages rather than from these cards.




Figure 5-3 Location and physical installation of the Physical Installation of ACC O-RAN BS @ UC testbed

PoC #3 DEMO CONTROL
Menu

Other apps — lifecycle

Remote scripts/*-service.sh via API. Additional registered apps (if any) appear below.

Validate SSH (all apps)

Model lifecycle (PoC#3)

Model creation
Actualisation

Implementation

No additional registered apps beyond the three xApps.

start
stop
restart
status
logs

Modelling

Connected UEs 1

Per-UE SSB (30s average) i

NR cell ID	UE ID	SSB RSRQ (dB)	SSB SINR (dB)	SSB RSRP (dBm)
1	4	-10.3	23.2	-71.7

1 UE · 2026-07-23T14:29:40Z

Intelligent Plane
start
stop
restart
status
logs

OptimizationCRTL
start
stop
restart
status
logs

Remote status & logs i Clear

Figure 5-4 Demo control website

5.2 Experimental validation of the PoC

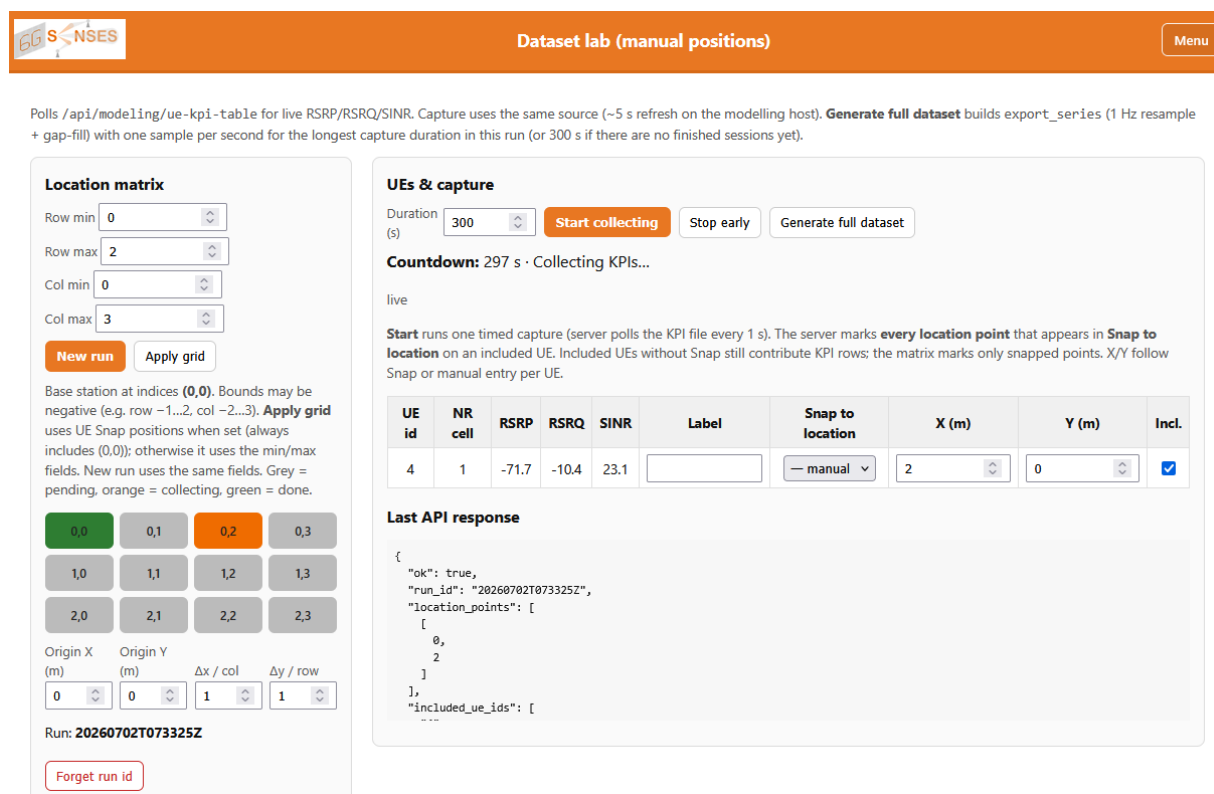
5.2.1 Final deployment and configuration

5.2.1.1 Setup 1 Data generation

The Dataset lab is a dedicated operator page on the same demo-control webservice (port 30880) used for manual, location-labelled KPI capture in the UC Santander lab. Its purpose is to build modelling datasets while a person places UEs (or moves with a UE) at known grid points, so that live SSB metrics from the Modelling xApp can be stored together with ground-truth location. The service reads the same live UE KPI table as the main dashboard which ultimately comes from Modelling's latest_ue_kpi_30s.

The UI is a two-column layout as described in Figure 5-5. On the left, a Location matrix defines an inclusive integer grid of row/column indices (bounds may be negative). Index (0,0) is treated as the base-station reference and is always kept inside the matrix. The operator creates a New run (which allocates a UTC run_id under executions/dataset_runs/<run_id>/ on the control host), optionally adjusts row/col min-max, and Applies the grid. Cells are colour-coded: grey pending, orange collecting, green done. Optional grid metrics (origin X/Y and $\Delta x/\Delta y$ in metres) map matrix indices to physical coordinates for snapped positions. On the right, UEs & capture shows the live table (UE id, NR cell, RSRP, RSRQ, SINR) plus operator fields: label, Snap to location (row_col), X/Y in metres, and an Include checkbox. Capture controls set Duration (5–7200 s, default 300), Start collecting, Stop early, and Generate full dataset.

A typical session would be: create a run, set or snap UEs to distinct grid points, start a timed capture while the server polls Modelling KPIs once per second, when finished, move to the next points and repeat, finally generate and download a full export. Each capture writes raw sessions, exports builds export with a 1 Hz export_series per included UE (RSRP/RSRQ/SINR arrays and fill flags for gaps). Export length is taken from the longest session duration in the run (or 300 s if none). The run id is remembered in the browser until Forget run id.



Dataset lab (manual positions) Menu

Polls /api/modelling/ue-kpi-table for live RSRP/RSRQ/SINR. Capture uses the same source (~5 s refresh on the modelling host). **Generate full dataset** builds export_series (1 Hz resample + gap-fill) with one sample per second for the longest capture duration in this run (or 300 s if there are no finished sessions yet).

Location matrix

Row min: 0, Row max: 2, Col min: 0, Col max: 3

New run **Apply grid**

Base station at indices (0,0). Bounds may be negative (e.g. row -1...2, col -2...3). **Apply grid** uses UE Snap positions when set (always includes (0,0)); otherwise it uses the min/max fields. New run uses the same fields. Grey = pending, orange = collecting, green = done.

Grid: 0,0 (green), 0,1 (grey), 0,2 (orange), 0,3 (grey), 1,0 (grey), 1,1 (grey), 1,2 (grey), 1,3 (grey), 2,0 (grey), 2,1 (grey), 2,2 (grey), 2,3 (grey)

Origin X (m): 0, Origin Y (m): 0, Δx / col: 1, Δy / row: 1

Run: 20260702T073325Z

Forget run id

UEs & capture

Duration (s): 300 **Start collecting** **Stop early** **Generate full dataset**

Countdown: 297 s · Collecting KPIs...

live

Start runs one timed capture (server polls the KPI file every 1 s). The server marks **every location point** that appears in **Snap to location** on an included UE. Included UEs without Snap still contribute KPI rows; the matrix marks only snapped points. X/Y follow Snap or manual entry per UE.

UE id	NR cell	RSRP	RSRQ	SINR	Label	Snap to location	X (m)	Y (m)	Incl.
4	1	-71.7	-10.4	23.1		— manual	2	0	☒

Last API response

```
{
  "ok": true,
  "run_id": "20260702T073325Z",
  "location_points": [
    {
      "x": 0,
      "y": 2
    }
  ],
  "included_ue_ids": [
    ...
  ]
}
```

Figure 5-5 Date set generation interface.

The Dataset lab is the Model Creation data path on the real network, labelling radio samples at known locations for offline or partner training. It depends on Modelling being up and producing live KPIs; it does not itself run Inference or change Tx power. Trained artefacts and the NDT CSV later enter the closed loop through Control application.

Using the tool, a simple validation measurement campaign was conducted in the lab where the O-RU is located. The laboratory was divided according to a grid that encompassed the entire facility, as well as some outdoor areas. During the campaign, UEs implemented with Quectel modems were deployed following the defined grid and samples of the RSRP, RSRQ and SINR were gathered. From each points a total of 300 samples were collected with an interval period of 1 second.

Figure 5-6 shows the data generated by the dataset generation service. The generated file gathers together the measurement obtained during different rounds. It also indicates the points of the grid where data was generated, as well as the missing ones. Furthermore, for each measurement points it indicates the measuring time, location, a list with the values gathered and averaged values every 5 seconds.

```
{
  "schema_version": 1,
  "run_id": "20260706T102343Z",
  "generated_at": "2026-07-06T13:54:41Z",
  "source": "demo_control_phase_a",
  "source_note": "Samples from SSH cat of latest_ue_kpi_30s.json (~5s refresh, 30s rolling mean on host). export_series: per-UE 1 Hz bucket + gap-fill (dataset_resample.py) for SSB KPIs; dlBitRate, ulBitRate, QCI, rssi reserved (empty) for future sources.",
  "grid": {
    "row_min": 0,
    "row_max": 4,
    "col_min": -2,
    "col_max": 4,
    "rows": 5,
    "cols": 7,
    "locationPoint": [...]
  },
  "export_series": [
    {
      "ue_id": "19",
      "": "10"
```

Figure 5-6 Snippet of dataset (JSON format) generated during the validation measurement campaign

5.2.1.2 Setup 2 Automatic Data Acquisition

The **PoC#3** demo-related activity in the Greek site has been performed at the live **OTE** Lab, where the Accelleran dRAX RAN and ATHONET 5GCN are deployed in PROXMOX virtualized environment.

The pre-trial of PoC#3 at OTE in Athens aimed to pre-test the DT rApp and generate network performance datasets that will be used for estimating the performance of the NDT(s). The environment setup has been the same as the trial held in Santander, so that datasets could be portrayed between demo sites for the training of the NDT xApp. For this purpose, an ACC O-RAN gNB was deployed to the 5GCN of OTE Testbed.




Figure 5-7 ACC O-RAN solution in OTE testbed



a)



b)

Figure 5-8 Physical Installation of ACC O-RAN BS @ OTE testbed

The configuration of the **ACC** O-RAN gNB is as listed in Table 5-1.

Table 5-1 Configuration @ OTE testbed

Parameter	Value
SITE ID	42473 ACCORAN1
cellLocalId	131
EARFCN DL	636666
EARFCN UL	636666
Frequency Band	3500-1
cell Name	N1ACCORAN1E1
ssbFrequency	634080
nRPCI	3
rachRootSequence	66
bSChannelBwDL	100 (MHz)
bSChannelBwUL	100 (MHz)
Tx Power Leve	(a) 24dBm, (b) 19dBm, and (c) 14dBm

The RU's Tx power has been set to three levels as follows: (a) 24dBm, (b) 19dBm, and (c) 14dBm, for the purpose of obtaining performance and coverage data under different radio (energy/ power related) configurations. To this end, telemetry data was collected either from inherent network monitoring functions and from user collected and reported data. For this purpose, the OTE-inhouse developed performance and coverage monitoring tool named "COSMOTOOLS" was used. The tool captured and stored datasets -in the .csv and .json format - of network quality and performance metrics over time, location, conditions, etc. as received and reported from UEs registered to the COSMOTOOLS platform. The parameters that were captured are those listed in Table 5-2.

Table 5-2 COSMOTOOLS – Parameters / Data Received

Name of Parameter	Description / Data retrieved
GPS	Global Positioning System, used for location tracking
JITTER	The average delay time between two consecutive packets.
LATENCY	The time it takes for data to travel from source to destination (measured in milliseconds)
DL BITRATE	Download bitrate, the speed at which data is received from a network (measured in Mbps)
UL BITRATE	Upload bitrate, the speed at which data is sent to a network (measured in Mbps)
RSRP	Reference Signal Received Power (RSRP), indicates the power level of the received 5G signal
RSRQ	Reference Signal Received Quality (RSRQ), measures 5G signal quality by evaluating interference and power levels
RSSNR	Reference Signal Signal-to-Noise Ratio (RSSNR), indicates the strength of the signal compared to noise.

An indicative measurement in .json format has been the following:

For DOWNLOAD:

```
{
  "timestamp": 1783508747734,
  "latitude": 38.0499735,
  "longitude": 23.788279,
  "throughputMbps": 537,
  "direction": "DOWNLOAD",
  "host": "172.25.2.69",
  "command": "iperf3 -c 172.25.2.69 -p 5201 -t 5 --forceflush -R -u -b
1000M",
  "jitterMs": 0.021,
  "lostPackets": 40969,
  "totalPackets": 92329,
  "isSummary": false,
  "serverOutput": null,
  "rsrp": -78,
  "rsrq": -10,
  "rssnr": 34,
  "cellId": 1,
  "pci": 1,
  "tac": 1,
  "plmn": "20202",
  "carrierName": null,
  "networkName": "Leonardo",
  "isRoaming": false,
  "networkType": "NR (5G) "
},
```

For UPLOAD:

```
{
  "timestamp": 1783508729827,
  "latitude": 38.0499737,
  "longitude": 23.7882764,
  "throughputMbps": 23.7,
  "direction": "UPLOAD",
  "host": "172.25.2.69",
  "command": "iperf3 -c 172.25.2.69 -p 5201 -t 5 --forceflush -u -b
1000M",
  "jitterMs": null,
  "lostPackets": null,
  "totalPackets": null,
  "isSummary": false,
  "serverOutput": null,
  "rsrp": -85,
  "rsrq": -10,
  "rssnr": 33,
  "cellId": 1,
  "pci": 1,
  "tac": 1,
  "plmn": "20202",
  "carrierName": null,
  "networkName": "Leonardo",
  "isRoaming": false,
  "networkType": "NR (5G) "
},
```

The updated for 5G GUI of the tool is that included in Figure 5-9.




Figure 5-9 5G COSMOTOOLS GUI used for capturing datasets for NDT training

5.2.1.3 Inference

Control application: The Control application is the operator surface for the closed loop on the Optimisation Control xApp: live radio metrics → Inference → NDT lookup → optional Tx apply via dRAX, with logging and spatial visualisation of predicted locations. Scheduling of the full cycle (Once / every minute / every five minutes) runs on the demo-control host; the browser mainly displays state and triggers actions.

Data metrics: The loop starts from live UE radio metrics produced by the Modelling xApp. Modelling consumes RIC UE measurements and maintains rolling 30-second averages of SSB RSRP, RSRQ and SINR in latest_ue_kpi_30s.json. Demo-control reads that file over SSH and exposes it to the Control page as the per-UE SSB table (UE id, NR cell, RSRQ, SINR, RSRP). The control cycle uses a selected UE row (typically the first / strongest in the table) as Inference input. In parallel, a background sampler on demo-control polls configured Tx power from the Optimisation Control path (dRAX maximum_eirp_dbm on cell-1) about every 30 seconds and stores history for the Tx-over-time and UE-metric charts. The UI itself refreshes state about every 10 seconds so operators see metrics, Inference/NDT results and logs without overloading SSH. Metrics therefore play two roles: they feed the next Inference, and they provide continuous situational awareness between cycles. Figure 5-10, shows the interface for the data metrics and the control cycle page.




Figure 5-10 Control application metrics from RU and UEs

Inference box: When a control cycle runs, the Optimisation Control pipeline normalises the selected UE's RSRP/RSRQ/SINR and passes them to the location estimator. As an validation example a simple K-Nearest neighbors (KNN) estimator has been implemented using the data collected at the **UC**, as described in section 5.2.1.1. In particular 80/20% of the samples were used for training/testing, respectively. The results, shown in Figure 5-11, illustrate a reasonable location prediction accuracy for most of the measurement points, except in one case. It is worth highlighting that this predictor is intended for validation of the Intelligent Plane flow, rather than to be a closed solution.




Figure 5-11 Control application metrics from RU and UEs

Inference

Location **x_-10_y_000** X-10 Y0
 Mode **poc3_knn**

Inference model
Apply Model
Load...
Download
Reset

PoC3 KNN package · active · model 388136 B · scaler 639 B · estimator 2316 B · mtime 2026-07-17T11:40:30Z · sha a9dfc441

Load a .zip (knn_mod1.pkl + scaler.save + poc3_loc_estimator.py), then **Apply Model** to activate.

Download started (poc3_inference_package_2026-07-23T072042Z.zip)

► Advanced: load one artifact

Figure 5-12 Inference box for the control application

Once trained, the model can be called with new samples (i.e. RSRP, RSRQ and SINR values), see Figure 5-12, and it returns a location label such as “x_010_y_008”, which is parsed into grid coordinates (x, y). The Inference card on the UI shows that location string, the parsed X and Y, and the mode (e.g. poc3_knn; a dummy fallback only if the model or dependencies are missing). Operators can also manage the Inference package from the same card: **Load** stages a new model set, **Apply** activates it on optctrl, **Download** retrieves the active package, and **Reset** restores the factory artefacts. Inference does not change the radio; it only produces the location used by the next stage. The solution accepts models implemented as Python classes that only need to implement a simple interface consisting on two functions: a constructor (`__init__` function) accepting a model filename and a scaler name (both in PKL format), and a **Predict** function which is passed the new sample and return the location label mentioned above. As can be seen, the implemented solution provides high flexibility to update both the model or the whole Python script if additional operations need to be performed over a trained model.

Inference log, predicted path and location heatmap: Below the twin cards, the page keeps a browser-side record of how Inference behaves over time during a session (see Figure 5-14). The Inference log appends a row when values change: time, actual configured Tx, RSRP, RSRQ, SINR, and predicted location. Rows can be exported as CSV or cleared; they are kept in session storage for the tab (no xApp restart required). From those location labels the UI builds a predicted path: successive (x, y) visits on the UC grid (X from -10 to 20, Y from 0 to 8), drawn over the lab floor-plan underlay, with visit-share labels and an antenna marker at (0,0).

The location heatmap in Figure 5-13 uses the same visit set: empty cells stay light, and colour intensity rises with the share of visits. Together, turn the closed loop into an observable spatial story—metrics drive Inference, Inference drives NDT/Control, and the lower panels show where the estimator has placed the UE across the trial.




Figure 5-13 Inference Heat map for inferred location.

Inference log (testing) 13 rows Export CSV Clear

Appends when values change. Tx = actual RU configured power (not NDT optimal).

Time	Tx power (dBm)	RSRP (dBm)	RSRQ (dB)	SINR (dB)	Predicted location
2026-07-23T13:30:30.605Z	19	-78.7	-10.2	23.2	x_010_y_002
2026-07-23T13:30:20.586Z	19	-78.7	-10.3	23.2	x_-10_y_000
2026-07-23T13:30:05.917Z	19	-75	-10.3	23.2	x_-10_y_000
2026-07-23T13:29:55.994Z	24	-73.9	-10.2	23.2	x_005_y_006
2026-07-23T13:29:46.031Z	24	-76.6	-10.4	23.1	x_005_y_008
2026-07-23T13:29:33.316Z	24	-78.2	-10.4	23.1	x_005_y_008
2026-07-23T13:29:23.341Z	19	-79.2	-10.4	23.1	x_005_y_008
2026-07-23T13:29:13.260Z	19	-78.2	-10.2	23.2	x_-05_y_002
2026-07-23T13:29:03.282Z	19	-78.6	-10.2	23.2	x_010_y_002
2026-07-23T13:28:53.280Z	19	-79.1	-10.2	23.2	x_010_y_002

Figure 5-14 Inference logs for PoC#3

5.2.1.4 NDT

Once the location is predicted, the Intelligent Plane flow exploits that information to optimise the network configuration. Concretely, the current implementation adapts the transmission power of the O-RU to foster energy efficiency, while keeping the given certain QoS requirements. To that end an additional box is available for the control application (see Figure 5-15).

NDT box: The NDT stage takes the inferred (x, y) and looks up an optimal Tx power in a matrix CSV (ndt_tx_lookup.csv) held on the Optimisation Control host. The NDT + Control card shows the point @x,y, the currently configured Tx (same source as the RU table), the optimal Tx from the table, and a control LED. If the lookup misses, there is no optimal value and no automatic Tx send (LED stays idle/grey). If an optimal value exists and differs from configured Tx beyond a small tolerance (0.2 dBm), Control applies it through the dRAX Configuration API and confirms within about 10 seconds (LED orange while applying, green when matched, red on failure). Like Inference, the NDT package supports Load / Apply / Download / Reset so partners can replace the lookup table without redeploying the whole stack. In process terms: Inference answers “where is the UE?”, NDT answers “what Tx should we use there?”, and Control enacts that decision on the live cell when needed.



Figure 5-15 NDT box for the control application




Figure 5-16 System metrics for the control NDT and RAN behaviour

As a functional validation of the Intelligent Plane flow, an NDT has been implemented using a hybrid simulation approach, as depicted in Figure 5-17. It uses the widely adopted 5G-LENA simulator, based on ns-3 simulation framework. It allows the configuration of scenario simulation embracing multiple nodes, each implanting the whole protocols stack, including applications. In order to provide higher realism for specific scenarios, the latest version of 5G-LENA has been integrated the SIONNA Ray Tracing (SIONNA-RT) simulator. The interaction between both elements is performed within the simulator, by using specialised libraries that allow using Python code (i.e. SIONNA-RT) from C++ (ns-3). This new feature allows the ingestion of default and customised scenarios defined as Mitsuba files.

The validation has included the modeling of the UC lab where the O-RU (Benetel) is located, as shown in Figure 5-17, using the Blender computer graphics software toolkit. Later on, the generated scenario has been used within ns-3 to obtain the QoS parameters, such as SINR, for each of the measurement points described

Section 5.2.1.3. The simulation has been repeated for the different transmission power values that can be configured in the O-RU.




Figure 5-17 Functional architecture of the NDT, based hybrid simulation, to feed the Intelligent Plane

Although this simulation approach allows the generation of a rather precise NDT, it comes at the expense of high computation requirements and long simulation periods, according to the specific scenarios. In that sense, the developed NDT would be suitable for network planning solutions, or it can be exploited to generate configuration look-up tables that can be used real-time. More information related to these last topics will be provided at the Final Project Review.

6 Summary and Conclusions

This deliverable has presented the final implementation, demonstration, and experimental validation activities carried out in **WP5** of the **6G-SENSES** project. Building upon the architectural design established in previous work packages and the integration activities reported in **D5.2**, the work described in this document demonstrates the feasibility of combining advanced communication, sensing, and AI technologies within a common 6G-oriented architecture. The three Proofs of Concept (PoCs) collectively validate key technological building blocks of the **6G-SENSES** vision while providing valuable experience on the practical challenges associated with integrating heterogeneous hardware platforms, software frameworks, and AI-driven network intelligence.

PoC#1 demonstrated the successful integration of heterogeneous sensing technologies into an O-RAN-compliant Radio Access Network. The implementation combined sensing-capable 5G radio units and non-3GPP sensing devices within a common control framework, enabling sensing information to be exposed through standardized O-RAN interfaces and exploited by xApps to support communication services. A particularly relevant achievement was the implementation of a monostatic sensing-enabled radio unit based on OpenAirInterface and commercial SDR platforms, exploiting the flexibility of the 5G NR frame structure to embed sensing transmissions without modifying the standard communication framework. The experimental validation confirmed that heterogeneous sensing information can effectively support network optimization while also highlighting the need for improved mechanisms to transport high-volume sensing data within future O-RAN systems.

PoC#2 validated the integration of advanced Cell-Free MIMO technologies into experimental 5G platforms. The work demonstrated that distributed beamforming and OTA reciprocity calibration can be incorporated into practical software-defined radio platforms while providing measurable performance improvements over conventional small-cell deployments. At the same time, the implementation revealed the challenges of extending existing 5G software stacks towards large-scale distributed antenna systems. The migration from legacy proprietary development environments towards open-source implementations provided valuable insight into the current limitations of software architectures for experimental massive MIMO research and confirmed the need for more scalable Layer-1 processing and beamforming support in future open-source RAN platforms.

PoC#3 demonstrated the integration of AI-native network intelligence through the implementation of an Intelligent Plane supported by an NDT. By combining real network telemetry, AI model lifecycle management, and hybrid simulation environments based on ns-3 and Sionna RT, the project established a complete workflow for creating, validating, and deploying AI models capable of supporting network optimization. The work demonstrated that DTs can become an effective component of future network management frameworks, although computational complexity and synchronization between real and virtual environments remain important challenges before such approaches can be deployed for fully real-time operation.

An important outcome of the integration activities is the confirmation that the **6G-SENSES** architecture provides a flexible framework capable of accommodating multiple sensing technologies, advanced radio access techniques, and AI-native control mechanisms within a common architectural vision. While the three PoCs addressed complementary aspects of the architecture, they also demonstrated that integrating innovations across the physical layer, RAN, transport network, and Intelligent Plane introduces significant engineering complexity. Rather than pursuing a monolithic end-to-end implementation, the project adopted a modular integration strategy in which common software platforms—including **BR** for the ISAC platform and the **ACC** Open 5G/dRAx framework for the Intelligent Plane—enabled independent technological developments to converge towards a coherent architecture while maintaining manageable implementation complexity.

Beyond the individual technical achievements, the experimentation campaign provided several important lessons for future 6G systems. First, practical ISACC solutions can be realized by extending programmable software-defined radio platforms and existing 5G frame structures instead of introducing entirely new communication paradigms. Second, open-source software ecosystems such as OAI provide an excellent basis for standards-compliant experimentation but still require significant architectural evolution to efficiently support large-scale distributed MIMO and advanced beamforming techniques. Finally, AI-native network operation requires much more than inference engines; it depends on complete model lifecycle management, trustworthy data pipelines, and realistic DTs capable of accurately representing network behaviour.

Overall, the results presented in this deliverable provide strong evidence that the technological concepts proposed by 6G-SENSES are technically feasible and can be successfully integrated into representative laboratory environments. At the same time, the project has identified several research and engineering challenges that remain open, including scalable sensing data management, efficient support for distributed antenna systems, and computationally efficient DT frameworks. These lessons constitute an important contribution to the evolution of future AI-native 6G architectures and provide a solid foundation for the continued development of integrated communication and sensing technologies beyond the lifetime of the project.

7 References

- [1] 6G-SENSES deliverable D5.1, “Report on the testing methodologies and testbed setup”, November 2025, [link](#)
- [2] 6G-SENSES deliverable D5.2, “Report on the integration of 6G-SENSES technologies”, June 2026, [link](#)
- [3] 6G-SENSES deliverable D2.1, “Report on 6G-SENSES use cases, network architecture, KPIs and supported RAN functions”, September 2024, [link](#)
- [4] 6G-SENSES deliverable D2.2, “System architecture and preliminary evaluations”, February 2025, [link](#)
- [5] 6G-SENSES deliverable D2.3, “Final report on 6G-SENSES network architecture evaluation”, January 2026, [link](#)
- [6] 6G-SENSES deliverable D3.1, “Initial report on the development of 6G-SENSES infrastructure building blocks”, January 2025, [link](#)
- [7] 6G-SENSES deliverable D3.2, “Final report on the development of 6G-SENSES infrastructure building blocks”, January 2025, [link](#)
- [8] 6G-SENSES deliverable D4.1, “Initial SoTA and design of wireless edge caching solutions”, January 2026, [link](#)
- [9] 6G-SENSES deliverable D4.2, “Final design and implementation of Wireless Edge Caching solutions”, January 2026, [link](#)
- [10] 6G-SENSES deliverable D4.3, “AI/ML techniques for enhanced network performance”, May 2026, [link](#)
- [11] 6G-SENSES deliverable D4.4, “Design and implementation of the 6G-SENSES Intelligent Plane”, July 2026
- [12] T. X. Vu, E. Baştuğ, S. Chatzinotas, and T. Q. S. Quek, “Wireless Edge Caching: Modeling, Analysis, and Optimization”, in *Cambridge University Press*, 2021.
- [13] K. S. Tharakan, B. N. Bharath, V. Bhatia, J. Nebhen, M. Dobrovolny and T. Ratnarajah, "Wireless Edge Caching and Content Popularity Prediction Using Machine Learning," in *IEEE Consumer Electronics Magazine*, vol. 13, no. 4, pp. 32-41, July 2024, doi: 10.1109/MCE.2022.3160585.
- [14] P. Yang, N. Zhang, S. Zhang, L. Yu, J. Zhang and X. Shen, "Content Popularity Prediction Towards Location-Aware Mobile Edge Caching," in *IEEE Transactions on Multimedia*, vol. 21, no. 4, pp. 915-929, April 2019, doi: 10.1109/TMM.2018.2870521.
- [15] Massive Beams, “A NEW ERA IN THE WORLD OF SOFTWARE-DEFINED RADIOS”, 2025, <https://www.massivebeams.com/>
- [16] J. Sanson et al. "Extracting Range–Doppler Information of Moving Targets from Wi-Fi Channel State Information." To appear in Proc. IEEE GLOBECOM (2025). <https://arxiv.org/abs/2508.02799>
- [17] V. Frasca, D. Cavalcanti, and R. Shah. "Wi-Fi Evolution: The Path Towards Wi-Fi 7 and its impact on IIoT." *Journal of Mobile Multimedia* 19, 01 (2022). doi:10.13052/jmm1550-4646.19113

- [18] J. Gutiérrez et al. "Seamless Integration of Efficient 6G Wireless Technologies for Communication and Sensing Enabling Ecosystems." AIAI 2024 IFIP WG 12.5 Int. Workshops.
- [19] A. Tzanakaki and M. Anastasopoulos, "Optical Transport Networks Supporting Integrated Communications and Sensing in 6G," 2025 European Conference on Optical Communications (ECOC), Copenhagen, Denmark, 2025, pp. 1-4, doi: 10.1109/ECOC66593.2025.11263441.
- [20] C. Shepard, H. Yu, N. Anand, E. Li, T. Marzetta, R. Yang, and L. Zhong, "Argos: practical many-antenna base stations". In Proceedings of the 18th annual international conference on Mobile computing and networking (Mobicom '12), 2012.
- [21] U. K. Ganesan, R. Sravendranath, E. G. Larsson, "BeamSync: Over-The-Air Synchronization for Distributed Massive MIMO Systems," in IEEE Transactions on Wireless Communications, 2024, doi:10.1109/TWC.2023.3335089.

8 Acronyms

Acronym	Description
3GPP	3rd Generation Partnership Project
5G NR	5G New Radio
5QI	5G Quality of Service Identifier
ACC	Accelleran (6G-SENSES Beneficiary)
AI	Artificial Intelligence
AI-P	Antenna-in-Package
AoA	Angle of Arrival
AP	Access Point
BiCMOS	Bipolar-CMOS
BI	<i>Barkhausen Institut</i> (6G-SENSES Beneficiary)
BR	BubbleRAN (6G-SENSES Beneficiary)
BS	Base Station
BW	Bandwidth
CF-MIMO	Cell-Free MIMO
CMOS	Complementary Metal-Oxide-Semiconductor
CN	Core Network
CI/CD	Continuous Integration/Continuous Deployment
CIR	Channel Impulse Response
COTS	Commercial-Off-The-Shelf
CPRI	Common Public Radio Interface
CQI	Channel Quality Indicator
CSI	Channel State Information
CU	Central Unit
DB	DataBase
DC	Direct Current
DFT	Discrete Fourier Transform
DL	Downlink
DN	Data Network
DoW	Description of Work
DPP	Drift-plus-penalty
DU	Distributed Unit
DUT	Device Under Test

e2e	End-to-End
eCPRI	Enhanced CPRI
EPC	Evolved Packet Core
E2AP	E2 Application Protocol
E2SM	E2 Service Model
FFT	Fast Fourier Transform
FL	Federated Learning
GBR	Granted Bit Rate
gNB	gNodeB
GNSS	Global Navigation Satellite System
GNU	GNU is Not Unix
GTEC	Communications and Electronics Technology Group
GTIS	GTEC Testbed Interface Software
GUI	Graphical User Interface
PoC	Proof of Concept
HRLLC	Hyper Reliable and Low-Latency Communication
IASA	Institute of Accelerating Systems and Applications (6G-SENSES Beneficiary)
IHP	<i>IHP – Leibniz-Institut für innovative Mikroelektronik</i> (6G-SENSES Beneficiary)
INT	Intel Deutschland GmbH (6G-SENSES Beneficiary)
IP	Internet Protocol
ISAC	Integrated Sensing and Communication
IMT	International Mobile Telecommunications
I/Q	In-phase and Quadrature components
ITU	International Telecommunication Union
ITU-R	ITU Radiocommunication Sector
KPI	Key Performance Indicator
LAC	Location Area Code
LCM	Life Cycle Management
LENA	LTE/EPC Network simulator
LLC	Lower Layer Control
LTE	Long-Term Evolution
MAC	Medium Access Control
MC-LAG	Multi-Chassis Link Aggregation Group
MCS	Modulation and Coding Scheme
MEC	Multi-Access Edge Computing

MIMO	Multiple-Input Multiple-Output
ML	Machine Learning
mmWave	Millimeter Wave
Multi-RAT	Multiple Radio Access Technology
N3IWF	Non-3GPP Interworking Function
NDT	Network Digital Twin
Near-RT	Near-Real-Time
NGAP	Next Generation Application Protocol
NI	National Instruments
NIC	Network Interface Card
NN	Neural Network
Non-RT	Non-Real-Time
NR	New Radio
NS3	Network Simulator 3
N3IWF	Non-3GPP Interworking Function
NTU	Nottingham Trent University
OAI	OpenAirInterface
ODL	OpenDayLight
O-DU	Open-DU
OFDM	Orthogonal Frequency-Division Multiplexing
O-RAN	Open-RAN
O-RU	Open-RU
OS	Operative System
OTA	Over-The-Air
OTE	<i>ORGANISMOS TILEPIKOINONION TIS ELLADOS OTE AE (6G-SENSES Beneficiary)</i>
PCB	Printed Circuit Board
PoC	Proof of Concept
PoE	Power over Ethernet
Pout	Outage Probability
PHY	Physical Layer
PPS	Pulse Per Second
PRB	Physical Resource Block
PTP	Precision Time Protocol
PUCCH	Physical Uplink Control Channel
PUSCH	Physical Uplink Shared Channel

QoS	Quality of Service
RAN	Radio Access Network
RAT	Radio Access Technology
RB	Resource Block
REF	Reference Clock
RFSoc	Radio Frequency System-on-Chip
RIC	RAN Intelligent Controller
RIS	Reconfigurable Intelligent Surface
ROS	Robot Operating System
RSRP	Reference Signal Received Power
RSRQ	Reference Signal Received Quality
RSSI	Received Signal Strength Indicator
RSSNR	Reference Signal Signal-to-Noise Ratio
RT-RIC	Real Time RIC
RTT	Round-Trip Time
RU	Radio Unit
SCTP	Stream Control Transmission Protocol
SDR	Software-Defined Radio
SINR	Signal-to-Interference-plus-Noise Ratio
SMD	Surface Mount Device
SMO	Service Management and Orchestration
SNDR	Signal-to-Noise-and-Distortion-Ratio
SNR	Signal-to-Noise Ratio
SotA	State-of-the-Art
SRS	Sounding Reference Signal
sRU	sensing-Radio Unit
SUT	System Under Test
TAC	Tracking Area Code
TCO	Total Cost of Ownership
TDD	Time Division Duplex
TDoA	Time Difference of Arrival
TLV	Type-Length-Value
TRP	Transmission Reception Points
TUBS	<i>Technische Universität Braunschweig (6G-SENSES Beneficiary)</i>
UC	<i>Universidad de Cantabria (6G-SENSES Beneficiary)</i>

UDC	University of A Coruña
UE	User Equipment
UHD	USRP Hardware Driver
UL	Uplink
ULA	Uniform Linear Array
UNIROMA1	<i>Sapienza Università di Roma (6G-SENSES Beneficiary)</i>
UPF	User Plane Function
USB	Universal Serial Bus
USRP	Universal Software Radio Peripheral
VM	Virtual Machine
WAT	Wireless Access Technology
WEC	Wireless Edge Caching
Wi-Fi	Wireless Fidelity
WP	Work Package