

D4.3: Final deployment of beyond 5G RAN, core, and open-source networks, disruptive RAN technologies and trial controller

Revision: 1.0

Work package	WP4
Task	T4.1, T4.2, T4.3, T4.4
Due date	30/06/2025 (M30)
Submission date	30/06/2025
Deliverable lead	IDE
Version	1.0
Authors	Nokia: Olli Liinamaa VTT: Jarno Pinola, Olli Apilo Vicomtech: Roberto Viola CGE: Javier Godas, Enrique Rubio, Aurora Ramos I2CAT: Daniel Camps Mur, Amr AbdelNabi UOulu: Antti Pauanne, Jussi Haapola, Mahdi Salmani, Jani Pellikka, Marko E. Leinonen IT: Hemalatha Vulchi, Jonathan Rodriguez, Manuel Violas Interdigital: Javier Lorca Hernando Ericsson: Diego San Cristobal IMEC: Miguel Glassee, Björn Debaille
Reviewers	IT: Mohammed Al-Rawi (al-rawi@av.it.pt) CGE: Jose Soriano (jose.a.soriano@capgemini.com)

Abstract	This document completes the final design and performance results of the XR enablers for 3GPP, O-RAN and Disruptive paths towards 6G. It also reports the final implementation and results of the Trial Controller.
Keywords	Extended Reality (XR), Holographic Communications, 3D Digital Twin, Energy optimization, O-RAN, Trial controller, Network controlled repeaters, Reconfigurable intelligent surfaces.

Document Revision History

Version	Date	Description of change	List of contributor(s)
V0.1	17/01/25	1st version of the ToC for comments	IDE
V0.2	28/02/25	1 st version of the Final Solution Design	See Authors
V0.3	02/05/25	1 st version of the Final Results	See Authors
V0.4	20/06/25	Final version ready for external review	See Authors

DISCLAIMER



Co-funded by
the European Union



Project funded by

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation

Federal Department of Economic Affairs,
Education and Research EAER
State Secretariat for Education,
Research and Innovation SERI

6G-XR (6G eXperimental Research infrastructure to enable next-generation XR services) 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 No 101096838. This work has received funding from the [Swiss State Secretariat for Education, Research, and Innovation \(SERI\)](#).

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

COPYRIGHT NOTICE

© 2023 - 2025 6G-XR Consortium

Project co-funded by the European Commission in the Horizon Europe Programme		
Nature of the deliverable:	R	
Dissemination Level		
PU	Public, fully open, e.g. web (Deliverables flagged as public will be automatically published in CORDIS project's page)	✓
SEN	Sensitive, limited under the conditions of the Grant Agreement	
Classified R-UE/ EU-R	EU RESTRICTED under the Commission Decision No2015/ 444	
Classified C-UE/ EU-C	EU CONFIDENTIAL under the Commission Decision No2015/ 444	
Classified S-UE/ EU-S	EU SECRET under the Commission Decision No2015/ 444	

* R: Document, report (excluding the periodic and final reports)

DEM: Demonstrator, pilot, prototype, plan designs

DEC: Websites, patents filing, press & media actions, videos, etc.

DATA: Data sets, microdata, etc.

DMP: Data management plan

ETHICS: Deliverables related to ethics issues.

SECURITY: Deliverables related to security issues

OTHER: Software, technical diagram, algorithms, models, etc.

EXECUTIVE SUMMARY

This is the 3rd and final deliverable of Work Package 4 (WP4) – “Experimental RAN infrastructure”. WP4 focuses on the deployment of 6G-XR beyond 5G cellular network infrastructure. Deliverable D4.3 describes final deployment and performance evaluation of five 3GPP enablers, three O-RAN enablers, and three Disruptive enablers. The deliverable also reports on a final implementation of the Trial Controller. New enablers are integrated and tested at the project’s North Node (UOULU 5GTN and VTT 5GTN) and South Node (5GBarcelona and 5TONIC), and performance results are also documented if available. Deployment includes also XR trial controller meant to ease the automatic deployment of XR-related experiments and remote use of testbed.

Among the 3GPP XR enablers, advanced simulations proved the ability of network-controlled repeaters in indoor environment to boost indoor XR network performance and ensure robust QoE under demanding usage conditions. The Network-assisted Rate Control API was designed to enhance XR service rate adaptation by integrating explicit rate recommendations from the network, which will be tested in I2CAT laboratory and integrated within the South Node infrastructure. ATSSS-based capacity enhancement was proved to effectively perform traffic steering and switching in maintaining service quality for latency- and bandwidth-sensitive XR applications across heterogeneous access networks. Energy and application-aware management of 3GPP infrastructure highlighted the importance of both energy and application awareness when making optimization decisions for the infrastructure. Finally, the upgrades and evaluation of experimental RAN infrastructure in the South and North Nodes demonstrated the state-of-the-art performance available in the test facilities setup by the 6G-XR project.

Among the O-RAN XR enablers, congestion aware load balancing enabled proactive mitigation of congestion, efficient user distribution across cells, and sustained service quality for XR applications within O-RAN-enabled networks. Energy-aware end-to-end resource management proved that the maximum power consumption of the overall O-RAN infrastructure can be controlled by limiting the allowed data rates of lower priority users. Finally, O-RAN enabled slicing to support XR services proved slice isolation prevents extra latency spikes from appearing when another slice is overloaded with data.

Among the Disruptive XR enablers, the High-Frequency transceivers for THz-RIS and ISAC at 140 GHz and 300 GHz demonstrated successful establishment of wireless links capable of reliably delivering a high throughput with beamsteering functionality, mostly aligning with simulations, and in some cases requiring the development of multiple new measurement capabilities during the project. Baseband implementation for THz-RIS and ISAC based on SC-FDE demonstrated how to build a full disruptive system by splitting it into their baseband and frontend components that are offline integrated in the lab, and theoretically studied a technique to estimate and equalize any type of doubly-dispersive channel for high-speed applications. Finally, Deep Reinforcement Learning (DRL) for THz-RIS demonstrated significant advancements and improvements in the performance of the DRL algorithm compared to previously employed methods to benefit from coverage improvements with very large RIS structures.

The final Trial Controller implementation comprises a Unified Web Portal (designed with the aim of managing the trial on both the North and South web portals) and the corresponding web portals and adapters of the North and South Nodes. The Trial Controller used Qosium as the KPI collection engine for North Node and South Node monitoring.

In addition to the results and conclusions, next steps are hinted for each of the above-mentioned enablers and Trial Controller to successfully leverage the results obtained in the 6G-XR Project for future lines of research.

TABLE OF CONTENTS

DISCLAIMER	3
COPYRIGHT NOTICE	3
1 INTRODUCTION	15
1.1 Objective of the deliverable	15
1.2 Structure of the deliverable	15
1.3 Target audience of the deliverable	16
2 3GPP XR ENABLERS	17
2.1 Summary of 3GPP XR enablers in D4.2	17
2.2 3GPP XR proposed enablers: design and final implementation results	17
2.2.1 Network-Controlled Repeaters in indoor environment	17
2.2.2 Network assisted Rate Control API	22
2.2.3 ATSSS based capacity enhancement	26
2.2.4 Energy and Application aware management of 3GPP infrastructure	29
2.2.5 Upgrade and evaluation of experimental RAN infrastructure in SN and NN	33
3 O-RAN XR ENABLERS	45
3.1 Summary of O-RAN enablers in D4.2	45
3.2 O-RAN XR proposed enablers	45
3.2.1 Congestion aware load balancing	45
3.2.2 Energy-aware end-to-end resource management	53
3.2.3 O-RAN enabled slicing to support XR services	57
4 6G DISRUPTIVE XR ENABLERS	63
4.1 Summary of disruptive enablers proposed in D4.2	63
4.2 6G disruptive XR proposed enablers	63
4.2.1 High-Frequency Transceivers for THz-RIS and ISAC	63
4.2.2 Baseband implementation for THz-RIS and ISAC based on SC-FDE	71
4.2.3 Deep Reinforcement Learning for THz-RIS	77
5 THE 6G-XR TRIAL CONTROLLER	83
5.1 Final Trial Controller implementation	83
5.1.1 Unified Web Portal	84
5.1.2 North Node Web Portal and Adapter	92
5.1.3 South Node Web Portal and Adapter	97
5.2 KPIs collection and visualization	102

5.2.1	North Node monitoring.....	103
5.2.2	South Node monitoring.....	105
6	SUMMARY	107
7	REFERENCES.....	111
8	ANNEXES	113
8.1	Description of xApps and control signalling for the energy-aware end-to-end resource management O-RAN enabler.....	113
8.2	Channel Estimation and Equalization in Doubly-Dispersive Channels.....	115
8.3	Unified Web Portal User Guide	119
8.4	North Web Portal User Guide.....	121
8.5	South Web Portal User Guide.....	124

LIST OF FIGURES

FIGURE 1: NCR METHODOLOGY FLOWCHART.....	18
FIGURE 2: NCR PERFORMANCE UNDER THE BAND FR1, BW-50MHZ, SCS=15KHZ, TRAFFIC MODEL TYPE-XR GAMING, AND AVEIRO STREET CITY MAP SCENARIO.....	19
FIGURE 3: (A) IMPACT OF ADJUSTED NCR POSITION ON SNR AND BLER UNDER FR1 BW-(100 MHZ, SCS = 15 KHZ) NCR [-137.6145; -67.8227] IN THE I2CAT INDOOR SCENARIO (B) EVALUATING NCR PERFORMANCE WITH SYNTHETIC TRAFFIC MODELS IN THE SIMULATOR AND EXPERIMENTAL TRAFFIC MODEL DATA COLLECTED THROUGH USE CASES.....	21
FIGURE 4: NETWORK ASSISTED RATE CONTROL API ARCHITECTURE (NEXTXRATE) WITH RED ARROW INDICATING THE NEW PROPOSED ANALYTICS.	23
FIGURE 5: PERFORMANCE COMPARISON OF NETXRATE AND OTT WITH EQUAL PRIORITY FOR XR AND EMBB USERS.	25
FIGURE 6: PERFORMANCE COMPARISON OF NETXRATE AND OTT WITH XR PRIORITIZED OVER EMBB USERS.....	25
FIGURE 7: ATSSS API ARCHITECTURE FEATURING RULE-BASED TRAFFIC FORWARDING	27
FIGURE 8: COMPARATIVE ANALYSIS OF CPE UPLINK THROUGHPUT OVER TIME IN WIFI AND 5G ACCESS NETWORKS	28
FIGURE 9: 3GPP NG-RAN TEST SETUP FOR ENERGY AND APPLICATION AWARE RESOURCE CONTROL. ..	30
FIGURE 10: MEASURED MAXIMUM THROUGHPUTS FOR 6G-XR UC1-4 WITH DIFFERENT 3GPP RAN ENERGY SAVING STATES.	32
FIGURE 11: ESTIMATED NUMBER OF SIMULTANEOUS XR USERS SUPPORTED IN DIFFERENT 3GPP RAN ENERGY SAVING STATES.	33
FIGURE 12: MAP OF TARGET TEST LOCATIONS WITH RED SQUARES SHOWING THE LOCATION OF THE ANTENNAS.....	34
FIGURE 13: FR1 OUTDOOR AND INDOOR INSTALLATION, UOULU 5GTN.....	36
FIGURE 14: FR1 AND FR2 GNB ANTENNAS (TRIANGLES) AND MEASUREMENT LOCATIONS (CROSSES) AT VTT 5GTN TEST FACILITY.	37
FIGURE 15: DOWNLINK AND UPLINK THROUGHPUT SATURATION BENCHMARK OVER FR1 AT VARYING DISTANCES FROM THE ANTENNA.....	39
FIGURE 16: DOWNLINK AND UPLINK THROUGHPUT SATURATION BENCHMARK OVER FR2 AT VARYING DISTANCES FROM THE ANTENNA.....	39
FIGURE 17: COMPARATIVE ANALYSIS OF DOWNLINK LATENCY AT DIFFERENT LOCATIONS.....	40
FIGURE 18: MAX THROUGHPUT MEASUREMENTS FOR FR1 (N77)	41
FIGURE 19: DELAY MEASUREMENTS FOR SLICE/CHANNEL OVERLOAD (FR1)	42
FIGURE 20: MAXIMUM THROUGHPUT MEASUREMENTS FOR TCP AND UDP TRAFFIC AT VTT 5GTN TEST FACILITY.	43
FIGURE 21: ONE-WAY LATENCY MEASUREMENTS FOR UDP TRAFFIC AT VTT 5GTN TEST FACILITY.	44
FIGURE 22: DESIGN OF O-RAN CDF FUNCTION TO ENFORCE HO UPON CONGESTION	47
FIGURE 23: SNAPSHOT OF THE HOLOGRAPHIC CALL WITH/WITHOUT CDF ALERTS	48
FIGURE 24 XGBOOST AND LSTM MODELS COMPARISON TABLE	50

FIGURE 25 XGBOOST REFITTED MODEL RECALL VS CONSTANT MODEL RECALL.....	50
FIGURE 26 LSTM REFITTED MODEL RECALL VS CONSTANT MODEL RECALL.....	51
FIGURE 27: PLOT OF UPLINK PRB UTILIZATION VS TIME FOR BOTH CELLS	52
FIGURE 28: PLOT OF UE UPLINK THROUPTUT VS TIME FOR BOTH CELLS.....	52
FIGURE 29: O-RAN TEST SETUP AND INFORMATION FLOWS.....	53
FIGURE 30: MEASURED MAXIMUM DL THROUGHPUTS FOR UC3 AND UC4.	55
FIGURE 31. MEASURED POWER CONSUMPTION OF THE O-RAN INFRASTRUCTURE FOR UC3 AND UC4.	55
FIGURE 32. EXAMPLE OF THE CHANGES IN THE OVERALL NETWORK LOAD AND O-RAN INFRASTRUCTURE POWER CONSUMPTION WITH HIGH PRIORITY XR UE (UE1) AND LOW PRIORITY UE (UE2).....	56
FIGURE 33: O-RAN-ENABLED SLICING TO SUPPORT XR SYSTEMS.....	58
FIGURE 34: 140 GHZ COMMUNICATION LINK SET-UP	64
FIGURE 35: D-BAND LINK IN ABSORBER CAGE	65
FIGURE 36: TX BEAM PATTERN FOR VARIOUS BEAMSTEERING CONFIGURATIONS	66
FIGURE 37: RX BEAM PATTERN FOR VARIOUS BEAMSTEERING CONFIGURATIONS	66
FIGURE 38: BASEBAND SPECTRUM AND CONSTELLATION DIAGRAM OF HIGHEST THROUGHPUT MODE.....	67
FIGURE 39: BASEBAND SPECTRUM AND CONSTELLATION DIAGRAM OF MOST SPECTRALLY EFFICIENT MODE.....	67
FIGURE 40: RF ARCHITECTURE OF THE DESIGNED 300 GHZ SLIDING IF RECEIVER AND PHOTOGRAPH OF THE MANUFACTURED RFIC [SNM+2023].....	68
FIGURE 41: SINEWAVE SIGNAL MEASUREMENT SETUP AND 300 GHZ RECEIVER'S GAIN CURVE.....	69
FIGURE 42: MODULATED SIGNAL MEASUREMENT SETUP FOR 300 GHZ RECEIVER TESTING [SIN2024] ..	69
FIGURE 43: MODULATED SIGNAL MEASUREMENT SETUP FOR 300 GHZ RECEIVER TESTING [SPR+2025].....	70
FIGURE 44: BLOCK DIAGRAM AND PHOTOGRAPH OF 300 GHZ OTA MEASUREMENT SETUP AT UNIVERSITY OF OULU [HRM+2025]	70
FIGURE 45: REFERENCE SIGNAL STRUCTURE BASED ON INCREMENTAL CSI ESTIMATION	73
FIGURE 46: RF OUTPUT OF THE TEST SIGNAL AT 300 GHZ MEASURED BY UOULU AFTER DOWN CONVERSION FREQUENCY EXTENDER.....	74
FIGURE 47: RF SPECTRUM OF THE TEST SIGNAL AT 140 GHZ AS MEASURED BY IMEC.	74
FIGURE 48: BIT ERROR RATE PERFORMANCE OF PIECEWISE ESTIMATION AND EQUALIZATION IN A HIGHWAY NLOS CHANNEL MODEL AT 300 KM/H AND $M = 32$	75
FIGURE 49: BIT ERROR RATE PERFORMANCE OF PIECEWISE ESTIMATION AND EQUALIZATION IN A 3GPP TDL-C (300) CHANNEL MODEL AT 300 KM/H AND $M = 64$	76
FIGURE 50: NMSE AND SQUARED BOUND FOR HIGHWAY NLOS AND TDL-C (300) CHANNEL MODELS AT 120 KM/H AND 300 KM/H.	76
FIGURE 51: FIRST ADVANCED IMPROVED DRL RESULTS SHOWING THE EFFICACY OF DRL	79
FIGURE 52: AVERAGE SPECTRAL EFFICIENCY (SE) VERSUS NUMBER OF RIS ELEMENTS N FOR MU-MIMO AND MU-MISO SYSTEMS ACROSS LOW, HIGH, AND HYPER N REGIMES. RESULTS ARE SHOWN FOR RANDOM AND COSINE PHASE CONFIGURATIONS AT 140 GHZ AND 300 GHZ.....	80
FIGURE 53: SPECTRAL EFFICIENCY VS RIS ELEMENTS CALCULATED WITH DIFFERENT PHASE OPTIMIZATION APPROACHES SUCH AS DRL, COSINE ALGORITHM AND RANDOM PHASES.	82

FIGURE 54. TRIAL CONTROLLER ARCHITECTURE.....	84
FIGURE 55. USER LOGIN PAGE IN THE UNIFIED WEB PORTAL	86
FIGURE 56. USER SIGNUP PAGE IN THE UNIFIED WEB PORTAL	86
FIGURE 57. UNIFIED WEB PORTAL.....	87
FIGURE 58. UNIFIED WEB PORTAL- TRIAL DEFINITION	87
FIGURE 59. USER PROFILE IN THE UNIFIED WEB PORTAL	88
FIGURE 60. TRIAL MANAGEMENT– NORTHERN PORTAL CONNECTION.....	89
FIGURE 61. TRIAL MANAGEMENT–SOUTHERN PORTAL CONNECTION.....	90
FIGURE 62. USER MANAGEMENT ACROSS UNIFIED AND NORTHERN PORTAL	91
FIGURE 63. USER MANAGEMENT ACROSS UNIFIED AND SOUTHERN PORTAL	92
FIGURE 64. USER LOGIN PAGE IN THE NORTH WEB PORTAL	93
FIGURE 65. NORTH WEB PORTAL	94
FIGURE 66: NORTH WEB PORTAL-EXPERIMENT DEFINITION	94
FIGURE 67: USER PROFILE IN THE NORTH WEB PORTAL.....	95
FIGURE 68: EXPERIMENT MANAGEMENT IN THE NORTHERN WEB PORTAL	96
FIGURE 69: REGISTER FORM (LEFT) AND LOGIN FORM (RIGHT) OF THE SOUTH NODE WEB PORTAL	97
FIGURE 70: DASHBOARD SECTION OF THE SOUTH NODE WEB PORTAL	98
FIGURE 71: NST SECTION OF THE SOUTH NODE WEB PORTAL.	99
FIGURE 72. POP-UPS TO SHOW FURTHER INFORMATION OF THE EXPERIMENT AT THE SOUTH NODE...	99
FIGURE 73. FORM TO CREATE A NEW EXPERIMENT IN THE SOUTH NODE	100
FIGURE 74. EXPERIMENT MANAGEMENT WITH UNIFIED WEB PORTAL, SOUTH NODE WEB PORTAL AND SOUTH NODE ADAPTER	101
FIGURE 75 INTEGRATION BETWEEN UNIFIED WEB PORTAL AND SOUTH NODE PORTAL	102
FIGURE 76. REAL-TIME VISUALIZATION VIA QOSIUM AND GRAFANA.....	104
FIGURE 77. GRAFANA DASHBOARD IN THE NORTH WEB PORTAL.....	105
FIGURE 78. CONFIGURATION INTERFACE TO ENABLE THE SN EXPERIMENTER TO RETRIEVE RESULTS ..	106
FIGURE 79. EXAMPLE OF METRICS RETRIEVED THROUGH THE SN METRIC COLLECTION SYSTEM	106
FIGURE 80: ILLUSTRATION OF THE DECOMPOSITION OF THE TIME-VARIANT CHANNEL TAPS $b_{i,j}$ INTO A SET OF PIECEWISE TIME-INVARIANT CHANNELS WHOSE CHANNEL RESPONSES ARE ESTIMATED BY SAMPLING THE TIME-VARYING TAPS AT EACH OF THE DEFINED N_L SYMBOL INTERVALS.....	116

LIST OF TABLES

TABLE 1: THROUGHPUT, UPLINK AND DOWNLINK MEASUREMENTS (BANDWIDTH = 50 MHZ, XR GAMING TRAFFIC).....	19
TABLE 2: TESTED 3GPP RAN ENERGY SAVING METHODS/CONFIGURATIONS.....	30
TABLE 3: UTILIZED TEST TRAFFIC PATTERNS.....	31
TABLE 4: LIST OF MEASUREMENT ENVIRONMENT EQUIPMENT.....	34
TABLE 5: LIST OF NETWORK CONFIGURATION PARAMETERS FOR FR1 AND FR2.....	35
TABLE 6: LIST OF MEASUREMENT ENVIRONMENT EQUIPMENT AT UOULU 5GTN TEST FACILITY.....	35
TABLE 7: LIST OF NETWORK CONFIGURATION PARAMETERS FOR FR1 AND FR2 AT UOULU 5GTN TEST FACILITY	35
TABLE 8: LIST OF MEASUREMENT ENVIRONMENT EQUIPMENT AT VTT 5GTN TEST FACILITY	36
TABLE 9: LIST OF NETWORK CONFIGURATION PARAMETERS FOR FR1 AND FR2 AT VTT 5GTN TEST FACILITY	36
TABLE 10: SUMMARY OF MAIN CHARACTERISTICS AND ACHIEVEMENTS OF THE XR ENABLERS AND TRIAL CONTROLLER.....	107
TABLE 11: XAPPS AND THEIR FUNCTIONALITY	113
TABLE 12: MONITORING AND CONTROL MESSAGES.....	114

ABBREVIATIONS

3GPP Release X Rx	
6G	6 th generation mobile network
5G	5th generation mobile network
5GC	5G Core Network
5GTN	5G Test Network
3D	3 Dimensional
3GPP	Third Generation Partnership Project
ADC	Analog to Digital Converter
AF	Application Function
AMBR	Aggregate Maximum Bit Rate
AMC	Adaptive Modulation and Coding
ANBR	Access Network Bitrate Recommendation
API	Application Programable Interface
AS	Application Server
ATSSS	Adaptive Traffic Switching Steering Splitting
BLER	Block Error Rate
BSoTA	Beyond-state-of- the-art
BW	Bandwidth
CAD	Computer Aided Design
CCC	Cell Configuration and Control
CDF	Congestion Detection Function
CG	Cloud Gaming
C-Link	Control Link
CNC	Computerized Numerical Control
CO ₂	Carbon Dioxide
COTS	Commercial Off-the-Shelf
CP	Control Plane
CPU	Central Processing Unit
CSS	Cascading Style Sheets
CU	Central Unit
DAC	Digital to Analog Converter
DC	Data Channel
DCS	Data Channel Server
DDPG	Deep Deterministic Policy Gradient
DF	Decode and Forward
DL	Downlink
DNN	Deep Neural Network
DRL	Deep Reinforcement Learning
DRX	Discontinuous Reception
DSP	Digital Signal Processing
DT	Digital Twin
DTX	Discontinuous Transmission
DU	Distributed Unit
DWDM	Dense Wavelength Division Multiplexing
E-UTRA	Evolved Universal Terrestrial Radio Access
E2E	End-to-End
EARFCN	E-UTRA Absolute Radio Frequency Channel Number
EI	External Interface

eMBB	enhanced Mobile Broadband
EPC	Evolved Packet Core
ETSI	European Telecommunications Standards Institute
FDD	Frequency Division Duplexing
FoV	Field of View
FMI	Finnish Meteorological Institute
FPS	Framerate Per Second
FRx	Frequency Range x
gNB	5G NodeB
GPS	Global Positioning System
GPU	Graphics Processing Unit
GUI	Graphical User Interface
HMD	Head-Mounted Display
HW	Hardware
IAB	Integrated Access and Backhaul
IF	Intermediate Frequency
InH	Indoor Hotspot
ISAC	Integrated Sensing and Communications
KPI	Key Performance Indicator
KPM	Key Performance Measurements
L1/L2	Layer 1/Layer 2
LO	Local Oscillator
LoS	Line of Sight
MBB	Mobile Broadband
MIMO	Multiple-Input Multiple-Output
ML	Machine Learning
mmWave/mmW	Millimeter-wave
MQTT	Message Queuing Telemetry Transport
MR	Mixed Reality
MTU	Maximum Transmission Unit
NCR	Network Controlled Repeater
NCR-Fwd	NCR Forwarding
NCR-MT	NCR Mobile Termination
NG-RAN	Next Generation Radio Access Network
NI	Network Interface
NLoS	Non Line of Sight
NN	North Node
NR	New Radio
O2I	Outdoor-to-Indoor
O-CU	O-RAN Central Unit
O-DU	O-RAN Distributed Unit
O-RAN	Open Radio Access Network
O-RU	O-RAN Radio Unit
PAPR	Peak to Average Power Ratio
PRB	Physical Resource Block
PTP	Precision Time Protocol
QoE	Quality of Experience
QoS	Quality of Service
RAN	Radio Access Network
RB	Resource Block
RC	RAN Control

RF	Radio Frequency
RI	Research Infrastructure
RIC	RAN Intelligent Controller
RIS	Reconfigurable Intelligent Surface
RL	Reinforcement Learning
RMSE	Root Mean Squared Error
RRH	Remote Radio Head
RRM	Radio Resource Management
RSRP	Reference Signal Received Power
RSRQ	Reference Signal Received Quality
RT	Real Time
RU	Radio Unit
SC-FDE	Single-Carrier Frequency Domain Equalization
SCS	Subcarrier Spacing
SDR	Software-Defined Radio
SINR	Signal-to-Interference plus Noise Ratio
SMO	Service Management and Orchestration
SN	South Node
SNR	Signal-to-Noise Ratio
SotA	State-of-the-Art
SW	Software
T	Task
TCP	Transmission Control Protocol
TDD	Time Division Duplexing
THz	Terahertz
TRX	Transmit-Receive
UDP	User Datagram Protocol
UE	User Equipment
UI	User Interface
UL	Uplink
UMa	Urban Macro
UP	User Plane
UPF	User Plane Function
uRLLC	ultra-Reliable and Low Latency Communications
VR	Virtual Reality
VV	Volumetric Video
WP	Work Package
XR	Extended Reality
XRM	XR and Media Services

1 INTRODUCTION

XR as an umbrella term covering Virtual Reality (VR), Augmented Reality (AR) and Mixed Reality (MR), which provides immersive experiences from a blend of the virtual and real worlds and finds applications in several emerging use cases across different domains and sectors. XR requires performance enablers that can support stringent requirements with respect to low latency, high capacity, low power consumption and high reliability, among other KPIs [\[Paymard+2023\]](#).

This deliverable D4.3 introduces the final design done in the project, and final evaluation of identified new XR enabler technologies that have been identified in the project and documented in [\[D4.1\]](#) and [\[D4.2\]](#). Structure follows the categorization of XR enablers to three alternative paths from SoTA to 6G - 3GPP based evolution path, OPEN-RAN (O-RAN), and DISRUPTIVE. In addition to XR enablers contributing to three alternative paths, this document describes also the deployment of Trial Controller, which is implemented to ease the automatic deployment of XR-related experimentations in the two test facilities of the project (South Node and North Node) [\[D4.2\]](#).

1.1 OBJECTIVE OF THE DELIVERABLE

This deliverable has three main objectives for each new XR enabler and for the Trial Controller: (1) Justification of the value, or motivation, of each enabler supporting the respective path towards 6G, (2) Solution design of the implemented enabler giving the reader a sufficient level of details of the structure of each new asset to support the utilization of XR based services in a mobile environment and (3) final validation results of the deployments if available.

1.2 STRUCTURE OF THE DELIVERABLE

The 6G-XR project aims to demonstrate technological feasibility of innovative radio spectrum technologies of beyond 5G and 6G spectrum and validate an end-to-end beyond 5G architecture including end-to-end service provisioning supporting evolution from current architectures and with slicing capabilities, and at cloud implementation level (Open RAN).

Therefore, the deliverable D4.3 is covering the state of the art for RAN, beyond 5G RAN, core, and open-source networks, as well as disruptive RAN XR enabling technologies. Those enabler technologies are categorized to support alternative paths from 5G networks towards 6G as described above.

The deliverable D4.3 is structured as follows:

- Chapter 2 develops the five identified XR enablers following the evolution of 3GPP defined extensions for XR support in each 3GPP release.
- Chapter 3 focuses on developing three enablers relevant to O-RAN architecture, and utilizing customized xAPPs/rAPPs allowing modifying the RAN behaviour to maximize the operational efficiency based on the service needs.

- Chapter 4 develops three DISRUPTIVE enablers relevant for high-frequency solutions and RIS, thus assuming mobile connectivity for tailored network use.
- Chapter 5 develops the trial controller common design components and selected deployments to north and south experimentation sites.
- Chapter 6 concludes the document.

1.3 TARGET AUDIENCE OF THE DELIVERABLE

This deliverable is a public report which targets the project consortium, stakeholders, academic and research organizations, EU commission services, and the public.

2 3GPP XR ENABLERS

2.1 SUMMARY OF 3GPP XR ENABLERS IN D4.2

In order to fill some of the non-solved gaps in the XR capabilities of 5G New Radio (NR) in R15 and R16, [\[D4.2\]](#) presented five technology enablers supporting XR functionalities in 3GPP. [\[D4.2\]](#) described the interim enabler research results with the value or motivation of each enabler, explained the solution design and provided initial validation results of the deployments, if available.

This chapter completes the design and implementation results obtained with the five technical enablers by providing their final solution design, final evaluation results, and main takeaways and conclusions.

2.2 3GPP XR PROPOSED ENABLERS: DESIGN AND FINAL IMPLEMENTATION RESULTS

2.2.1 Network-Controlled Repeaters in indoor environment

2.2.1.1 Motivation

The adoption of 3GPP- Network-Controlled Repeaters (NCRs) were proposed from critical indoor deployment challenges which were detailed in deliverables [\[D4.1\]](#) and [\[D4.2\]](#). As documented in [\[D4.1\]](#), these adverse conditions notably diminish user experience (QoE) for XR services, requiring novel solutions beyond traditional deployments [\[Sousa2024\]](#) [\[Xu2022\]](#). To overcome such impediments, [\[D4.2\]](#) showcased the effectiveness of NCR deployment by employing advanced beamforming and dynamic network control, facilitating enhanced coverage and substantial improvements in Signal-to-Noise Ratio (SNR) and user throughput within indoor scenarios. [\[Goyal2025\]](#).

The motivation for adopting 3GPP Network-Controlled Repeaters (NCRs) stems from critical indoor deployment challenges addressed in the deliverables. Building upon the validated effectiveness documented in [\[D4.2\]](#), D4.3 introduces notable advancements and expansions in the NCR deployment scenarios. The entire workflow of the methodology is as shown in (Figure 1), major evolution pertains firstly to the use of advanced simulation parameters specifically optimized for Frequency Range 1 (FR1) bands, aimed at enhancing NCR performance and ensuring comprehensive validation at these widely utilized frequencies.

2.2.1.2 Final Solution design

As highlighted in Deliverable [\[D4.1\]](#), Indoor XR applications typically encounter substantial signal attenuation due to structural blockages, high Outdoor-to-Indoor penetration losses, and propagation constraints, especially prominent at higher frequency bands (FR2). This is further validated in [\[D4.2\]](#) through initial deployments using system-level simulations.

Advancements in subcarrier spacing (SCS), BW Scenario: Building on the intermediate results documented in [\[D4.2\]](#), D4.3 introduces significant enhancements in the NCR solution deployment, particularly leveraging advanced parameters in the FR1 band. A key evolution includes explicitly

utilizing a SCS of 15 kHz and 50MHz as Bandwidth, combined with XR-optimized traffic modelling, specifically using gaming traffic profiles available from standard 3GPP XR traffic models. This represents a marked departure from the previously used generic full-buffer traffic model, thus ensuring closer alignment with realistic XR service demands. Additionally, scenario realism has been improved through updating the simulation environment to the detailed Aveiro street map scenario, enabling more accurate assessment of NCR performance in practical urban conditions.

Scaling NCR parameters: In this phase, the i2CAT scenario is considered by modifying the NCR deployment setup. Specifically, the NCR position was adjusted to be slightly closer to the user, enabling a more accurate emulation of indoor deployment characteristics. This change supports more representative performance analysis, especially in dense XR environments where user proximity critically impacts signal quality and reliability.

Traffic Models: In this phase, we evaluate the performance of Network-Controlled Repeaters (NCRs) under Frequency Range 1 (FR1) with a bandwidth of 50 MHz and subcarrier spacing (SCS) of 15 kHz. The scenario is based on the Aveiro street map and assumes XR gaming traffic to reflect a high demand use case.



Figure 1: NCR methodology flowchart.

2.2.1.3 Final evaluation results

The simulations focused on enhanced bandwidth conditions—50 MHz, 100MHz and utilized various traffic types, reflecting realistic XR use-cases. The evaluated scenarios involved both NCR-enabled and baseline (without NCR) conditions, at specific outdoor-to-indoor (O2I) wall losses of 10 dB, highlighting challenging yet realistic indoor coverage conditions for XR applications.




Figure 2: NCR performance under the band FR1, BW=50MHz, SCS=15kHz, traffic model type-XR gaming, and Aveiro street city Map scenario.

The obtained results clearly demonstrate substantial performance gains through NCR deployment (Figure 2). The empirical cumulative distribution function (ECDF) plots indicate notable improvements in user throughput when NCRs were employed. Specifically, the throughput values reached up to approximately 31 Mbps in NCR-enabled conditions, while scenarios without NCR exhibited throughput levels considerably lower, primarily distributed below 25 Mbps. Additionally, both DL and UL SINR/SNR results reflected meaningful enhancements when NCR was activated. The DL SINR distributions shifted distinctly toward higher values (around 40 dB with NCR), indicating better quality and more reliable XR experiences. Similarly, UL SNR distributions displayed significant improvement, enhancing from predominantly negative or marginal values to consistently positive ranges with NCR deployment.

Graph Interpretation and NCR Performance Analysis (Percentiles of ECDF) is as explained in Table 1.

Table 1: Throughput, Uplink and Downlink measurements (Bandwidth = 50 MHz, XR Gaming Traffic).

Metric	Scenario	Median (50th percentile)	Higher Percentile	Interpretation
Throughput (XR Gaming)	Without NCR	~22 Mbps	90th percentile: ~30 Mbps	NCR significantly improves throughput, especially for users below median, enhancing XR gaming experience.
	With NCR	~30 Mbps	90th percentile: >31 Mbps	
Downlink (DL) SINR	Without NCR	~25 dB	70th percentile: ~30 dB	NCR deployment notably shifts SINR higher, enhancing signal quality for better XR user experiences
	With NCR	~32 dB	70th percentile: ~35 dB	


Uplink (UL) SNR	Without NCR	~8 dB	80th percentile: ~20 dB	NCR significantly improves indoor UL conditions, crucial for uplink-intensive XR services.
-----------------	-------------	-------	-------------------------	--

Results on Scaling and Traffic models.


By repositioning the NCR closer to the user in the i2CAT indoor scenario, we observed consistent performance improvements across key metrics. The ECDF plots (a) in Figure 3 demonstrate a clear increase in throughput and SNR for both uplink and downlink when the NCR is present, compared to the no-NCR baseline. The Block Error Rate (BLER), a critical metric that reflects the proportion of transmission blocks that are received with errors, shows a significant reduction across all users.

A lower BLER indicates higher link reliability and fewer retransmissions, which directly translates to more efficient use of network resources and improved Quality of Service (QoS) for latency-sensitive XR applications. These findings highlight how small adjustments in NCR positioning can lead to measurable gains in signal quality and user experience within realistic indoor deployment scenarios.

Under the traffic models we evaluated NCR performance, under different traffic conditions, two sets of traffic models were used. The first graph Figure 3 (b) was generated using standard traffic models available in the simulator, including full buffer, gaming, and video streaming configurations. These were used to establish a baseline for NCR behaviour under generic and controlled simulation environments. In contrast, the second graph presents NCR performance based on experimental data collected from real XR use cases, including volumetric video over TCP, holographic communication over IMS data channels, and high-rate collaborative VR applications (Table 2). These real-world traces introduce varying transport protocols, packet sizes, and data rates—ranging from 4.1 Mbps to 88 Mbps—providing a comprehensive view of NCR behaviour under realistic XR traffic. The comparison highlights the gap between analytical modelling and actual application demands, reinforcing the importance of incorporating experimental traffic profiles for accurate system evaluation.



(a) Results on scaling



(b) Results on traffic models

Figure 3: (a) Impact of Adjusted NCR Position on SNR and BLER under FR1 BW-(100 MHz, SCS = 15 kHz) NCR [-137.6145; -67.8227] in the i2CAT Indoor Scenario (b) Evaluating NCR performance with synthetic traffic models in the simulator and experimental traffic model data collected through use cases.

2.2.1.4 Takeaways and conclusions

These advanced simulations utilizing expanded bandwidth, scaling and specific traffic models confirm NCR's practical effectiveness in real-world XR scenarios, emphasizing its critical role in substantially boosting indoor XR network performance and ensuring robust QoE under demanding usage conditions.

2.2.2 Network assisted Rate Control API

2.2.2.1 Motivation

As outlined in [\[D4.1\]](#) and [\[D4.2\]](#), XR applications demand high data transmission rates in both uplink and downlink directions to ensure seamless user experiences. However, since radio resources are shared with general Internet users, the increasing XR traffic places significant strain on network capacity, often leading to congestion and subsequent Quality of Service (QoS) degradation. Existing XR services primarily rely on over-the-top (OTT) capacity estimation techniques to adjust media flow rates, aiming to optimize user experience and minimize buffering. This is typically achieved through adaptive bitrate algorithms that monitor metrics such as throughput, latency, and packet loss. However, these methods are insufficient for interactive XR applications due to their stringent latency and buffering constraints. To guarantee reliable and high-quality XR experiences, more intelligent and adaptive mechanisms are required.

Recognizing this challenge, 3GPP has envisioned the network as an intelligent platform capable of providing real-time insights through standardized APIs to enhance service performance. In alignment with this vision, the CAMARA project has developed APIs to optimize service delivery, with the CAMARA Quality on Demand (QoD) API emerging as a critical enabler for XR applications. Within the 6GXR framework, leveraging network-assisted mechanisms is essential to facilitate dynamic media rate adaptation. One of the key enablers in this context is the rate control API, which integrates real-time XR client feedback with network-provided capacity insights to optimize media rate adjustments, ensuring both efficiency and quality in XR service delivery.

2.2.2.2 Solution design

To bridge the limitations of conventional rate control mechanisms, we propose NetXRate, an API designed to enhance XR service rate adaptation by integrating explicit rate recommendations from the network. The architectural design of NetXRate is illustrated in Figure 4. Unlike over-the-top (OTT) rate control mechanisms, which rely solely on inferred network conditions, NetXRate introduces a direct communication channel between the XR service and the network. This approach enables the rate control function of XR applications to subscribe to real-time network-supplied rate recommendations, ensuring a more adaptive and informed media rate adjustment. The Network Exposure Function (NEF) of the 5G Core facilitates this subscription, granting seamless access to network insights. A key innovation of NetXRate is its ability to provide rate recommendations at configurable intervals, allowing XR services to receive dynamic updates on the maximum capacity available for a given session at any moment. These recommendations act as an upper bound rather than a strict enforcement mechanism, ensuring that existing OTT rate control algorithms remain operational while benefiting from enhanced network visibility. Instead of replacing traditional XR rate adaptation methods, NetXRate complements them by setting a maximum threshold that prevents excessive bandwidth allocation beyond what the network can sustain. This approach enhances XR performance, mitigates congestion, and ensures a more reliable, high-quality user experience.

The NetXRate architecture illustrated in Figure 4 includes the key 5G Core components, specifically the Network Data Analytics Function (NWDAF) and the NEF, to facilitate network-assisted rate control. NWDAF processes real-time network data to generate rate recommendations, while NEF exposes this information to XR services. A critical component of NetXRate is the NetXRate xApp, deployed within the Near-Real-Time RAN Intelligent Controller (Near-RT RIC). It continuously monitors available

capacity per cell and computes dynamic rate recommendations, ensuring XR traffic adapts to fluctuating network conditions. The NetXRate xApp interfaces with NWDAF through the O-RAN Y1 interface, enabling real-time RAN insights to inform NWDAF's analytics. These recommendations are then relayed through the NEF to the XR rate control function, ensuring that XR applications receive network-aware guidance for media rate adjustments. The overall NetXRate architecture integrates the Policy Control Function (PCF) where user priorities can be stored and managed. There are two options to enforce user priority either by CAMARA QoS where user profiles with QoS requirements are used or by using preconfigured priorities in the PCF. NetXRate considers the latter case in the per user rate recommendation.




Figure 4: Network Assisted rate control API architecture (NextXRate) with red arrow indicating the new proposed analytics.

2.2.2.3 Final evaluation results

The Network-Assisted Rate Control API has been evaluated through the ns-3 simulator integrated with the 5G-LENA module is employed under an urban macrocell environment operating at 3.5 GHz with a 100 MHz bandwidth and 30 kHz subcarrier spacing. A proportional fair scheduler is utilized, optionally enhanced with user priority weighting (e.g., XR:6, eMBB:1). The simulation spans 40 seconds across 9 seeds. XR traffic adopts dynamic adaptation strategies, while eMBB traffic includes HTTP, FTP, and video streaming applications following standardized traffic models. Users are clustered into groups (10 eMBB, 6 XR), with eMBB traffic distributed in a 6/2/2 pattern and XR subgroups transmitting at 6, 12, and 18 Mbps to emulate increasing frame rate demands. The simulations are to assess the maximum potential gains achievable when a tight integration between the RAN and the analytics module is established. The analysis conducted in this deliverable has focused on realistic network scenarios considering the coexistence of both XR and eMBB users within the same cell. Additionally, the benefits of NetXRate are examined in scenarios where XR users are prioritized by the core network. Two rate adaptation strategies called PacketSize and IPA (*Inter-Packet Arrival*) are considered, these strategies adjust the effective data rate through different mechanisms: the PacketSize strategy modifies the size of each packet, while the IPA strategy alters the packet generation interval.

Figure 5 and Figure 6 illustrate the performance trends of XR and eMBB traffic flows, respectively, as a function of the number of clusters, with and without XR priority. The analysis encompasses the following scenarios: (i) NetXRate with a PacketSize adaptation strategy (blue), (ii) OTT with a PacketSize adaptation strategy (red), (iii) NetXRate with an IPA adaptation strategy (purple), and (iv) OTT with an IPA adaptation strategy (yellow).

Figure 5 (a) illustrates that XR users achieve only 10.6% or 17% of their requested resolution or FPS, depending on the adaptation strategy employed. NetXRate with the IPA adaptation strategy outperforms OTT up to four clusters but adopts a more conservative approach with five clusters, prioritizing stability over throughput. Figure 5(b) shows that OTT experiences an increase in outages as the number of clusters grows, while NetXRate maintains near-zero outages, reducing them by up to 13 percentage points. Figure 5(c) demonstrates that OTT applications endure longer outage durations, with an average delay 270 ms greater than NetXRate, which significantly affects XR performance. Figure 5(d) presents the complementary cumulative distribution function (CCDF) of point-to-point delay, highlighting that NetXRate substantially reduces latency compared to OTT, even with the same adaptation strategy. At the 99.9th percentile, NetXRate reports delays of 11 ms and 19 ms for PacketSize and IPA, respectively, whereas OTT experiences delay of 213 ms and 287 ms. This improvement is attributed to NetXRate's rate recommendations, which effectively mitigate network saturation.

Figure 6Figure 6Figure 6Figure 6

In Figure 6 (a), NetXRate and OTT exhibit very similar normalized throughput, with only minor variations observed in scenarios with 3 and 5 clusters. The disadvantage of NetXRate observed in the non-prioritized case diminishes when XR priority mechanisms are applied. This improvement is likely attributable to the enhanced capacity allocated to XR users, which stabilizes their data rates. Figure 6(b) illustrates that NetXRate continues to achieve fewer outages than OTT. However, in this experiment, the increased capacity allocated to XR traffic also reduces the outage rate of OTT, narrowing the gap between the two approaches to approximately 6 percentage points at most. This finding is consistent with the results in Figure 6(c), where NetXRate shows higher outage durations in configurations with 3 and 4 clusters, contrasting with the previous experiment where NetXRate consistently achieved lower outage times.

When comparing internal variants of NetXRate, the IPA-based prioritization strategy outperforms the PacketSize-based approach in terms of outage duration and latency. This indicates that IPA more effectively leverages traffic prioritization to maintain service continuity. Conversely, the PacketSize approach demonstrates superior performance in normalized throughput and outage occurrence, suggesting a trade-off between maximizing throughput and minimizing service disruption. Delay trends in Figure 6(d) remain consistent with earlier findings: NetXRate achieves lower delays—84 ms and 110 ms at the 99.99th percentile—compared to OTT's 217 ms and 241 ms, respectively. These delay reductions are likely a result of the additional capacity provisioned for XR traffic, which mitigates severe network congestion.




Figure 5: Performance Comparison of NetXRate and OTT with Equal Priority for XR and eMBB Users.




Figure 6: Performance Comparison of NetXRate and OTT with XR Prioritized Over eMBB Users

The results depicted are summary of the provided results in the paper submitted to IEEE Transactions on Broadcasting.

2.2.2.4 Takeaways and conclusions

The next phase involves testing a simplified version of the API in the i2CAT laboratory, utilizing an Amarisoft gNB. This experimentation is critical for defining the API communication framework between the network and the holoportation application developed in WP3. Finally, the API and the holoportation application will be integrated and validated within the South Node infrastructure, ensuring seamless interoperability.

2.2.3 ATSSS based capacity enhancement

2.2.3.1 Motivation

As highlighted in [\[D4.1\]](#) and [\[D4.2\]](#), Extended Reality (XR) services, particularly holoportation, impose stringent bandwidth requirements in both uplink and downlink directions, necessitating advanced solutions to ensure seamless and high-quality real-time holographic communication. XR applications are predominantly utilized in indoor environments, where Wi-Fi 7 is expected to be the primary connectivity technology. However, its performance degrades under suboptimal coverage conditions, necessitating the integration of additional capacity from public 5G-Advanced/6G networks. To address this challenge, the Adaptive Traffic Steering, Switching, and Splitting (ATSSS) enabler facilitates the aggregation of Wi-Fi and 5G NR capacities within a unified multi-access PDU session, employing various traffic steering mechanisms. By incorporating ATSSS functionality within XR devices (e.g., Head-Mounted Displays and capture systems) and network infrastructure (e.g., User Plane Function), this approach enables dynamic traffic distribution while maintaining transparency to XR applications. Furthermore, ATSSS empowers Mobile Network Operators (MNOs) to implement tailored traffic steering policies to optimize network utilization for XR services. The principal research challenge lies in designing an efficient ATSSS forwarding function capable of enhancing service quality by maximizing bandwidth utilization and minimizing latency, thereby addressing a critical bottleneck in the large-scale deployment of XR communications.

2.2.3.2 Final Solution design

The envisioned 6GXR ATSSS architecture establishes a unified multi-access framework to enable seamless traffic steering, switching, and aggregation across indoor Wi-Fi and public 5G networks. At the core of the architecture lies the ATSSS User Plane (UP) function, responsible for distributing and routing IP traffic intelligently based on network conditions and policy directives. This function operates within the Customer Premises Equipment (CPE) and intermediate network elements, enabling dynamic path selection without disrupting ongoing sessions.

To optimize XR service delivery, the ATSSS API architecture shown in Figure 7, leverages flow-based traffic steering, ensuring each hologram in a holographic communication session is transmitted via the most suitable network path based on capacity estimation. This REST API provides a programmable interface for defining policy-based routing rules in Linux environments. Unlike traditional routing, which primarily considers destination IP addresses, this API enables routing decisions based on flexible criteria such as source and destination IPs, ports, and transport protocols (TCP/UDP). The system leverages *nftables*, ip rules, firewall marks, and multiple routing tables to route traffic according to configurable policy definitions. First, a configuration file is provided. It defines the relationship between a technology (WiFi/5G), a network interface and a firewall mark (*fwmark*).

Based on these details, the system is initially configured:

- For each technology, an ip rule is added to match packets with the assigned *fwmark*. The rule directs matched packets to a dedicated routing table.
- Each routing table is configured with a default route via the specified interface. This ensures that packets with the corresponding *fwmark* are forwarded through the correct network path.

With the system set up, the API is ready to receive requests. The API supports dynamic creation, modification, and deletion of routing policies via HTTP endpoints. When a new routing policy is created, the following policy enforcement mechanism is applied:

- A corresponding *nftables* rule is generated based on the policy definition.
- This rule marks matching packets with the associated *fwmark*.
- Rules are inserted into the prerouting chain to ensure marking occurs prior to any routing decision.

This mechanism enables real-time, rule-based traffic forwarding aligned with application-level or network-level logic.




Figure 7: ATSSS API architecture featuring rule-based traffic forwarding

2.2.3.3 Final evaluation results

The experiments were conducted in a controlled laboratory environment designed to emulate realistic yet repeatable multi-access conditions for XR traffic flows. Two Customer Premises Equipment (CPE) units were deployed: one emulating an XR user generating uplink-intensive traffic, and the other acting as a synthetic interferer injecting high-volume traffic to induce controlled congestion. The experimental testbed integrated two heterogeneous access technologies: a Wi-Fi network and a 5G Standalone (SA) network deployed using an Amarisoft gNB connected to an Open5GS core network.

Both CPEs were capable of multi-access connectivity, enabling support for 3GPP's ATSSS functionalities.

In the initial phase, as illustrated in Figure 8, the XR user was connected to the Wi-Fi network and consistently achieved uplink throughput in the range of 40–50 Mbps—sufficient for immersive XR streaming. This served as the baseline throughput under uncongested conditions. The interferer was then activated to generate high-rate uplink traffic on the same Wi-Fi channel. Due to contention-based medium access, significant congestion occurred, leading to a substantial degradation in XR user throughput.

Upon detecting that the throughput dropped below a predefined quality threshold, a traffic switching event was triggered as per ATSSS policy rules. The XR traffic was seamlessly migrated from the congested Wi-Fi access to the 5G SA network. This switching process was executed through soft handover logic to ensure minimal service disruption. Following the transition, the XR user's uplink throughput recovered, validating the capacity of the 5G link to handle the XR session in its then-uncongested state.

To further evaluate the adaptability of the ATSSS mechanism, the interferer was switched to the 5G network. As uplink demand increased, the gNB's Physical Resource Block (PRB) utilization escalated, inducing observable congestion effects. This was reflected in a gradual decline in XR user throughput, again breaching the service quality threshold. Consequently, another ATSSS-based switching decision was made, redirecting the XR traffic back to the Wi-Fi access. This reverse handover successfully restored uplink performance, demonstrating the dynamic responsiveness of the ATSSS framework.




Figure 8: Comparative analysis of CPE uplink throughput over time in WiFi and 5G access networks

2.2.3.4 Takeaways and conclusions

These results underscore the effectiveness of ATSSS traffic steering and switching in maintaining service quality for latency- and bandwidth-sensitive XR applications across heterogeneous access networks. The ability to dynamically respond to changing congestion conditions demonstrated that ATSSS-driven traffic switching enables effective mitigation of uplink throughput degradation caused by localized congestion. The key conclusion is that seamless and policy-based transitions between Wi-Fi and 5G access networks, as enabled by the ATSSS framework, significantly enhance service robustness and continuity for uplink-sensitive applications like XR.

Future work will focus on advancing the traffic steering decision logic through predictive models such as reinforcement learning, incorporating multi-dimensional QoS metrics (e.g., latency, jitter, PRB utilization), and extending the approach to multi-user, multi-session environments. These enhancements aim to further optimize ATSSS responsiveness and scalability in real-world 5G deployments.

2.2.4 Energy and Application aware management of 3GPP infrastructure

2.2.4.1 Motivation

The transition times between gNB/cell level energy conservation methods in the current SoTA 3GPP-compliant RAN systems are usually in the order of $\gg 1s$ to minutes (slow control) depending on which methods are wanted to be activated or deactivated. Even if they are not very well suited for dynamic resource optimisations, they provide more significant energy saving potential as they operate on higher hierarchical level than the more dynamic methods. A handful of RAN configurations or energy saving states have been selected for testing. The XR applications from 6G-XR use cases UC1-4 are emulated and the impact of the different RAN energy saving states to the network's capability to support the related services are evaluated through measurements. As the time for the different configuration changes to take effect in the RAN and their impact on the overall energy consumption are known through WP5 tests and results, the feasibility of using different energy conservation methods is assessed from the point of view of the 6G-XR use case applications.

2.2.4.2 Solution design

Figure 9 presents the test setup used for the final measurements. The scalability of XR applications is assessed by measuring the maximum throughput of a single capacity or coverage cell using the iPerf3 [iPerf2025] traffic generator software. The generated test traffic emulates the key applications used in the 6G-XR use cases UC1-3 – Real-time Holographic Communications as well as in the use case UC4 – Collaborative 3D Digital Twin-like Environment. After measuring the maximum throughput achievable with the emulated application traffic patterns, the scalability of each 6G-XR use case is assessed by estimating the maximum number of simultaneous XR users supported by the different RAN energy saving configurations.

In the test setup in Figure 9, the 3GPP Next Generation RAN (NG-RAN) component are drawn with blue colour, the 5G Core Network (5GC) components are drawn with dark magenta, and network management platform is drawn with purple. The RAN state controller used to change the RAN configuration is drawn with yellow colour. It is a component developed in WP5 and utilised in the tests. The different databases containing energy and KPI measurement data as well as the traffic models based on the 6G-XR use cases are drawn with orange colour.




Figure 9: 3GPP NG-RAN test setup for energy and application aware resource control.

As discussed in [D4.2], based on the initial tests performed in the 6G-XR use case UC5 – Energy Measurement Framework for Energy Sustainability, the most suitable and effective energy saving methods in the 3GPP RAN infrastructure (including one coverage cell and one capacity cell) are Discontinuous Transmission (DTX), shutting down active transmitters (also known as MIMO muting), and putting the capacity cell to sleep mode (light or deep sleep), which results in a RAN infrastructure where only the coverage cell is active. The coverage cell utilised in the experiment operates at the n7 5G frequency band (FDD, BW: 10 MHz, MIMO: 4x1) and the capacity cell at the n78 5G frequency band (TDD, BW: 60 MHz, MIMO: 4x2). Both gNBs are installed into the same mast and the cells are overlapping. Intermediate results in WP5 indicate that each of the listed energy saving methods can potentially decrease the overall energy consumption at the RAN infrastructure level by additional 10 % when compared to a baseline configuration with no dynamic energy saving methods activated as shown in Table 2. The configuration changes between the RAN energy saving states are performed so that the methods on higher rows in Table 2 stay active when new methods from the lower rows are activated, i.e., DTX stays on when automatic shutdown of active transmitters is activated as an additional energy saving method. More detailed results on the actual achieved energy savings in the North Node test facilities will be reported in the 6G-XR deliverable D5.2 in the end of the project.

Table 2: Tested 3GPP RAN energy saving methods/configurations

Energy saving method	Power saving potential
Baseline configuration – No RAN energy saving	Baseline
Discontinuous Transmission – DTX	10 % saving compared to baseline
Shutting down active transmitters – MIMO muting	20 % saving compared to baseline
Putting a capacity cell to sleep mode (light or deep sleep) – Coverage cell only	30-50 % saving compared to baseline

Table 3 lists the XR applications emulated in the tests, maps them to the 6G-XR use cases, and provides the basic characteristics of the data traffic exchanged between the UE and the network. The traffic parameters have been attained by recording and analysing the data traffic traces of the real applications developed in WP3 and WP6. More detailed results on the final development versions of these applications will be reported in the 6G-XR deliverables D6.1 and D6.2 in the end of the project.

Table 3: Utilized test traffic patterns

Related use case	Emulated XR application	Data traffic characteristics (per user)
UC1-2 – Real-Time Holographic Communications (user plane optimizations)	Volumetric Video (VV)	Transport protocol: TCP Avg. packet size: 26000 B Avg. packet rate: 50 packets/s Avg. data rate: 10.4 Mbps
UC3 – Real-Time Holographic Communications (control plane optimizations)	Holographic call over IMS Data Channel	Transport protocol: UDP Avg. packet size: 1000 B Avg. packet rate: 510 packets/s Avg. data rate: 4.1 Mbps
UC4 – Collaborative 3D Digital Twin-like Environment	Virtual Reality (VR)	Transport protocol: UDP Avg. packet size: 1310 B Avg. packet rate: 8400 packets/s Avg. data rate: 88 Mbps

2.2.4.3 Final evaluation results

Figure 10 presents the measured maximum throughput values for the different 6G-XR use cases with different energy saving methods activated in the 3GPP RAN infrastructure containing the single capacity cell and a coverage cell. Each throughput value presented in the chart is the highest measured throughput after averaging over the whole duration of a three-minute test run. During the experiments, the maximum throughput value is defined by gradually increasing the number of emulated XR users for the different use cases step-by-step according to the traffic parameters show in Table 3. This means that, e.g., for UC4, each step added 88 Mbps (1310 B data packets and 8400 packets/s) to the overall load offered to the network. For UDP-based applications, the maximum is achieved when the throughput saturates and the one-way latency in the air interface starts to increase rapidly. For TCP-based applications, the congestion control mechanisms of the protocol throttle and saturate the throughput before the one-way latency is severely affected.




Figure 10: Measured maximum throughputs for 6G-XR UC1-4 with different 3GPP RAN energy saving states.

It can be seen from Figure 10 that the different transport layer protocol and packet sizes have a clear impact to the maximum throughput that can be achieved in the DL direction when no energy saving methods are active in the RAN infrastructure. For UC1-2, the maximum throughput of 448 Mbps is mainly limited by TCP utilised as the transport layer protocol. For UC3 and UC4, which both use UDP as the transport layer protocol, the difference between their respective maximum throughputs of 519 Mbps and 545 Mbps is caused by the difference in the data packet size. Smaller packet size adds both protocol and processing overhead into the system. It should be noted that the Maximum Transmission Unit (MTU) size that can be used in the VTT 5GTN test facility infrastructure is 1400 B, which means that the TCP packet size for UC1-2 in the experiments is smaller than indicated in Table 3.

As expected, the maximum DL throughput for each use case starts to decrease when more energy saving methods are activated in the RAN. However, the drops in the throughputs differ between the use cases. When DTX is activated, the relative drop in the UC1-2's throughput (4 %) is much smaller than the drop in the UC3's throughput (25 %) and UC4's UDP throughput (28 %). However, when MIMO muting is also activated, the relative drop in the UC1-2's throughput (20 %) is in turn larger than the drop in the UC3's throughput (13 %) and UC4's throughput (9 %). This indicates that the transport layer protocol used by the XR application is a factor when it comes to the impact of the different RAN energy saving methods to the application layer performance.

When the capacity cell is put to sleep and the XR users are served by the low-capacity coverage cell, the throughputs achieved for each use case are practically the same. As the tested energy saving methods are not impacting the UL data transfer, the same can be said for the UL throughputs of all use cases in both the capacity cell and coverage cell.

Figure 11 provides the estimated amounts of XR users the measured maximum throughputs are able to support of the different use cases. The estimate is calculated simply by dividing the measured maximum throughput values from Figure 10 with the data rates generated by individual XR users shown in Table 3. When assessed from the service scalability perspective, the differences between the use cases are much larger than the maximum throughput values suggested. The relative drops in the user amounts also differ greatly from the drops in the throughput. This highlights the need for use case or even application specific consideration when assessing the trade-off between the service or user

needs and RAN energy savings in real networks. In addition to the network load, it is also important to know what kind of applications are generating the load so that the RAN control algorithms can reliably estimate the real impact of the changed energy saving states into the service quality and user experience offered to the customers.




Figure 11: Estimated number of simultaneous XR users supported in different 3GPP RAN energy saving states.

2.2.4.4 Takeaways and conclusions

The performed measurements revealed interesting differences between the 6G-XR use cases when it comes to the achieved throughputs and use case scalability in different RAN energy saving states. The results highlight the importance of both energy and application awareness when making optimisation decisions for the infrastructure with configuration transition times spanning from tens of seconds to minutes. If a wrong optimisation decision is made and the configuration change executed, it takes a long time to correct the situation by reverting to the previous configuration. This can potentially cause serious service interruptions for the XR users served by the network.

2.2.5 Upgrade and evaluation of experimental RAN infrastructure in SN and NN

2.2.5.1 Motivation

In this section we summarize the 3GPP networks that have been deployed in the campus environments of the South Node (SN) and the North Node (NN), while providing a performance evaluation under realistic radio conditions. These results are expected to guide the design of the experiments that will be executed in the SN and the NN as part of the Open Call 3.

2.2.5.2 Deployed network architecture and proposed measurement methodology

2.2.5.2.1 3GPP SN deployment

Following the South Node depicted in [D4.1] and [D4.2]. The measurement campaign aims to evaluate network coverage and throughput at key locations by conducting performance validation tests on FR1 and FR2 cells. Measurements include TCP/UDP upload and download throughput and RTT latency, using iPerf3 and ICMP tests. The tests are performed at selected locations, with results averaged over multiple repetitions to ensure reliability. Data is exported for analysis, comparing performance trends

based on distance from the antenna to assess coverage and capacity. The target test locations for FR1 shown in Figure 12 are LAMPPOST1, LAMPPOST 2, FIB, and Vertex building while for FR2, the tests are conducted at the entrance, indoor and in front of i2cat building. The measurement tool is iPerf. Each test is performed 5 times per location then each metric is computed by the average per location (time-averaged and across 5 samples). The measurement tests are RTT ICMP tests (for latency evaluation), iPerf TCP Download (DL) throughput, iPerf TCP Upload (UL) throughput, iPerf UDP Download (DL) throughput, iPerf UDP Upload (UL) throughput, and one way latency with Qosium (DL+UL).




Figure 12: Map of target test locations with red squares showing the location of the antennas

- **Measurement Environment:**

Table 4: List of measurement environment equipment

5G CPE:	Askey NUQ3000M
5G gNB:	Ericsson gNB flight rack
5G Core:	Connected via VPN to 5TONIC
Edge Server:	5TONIC (Located in Madrid) and i2CAT (Located in Barcelona)

- **Network configurations:**

Table 5: List of network configuration parameters for FR1 and FR2

Parameter	FR1 (Mid-Band - n77)	FR2 (High-Band - n258)
Frequency Range	3700–3980 MHz	26 GHz
Band Name	n77	n258
Duplex Mode	TDD	TDD
TDD Pattern	DDDSU	DDDSU
Bandwidth Allocation	40 MHz	400 MHz (4 x 100 MHz carrier)
Power per Carrier	5 W	2.5 W per carrier

The list of KPIs measured are Latency, Uplink throughput, and Downlink throughput.

2.2.5.2.2 3GPP NN deployment

The North Node measurement campaign is based on the plans presented in [\[D4.2\]](#) chapter 2.2.5.4. Measurements focus on the verification DL capacity, UL capacity, and low communication latencies utilising the new gNBs installed into the North Node test facilities to enable optimum parameterisations for different use cases. Both FR1 and FR2 tests will be implemented. Data throughput capacity is measured on L3/IP-level for one-way delay separately for both DL and UL. Tests will be executed in different locations across both VTT and University of Oulu campus areas.

Similarly to the South Node tests, the utilised measurement tool is iPerf3. Each test is performed 5 times per location and each metric presented in the results is the averaged value for each location (time-averaged and across 5 samples). The measurement tests are iPerf3 TCP DL throughput, iPerf3 TCP UL throughput, iPerf UDP DL throughput, iPerf UDP UL throughput, and one-way DL and UL latency with 10 Mbps UDP stream. All measurements are done using the Qosium [\[Qosium2025\]](#) measurement tool.

UOULU 5GTN

Table 6: List of measurement environment equipment at UOULU 5GTN test facility

5G CPE:	Quectel RM500
5G gNB:	Nokia gNBs
5G Core:	Cumucore
Edge Server:	Local edge server at UOULU premises

Table 7: List of network configuration parameters for FR1 and FR2 at UOULU 5GTN test facility

Parameter	FR1 (Mid-Band - n77)
Frequency Range	3800-4000 MHz
Band Name	n77
Duplex Mode	TDD
TDD Pattern	DDDSU
Bandwidth Allocation	200 MHz (total), 100 MHz (used)
Power per Carrier	5 W (outdoor), 0.15 W (indoor)



Figure 13: FR1 outdoor and indoor installation, UOULU 5GTN

Figure 13 represents the location of both indoor and outdoor (FR1 n77) radio installations. Our outdoor deployment consists of three sectors, ensuring comprehensive area coverage. The outdoor radios are 64 TRX massive MIMO Nokia radios units. Indoor coverage is supplemented by Nokia Picos (AWHQW) 4T4R radio installations in labs and server rooms.

VTT 5GTN

The network and user equipment used during the VTT 5GTN test facility performance verification measurements is listed in Table 8. The network configuration parameters for the tested FR1 and FR2 gNBs are presented in Table 9.

Table 8: List of measurement environment equipment at VTT 5GTN test facility

5G CPE:	Sierra Wireless EM9293 (FR1) and Asus Snapdragon for Insiders (FR2)
5G gNB:	Nokia gNBs
5G Core:	Open5GS
Edge Server:	Local edge server at VTT Oulu premises

Table 9: List of network configuration parameters for FR1 and FR2 at VTT 5GTN test facility

Parameter	FR1 (Mid-Band - n77)	FR1 (Mid-Band - n78)	FR2 (High-Band - n258)
Frequency Range	3800-4000 MHz	3510-3570 MHz	26 GHz
Band Name	n77	n78	n258
Duplex Mode	TDD	TDD	TDD
TDD Pattern	DDDSU	DDDSU	DDDSU
Bandwidth Allocation	200 MHz (total), 100 MHz (used)	60 MHz	800 MHz (DL), 400 MHz (UL)

Power per Carrier	0.25 W (indoor)	5 W (outdoor)	0.125 W (outdoor)
-------------------	-----------------	---------------	-------------------

Figure 14 presents the locations of the tested gNB sites (i.e., installation locations of the gNB antennas) with small triangles. The green triangles represent the FR1 gNBs and red triangles the FR2 gNBs. The green and red crosses in Figure 14 mark the locations where the UEs were during the FR1 and FR2 measurements, respectively. All FR1 measurements were performed from the indoor lab space near to the FR1 gNBs. The antennas of the FR1 cell #1 operating at the n77 frequency band are installed indoors in the lab where the UE is also located, whereas the antennas of the FR1 cell #2 operating at the n78 frequency band are installed outdoors on the roof of the VTT building and pointing roughly to north. Both FR2 cells #1 and #2 operating at the n258 frequency band have their antennas installed outdoors next to the parking lot and pointing north. In all measurement locations (lab, parking lot, and sidewalk), the received signal strength and quality at the UEs was excellent with recorded values of Reference Signal Received Power (RSRP) = ≥ -73 dBm, Reference Signal Received Quality (RSRQ) ≥ -11 dB, and Signal-to-Interference-plus-Noise Ratio (SINR) ≥ 23 dB.




Figure 14: FR1 and FR2 gNB antennas (triangles) and measurement locations (crosses) at VTT 5GTN test facility.

2.2.5.3 Final network evaluation

2.2.5.3.1 SN performance evaluation

In Figure 15, the downlink (DL) and uplink (UL) throughput performance over FR1 was evaluated using iPerf-based measurements at two distinct locations within the test area: LAMPPOST1, situated near the radio access point, and the Vertex Building, being the furthest measurement point from the antenna system. At LAMPPOST1, the measured DL throughput achieved 420 Mbps for TCP and 320 Mbps for UDP, indicating a robust and stable radio link near the antenna. The UL throughput at the same location was recorded at 40 Mbps for TCP and 38 Mbps for UDP, demonstrating symmetrical uplink performance close to the theoretical limits under favorable radio conditions. Conversely, at the Vertex Building, a significant degradation in throughput was observed due to the increased path loss and reduced signal quality at this distant location. The DL throughput dropped to 40 Mbps for TCP and

20 Mbps for UDP, while the UL throughput deteriorated further, falling below 1 Mbps for both TCP and UDP traffic. This sharp decline is attributed primarily to the lower received signal strength and potential interference or fading effects encountered at greater distances from the transmission source.

For the evaluation of FR2 (mmWave) performance, the same set of iPerf-based throughput tests were conducted as in the FR1 case. The measurement campaign focused on two distinct coverage sectors, denoted Sector 1 (S1) and Sector 2 (S2), as illustrated in Figure 16.

- Sector 1 (S1) includes three representative test locations:
 - the entrance of the i2CAT building (outdoor),
 - the indoor area within the i2CAT premises,
 - and the pedestrian crossing adjacent to the building.
- Sector 2 (S2) corresponds to the Tables area; an open space located further from the building and oriented differently with respect to the antenna beams.

The system under test employed carrier aggregation across four 100 MHz carriers, resulting in an aggregated bandwidth of 400 MHz in the FR2 band. All throughput results presented reflect this aggregated configuration.

At the i2CAT entrance (S1), the downlink (DL) throughput reached 2.16 Gbps for TCP and 2.18 Gbps for UDP, demonstrating near-peak performance under line-of-sight (LoS) and favorable channel conditions. Similarly, at the Tables area (S2), the DL throughput was slightly lower, with 2.01 Gbps for TCP and 2.04 Gbps for UDP, suggesting minor degradation due to orientation, distance, or beamforming alignment.

In contrast, when the receiver was placed indoors (S1)—representing a non-line-of-sight (NLoS) environment—the DL throughput experienced a marked reduction, reaching 0.69 Gbps for TCP and 0.63 Gbps for UDP, respectively. This significant drop is attributable to high penetration loss associated with FR2 frequencies and the lack of reflective paths or efficient beam coverage indoors.

The Uplink (UL) throughput also exhibited strong performance at outdoor locations. At the entrance (S1), the UL throughput reached 411 Mbps for TCP and 420 Mbps for UDP, while at the Tables area (S2) it measured 360 Mbps for TCP and 373 Mbps for UDP. These results validate the uplink capability of FR2 systems under aggregated bandwidth conditions in well-aligned outdoor scenarios.

In summary, the FR2 throughput measurements demonstrate the substantial capacity of mmWave systems in open, LoS environments, while also highlighting the inherent challenges in indoor or NLoS scenarios due to severe path loss and signal blockage.




Figure 15: Downlink and Uplink Throughput Saturation Benchmark over FR1 at varying distances from the antenna.




Figure 16: Downlink and Uplink Throughput Saturation Benchmark over FR2 at varying distances from the antenna.

Latency Results

To characterize the latency performance of the wireless network under evaluation, a series of one-minute measurement sessions were conducted at a fixed physical location. The primary objective was to quantify end-to-end latency in downlink (DL) directions under controlled traffic conditions for FR1.

Traffic generation during each session was achieved using the iPerf tool, configured to transmit a UDP stream at 10 Mbps. This fixed-rate configuration ensures a uniform traffic load across all test sessions, thereby isolating the effects of network dynamics on latency behaviour without introducing variable-rate congestion artifacts.

To monitor the latency experienced during these transmissions, the *Qosium* probe-based measurement platform was employed. Qosium operates by passively tracing network traffic and timestamping each transmitted and received packet, enabling precise latency estimation on a per-sample basis. At each location, UDP test session of 10Mbps is conducted for duration of 1 minute. This is repeated for three times. Following data collection, post-processing was conducted to extract the average one-way latency for each individual test session. Subsequently, to obtain the representation of network performance at the measured location, an *average latency* was computed by aggregating the latency values across all *Test* sessions conducted at that location.

As illustrated in the downlink measurements shown in Figure 17, although with a less pronounced variation. At LAMPPOST1, the downlink latency is recorded at 6.952 milliseconds, which slightly increases to 7.167 milliseconds at Vertex. These results suggest that the downlink latency remains relatively stable, indicating more robust performance likely attributed to higher transmit power and optimized scheduling on the network side.




Figure 17: Comparative analysis of Downlink latency at different locations

2.2.5.3.2 NN performance evaluation

In this subsection, the North Node performance evaluation results are presented first for the UOULU 5GTN test facility and then for the VTT 5GTN test facility site.

UOULU 5GTN




Figure 18: Max throughput measurements for FR1 (n77)

Maximum throughput measurements are shown in the Figure 15. We have both indoor and outdoor coverage for our FR1 n77 base station. Our UE was positioned closer to the indoor pico, it is likely that UE was connected to this pico cell due to its proximity. Both UDP and TCP traffic were generated for network performance evaluation for different traffic conditions and potential overload scenarios. Maximum DL throughput for TCP traffic was 922.51 Mbps and 784.08 Mbps for UDP traffic. Similarly maximum UL throughput for TCP traffic was 156.4 Mbps and 153.6 Mbps for UDP traffic.




Figure 19: Delay measurements for slice/channel overload (FR1)

Figure 16 shows delay measurements for slice overload scenario. Two UEs were connected to the same slice and one of them used to overload the slice. Average DL delay is 7.8 ms for 10 Mbps UDP traffic and 17.5 ms for the overloaded channel. While average UL delay is 11.12 ms 10 Mbps UDP traffic and 24.47 ms for the overloaded channel.

VTT 5GTN

Figure 20 presents the maximum throughput values measured at the VTT 5GTN test facility. As discussed in subsection 2.2.5.2.2, all measurements except the ones performed against the FR1 n78 gNB were done in LoS conditions, and even the NLoS measurements had excellent signal strength and quality recorded at the UE. Hence, the presented results can be considered as the best case practical 5G performance achievable at the VTT 5GTN test facility in its current state of deployment. Each throughput value presented in the chart is the average value over five test runs each lasting five minutes, except FR2 n258 results which are averaged over two five-minute test runs, one performed in the parking lot and one in the sidewalk location marked in Figure 14.




Figure 20: Maximum throughput measurements for TCP and UDP traffic at VTT 5GTN test facility.

In FR1, the measured maximum TCP throughputs in the DL direction are 940 Mbps and 760 Mbps for n77 and n78, respectively. In the UL direction, the throughputs are 150 Mbps for n77 and 89 Mbps for n78. Without the congestion control functionality of TCP helping to throttle the offered load, the achieved averaged UDP throughputs in the DL direction are lower with the measured maximum of 821 Mbps for n77 and 667 Mbps for n78. In the UL direction where scheduling of the data transmission slots is done by the gNB based on requests by the UE, the difference between TCP and UDP is negligible. For n77, the maximum UDP throughput in the UL direction is 152 Mbps and for n78 it is 84 Mbps.

In FR2 at the n258 frequency band, the measured maximum TCP throughput in the DL direction is over 2.8 Gbps. For UDP, the measured maximum was approximately 2.5 Gbps. In the UL directions, the measured maximum throughputs were 214 Mbps for TCP and 209 Mbps for UDP. It should be noted that the peak throughput values especially for the DL direction are significantly higher reaching values of 3.5 Gbps. The main limitation of the current set up is that used FR2 UE (Asus smartphone) is not able to sustain the maximum throughput performance throughout longer test runs (> 15 minutes) due to overheating issues (Figure 20). Use of a larger form factor UE device with better cooling will make it possible to utilize the full capacity of the FR2 connectivity for long continuous data streaming tests. This will make it possible to utilize the full capacity of the FR2 connectivity for long continuous data streaming tests.

Figure 21 presents the measured one-way latencies for the FR1 and FR2 cells at the VTT 5GTN test facility. All presented values are measured with 10 Mbps UDP data traffic. Each latency value presented in the chart is the average value over five test runs each lasting five minutes, except FR2 n258 results which are averaged over two five-minute test runs, one performed in the parking lot and one in the sidewalk location marked in Figure 14. The measurements were done between the UE and a network edge server at VTT Oulu premises. The measurement end points are time synchronised using LinuxPTP [\[LinuxPTP2025\]](#) and hardware-based timestamping, which results in μ s-level synchronisation accuracy in the utilised test facility infrastructure.




Figure 21: One-way latency measurements for UDP traffic at VTT 5GTN test facility.

In FR1, the measured average DL latency is 7.6 ms for both n77 and n78. The measured average UL latency is 13.1 ms for n77 and 12.8 ms for n78. In FR2 at the n258, the values are 6.5 ms in the DL direction and 11.3 ms in the UL direction.

2.2.5.4 Takeaways and conclusions

The results provided in this section demonstrate the state-of-the-art performance available in the test facilities setup by the 6G-XR project. Multi-gbps peak data rates and sub-10 ms one-way latencies are available in all the 6G-XR test facilities, which provide a forward-looking baseline to validate future XR services.

3 O-RAN XR ENABLERS

3.1 SUMMARY OF O-RAN ENABLERS IN D4.2

Deliverable [\[D4.2\]](#) presented three O-RAN technology enablers for RAN optimization implemented as rApps or xApps, that can be used to enhance the performance of XR applications for different verticals, in relation to the 6GXR use cases. The analysis included a motivation, an initial solution design, and the initial results obtained for each of them.

This chapter completes the solution design, evaluation results, and main takeaways and conclusions of the three O-RAN enablers.

3.2 O-RAN XR PROPOSED ENABLERS

3.2.1 Congestion aware load balancing

3.2.1.1 Motivation

As highlighted in [\[D4.1\]](#) and [\[D4.2\]](#), the holoportation service in 6GXR Use Case 1 imposes significant capacity demands in both uplink and downlink, necessitating mechanisms to ensure service quality in O-RAN-enabled public networks. To address this challenge, the 6GXR congestion-aware load balancing enabler leverages O-RAN capabilities to optimize network resource allocation dynamically. In a typical deployment, holoportation users, primarily located indoors, share network resources with best-effort Internet users across multiple cells. While conventional cell selection relies on signal quality metrics (e.g., RSRP/RSRQ), Mobility Load Balancing (MLB) can further enhance network efficiency by steering users away from congested cells. The proposed solution integrates gNBs with a near-real-time RAN Intelligent Controller (near-RT RIC), utilizing the standardized interface E2 specified by the O-RAN Alliance that enables near-real-time control, monitoring, and optimization of RAN nodes (e.g., gNB, O-DU, O-CU) by the Near-RT RIC via service models. Specifically the E2 Key Performance Measurement (KPM) service model for congestion estimation and the E2 RAN Control service model for intelligent handover decisions. A dedicated congestion estimation rApp monitors PRB utilization across cells to detect potential congestion, while a load balancing rApp facilitates user handover to less congested cells, ensuring optimal holoportation service performance. This framework not only enhances XR service quality but also enables network-wide load balancing for all users, improving overall resource efficiency in O-RAN-enabled networks.

3.2.1.2 Solution design

The proposed congestion-aware load balancing mechanism optimizes radio resource allocation for XR services in O-RAN by dynamically mitigating network congestion. The Congestion Detection Function (CDF) rApp monitors PRB utilization, detects congestion, and triggers proactive load balancing, embedding congestion reporting for faster decision-making. The UE Location rApp tracks real-time UE positions, while the Handover Execution xApp running in the Near RT RIC ensures seamless transitions. Leveraging O-RAN's near-real-time RAN Intelligent Controller (nRT-RIC), the solution enhances network efficiency and reduces service disruptions. Experimental validation confirms improved handover execution and minimizes XR service interruptions, with future work focusing on predictive congestion management and scalability.

The proposed congestion-aware load balancing solution consists of the following core components:

3.2.1.2.1 CDF implementation

3.2.1.2.1.1 Reactive CDF

The **Congestion Detection Function (CDF) rApp** is a key component within the RAN Intelligent Controller (RIC) ecosystem, responsible for detecting and mitigating congestion in real time. By continuously monitoring Physical Resource Block (PRB) utilization via the standardized E2 Key Performance Measurement (KPM) service model, the CDF rApp identifies congestion scenarios that may degrade network performance. Once congestion is detected, the rApp autonomously makes load-balancing decisions to redistribute traffic across available resources, ensuring optimal service quality. Additionally, it incorporates built-in congestion reporting capabilities, eliminating reliance on external systems and streamlining the congestion management process within the near-real-time RIC (nRT-RIC) framework.

The **UE Location rApp** enhances network awareness by tracking User Equipment (UE) operating in the RRC_CONNECTED state. It communicates with the nRT-RIC to retrieve accurate, real-time location data for UEs within specified cells. This geolocation insight enables a range of applications, from resource optimization and mobility analytics to enhanced user experience strategies. The rApp also exposes a dedicated external API, allowing third-party applications or network functions to query current UE locations and their serving cell information. This supports advanced use cases such as targeted handovers, emergency response, and location-aware services within the RAN.

- Operational Workflow

- The CDF rApp continuously monitors PRB utilization to detect congestion in real time.
- If the congestion threshold is exceeded, affected UEs are identified and prioritized for load balancing.
- The CDF rApp queries the UE Location rApp to retrieve potential target cells with lower resource utilization.
- A congestion-aware handover decision is made based on PRB availability and UE proximity.
- The HO xApp requests UE measurement reports to confirm target cell suitability.
- If the target cell meets the criteria, a network-assisted handover is executed.
- The process is optimized to maintain XR service continuity with minimal latency.

3.2.1.2.1.2 Predictive CDF

To incorporate predictive capabilities to the CDF, an ML application has been designed and implemented which predicts whether there will be congestion in the next minute. The prediction is constant and can feed back new information, in background, to adapt to new congestion patterns that may arise. This process is performed for multiple cells at the same time. The initial ML model should therefore be chosen and trained for each cell.

The dataset, structured on a per-cell basis, was enriched through the application of network-level equations derived from the available telemetry variables. This process enabled the reconstruction of high-resolution temporal data with a granularity of one minute. A controlled congestion rate of approximately 50% per cell was established to balance class distribution. The resulting time series exhibited neither pronounced seasonality nor significant linear autocorrelation, thus supporting robust predictive modelling. The target variable for prediction is the uplink PRB load discharge percentage, which serves as the congestion indicator. A cell is in a congested state when this value reaches or exceeds the 80% threshold.

Historical data is used as features to capture complex relationships. It has been observed out that the models with the best performance are XGBoost and an LSTM neural network, trained with the initial 80% of the series and evaluated on the final 20%. They are tested with and without retraining. Recall (which measures how well a model finds all true congestion instances), F1_score (which shows how well a model detects true congestion while avoiding false congestion alerts) and computational weight are measured. In retrained models, the ability to learn new patterns and remember old ones is also assessed.

The LSTM model consists of three layers designed to learn temporal sequences with a window of a size according to the cell. It utilizes dropout to avoid overfitting and has an adjustment for possible imbalance between classes. For the XGBoost model, several hyperparameters are defined to increase accuracy and avoid overfitting

For feedback updates, the LSTM is refit with a reduced learning rate and a rolling window of 5 000 samples. The XGBoost model, by contrast, is retrained from scratch on a sliding history of 40 000 samples.

3.2.1.2.2 HO API implementation

The **Handover API (Amarisoft HO API)** as highlighted in [D4.1] and [D4.2], enables network-controlled mobility management to orchestrate seamless handovers. It proactively collects and processes UE measurement reports to evaluate radio conditions and neighbouring cell suitability prior to initiating handover procedures. The API incorporates congestion-awareness by triggering handovers when the serving cell experiences excessive load, thereby offloading traffic and maintaining service quality. Through predictive mobility strategies, it reduces service disruption and ensures continuity for delay-sensitive and high-priority applications in dynamic network environments.




Figure 22: Design of O-RAN CDF function to enforce HO upon congestion

3.2.1.3 Final evaluation results

3.2.1.3.1 CDF evaluation

3.2.1.3.1.1 Reactive CDF evaluation

The evaluation of the Congestion Detection Function (CDF) was conducted in a controlled 5G Standalone (SA) testbed, utilizing an Amarisoft gNB (Ultimate Box edition) connected to an Open5GS core network. The test scenario comprised two User Equipment (UEs) served by a single 5G cell: one functioning as the holographic user and the other as an interfering user. The interferer was configured to generate high traffic loads, thereby pushing the Physical Resource Block (PRB) utilization beyond a predefined congestion threshold.

Upon exceeding this threshold, the CDF emitted real-time congestion alerts to a centralized orchestration system, referred to as the holographic orchestrator. The orchestrator, upon receiving these alerts, triggered an intent-based mitigation response by instructing the holographic client to reduce its transmission bitrate. This adaptation helped mitigate congestion and stabilize the network's service quality.

Figure 23 the comparative impact of the orchestration decisions. When the orchestrator acted upon the CDF alerts, the system dynamically adjusted the holographic data rate, effectively preserving perceptual quality and preventing service degradation. In contrast, when the orchestrator disregarded the congestion alerts, sustained PRB overutilization led to a significant decline in the visual quality of the holographic stream, manifesting as reduced resolution and increased artifacts. The findings validate the role of the CDF as a critical enabler for proactive congestion management in holographic communication scenarios within 5G networks.

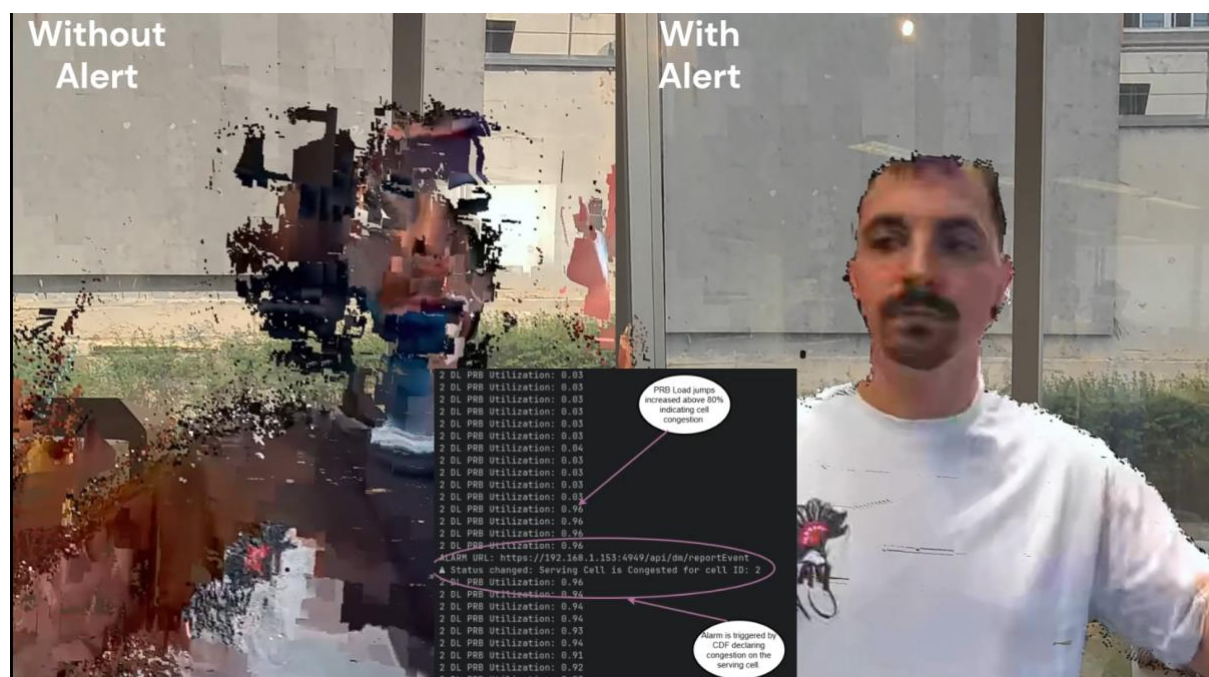


Figure 23: Snapshot of the holographic call with/without CDF alerts

3.2.1.3.1.2 Predictive CDF evaluation

The models for the predictive CDF that have worked best are the XGBoost with feature lags and, the LSTM with lag metrics as features.

Figure 24 depicts the results by looking at the Cell: CATX0133Q1A, whose behaviour is common in the group of cells studied.

	XGBoost	LSTM
Metrics for constant model (global)	0.9672 Recall 0.96223 F1_score	0.9426 Recall 0.9362 F1_score
Metrics for refitted model (global)	0.9708 Recall 0.9608 F1_score	0.96466 Recall 0.95317 F1_score
Computational weight	0.003 s predicting 0.73 s training	0.06 s predicting 13.3 s training
Code Complexity	Low	High
Learning new patterns in feedback	Guaranteed	Good at short times. In long times we don't know it.
Memory of old patterns	Forget where you don't train	In long times we have no security

Figure 24 XGBoost and LSTM Models Comparison Table

Figure 25 and Figure 26 show how the metrics differ between refitted (every 3000 new data) and constant versions of both models.




Figure 25 XGBoost Refitted model Recall vs Constant model Recall.




Figure 26 LSTM Refitted model Recall vs Constant model Recall.

Based on the metrics in the Figure 21, it is possible to conclude that the XGBoost model outperforms the LSTM model a little better, more concretely:

- Within each model, the refitted technique gives a slight improvement over the constant model. This does not happen for all cells, and the reason is that the data test does not always show new patterns from which retraining can learn.
- XGBoost quickly adapts to new patterns but can overwrite older ones, so we retrain it using an extensive historical window to preserve past knowledge.
- The LSTM shows good stability and improves in short retraining. However, due to its nature as a black box, it is not known how it balances the weights with the new information and how it forgets the old information.
- The computational weight (in terms of time) of the LSTM model is 18 times slower per retraining and 20 times slower in prediction than the XGBoost.

In view of these results, the most convenient model in this context and with these data is the refitted XGBoost model. Now, each cell can load this pre-trained model with its training data.

3.2.1.3.2 HO API evaluation

The performance evaluation was conducted in a controlled laboratory environment utilizing two Customer Premises Equipment (CPE) devices: one configured as an XR user and the other as an eMBB user generating uplink (UL) traffic. Both the serving and neighbouring cells were deployed using Amarisoft gNBs. Initially, the XR user operated exclusively on Cell 2, where the UL load remained below the congestion threshold of 80%. Upon activation, the eMBB user introduced substantial UL traffic, causing the UL load on Cell 2 to exceed the defined threshold. This congestion resulted in a noticeable degradation of the XR user's throughput, as illustrated in Figure 28. In response, the handover mechanism was triggered, and the XR user was migrated from Cell 2 to Cell 1. Following the handover, the XR user was able to achieve the target throughput, with UL load levels on both cells remaining within acceptable limits. Subsequently, the eMBB user ceased transmission, thereby alleviating congestion on Cell 2. Consequently, another handover was initiated to relocate the XR user back to Cell 2. These results demonstrate that increased UL load in the serving cell negatively impacts XR performance, whereas the handover to a less congested neighbouring cell effectively restores the required throughput by reallocating sufficient radio resources.




Figure 27: Plot of uplink PRB utilization vs time for both cells




Figure 28: Plot of UE uplink throughput vs time for both cells

3.2.1.4 Takeaways and conclusions

The integration of congestion-aware load balancing mechanisms within O-RAN-enabled networks, as proposed for the 6GXR Use Case 1, represents a pivotal advancement in supporting high-capacity services like holoportation. By leveraging near-real-time RIC functionalities—specifically the E2 KPM and RAN Control service models—alongside intelligent rApps for congestion estimation and load balancing, the solution dynamically adapts to network conditions. This enables proactive mitigation of congestion, efficient user distribution across cells, and sustained service quality for XR applications. Ultimately, the proposed architecture fosters improved spectral efficiency, enhanced user experience, and scalable support for next-generation immersive services within heterogeneous network environments.

As a next step, the framework will be extended to include rate recommendation mechanisms triggered upon congestion detection. This enhancement aims to guide holographic clients in adjusting transmission rates to maintain service continuity under constrained conditions. Validation and performance testing of this complete solution are planned to use the Amarisoft 5G testbed, enabling real-world assessment of its effectiveness in dynamic RAN scenarios.

3.2.2 Energy-aware end-to-end resource management

3.2.2.1 Motivation

In the O-RAN architecture, the control loops (fast control) operating at ~ 1 s timescale at the near-RT RIC can be utilised to perform dynamic RAN resource management. Extending the experimentation related to RAN energy efficiency optimization methods executed at the network and gNB/cell level in 3GPP networks at $\gg 1$ s timescale (slow control), this enabler studies the KPI monitoring and dynamic UE specific radio resource management functionality in the O-RAN architecture and assesses the potential benefits from the point of view of power consumption optimisation. Based on the change in the cell load and related E2 Node (comprised of a single CU, DU, and RU) power consumption when two different types of XR applications are used (App1 and App2), a data rate limiting xApp running in the near-RT RIC adjusts the UE specific resources to better serve the needs of the high priority XR user/application in an energy efficient manner. To be in line with the experiments performed in the 3GPP network infrastructure (see subsection 2.2.4, Table 3), the XR applications emulated in the O-RAN tests will be Collaborative VR representing a high bitrate App1 and Holographic Call over IMS Data Channel representing a low bitrate App2.

3.2.2.2 Solution design

Figure 29 presents the overall design to introduce the desired functionality and potential energy saving benefits into an O-RAN architecture. The O-RAN architecture's E2 Node components are drawn with blue colour, whereas the near-RT RIC platform is drawn with purple. The xApps providing the required monitoring and control functionalities for the enabler are drawn with yellow colour and the utilised external databases with orange. The UE is drawn with dark magenta colour. Figure 29 also shows the required information exchange between the different components. The green arrows indicate the flow of energy monitoring data, blue arrows the flow of KPI monitoring data, orange arrows the flow of decision algorithm input data between the xApps, and red arrows the control signalling data towards the E2 Node or UE.




Figure 29: O-RAN test setup and information flows.

The required monitoring and control functionality for the tested enabler can be implemented on top of the O-RAN platform as xApps. The xApps are divided into logical components providing a specific

capability into the overall monitoring and control functionality, which from the implementation perspective can also be deployed as a single xApp if preferred. Table 11 in Annex 8.1 describes the functionality of the xApps in more detail. Table 12 in Annex 8.1 lists the messages exchanged between the xApps, O-RAN architecture components, and energy measurement framework shown in Figure 29.

For the collection of O-RAN Key Performance Measurements (KPM), E2 interface services are accessed through the Accelleran dRAX APIs provided by the utilised O-RAN platform. The control signalling for the rate limiting functionality can be based on the modification of UE-specific Session Aggregate Maximum Bit Rate (Session-AMBR) parameter through 3GPP messaging procedures which, instead of 5GC, are initiated from by the near-RT RIC for more dynamicity. Alternatively, the rate limiting functionality can be based on the modification of the slice-specific Physical Resource Block (PRB) quota if the UEs running XR applications are using a dedicated slice and other UEs are using another slice.

The test setup used to test the potential energy saving benefits of the enabler utilises WP5 enablers in the form of the energy measurement framework including the Carlo Gavazzi EM111 energy analyser devices for monitoring the energy consumption in the O-RAN architecture as well as the related data collection and distribution pipelines based on InfluxDB and MQTT brokers. More information on the energy measurement framework can be found from the 6G-XR deliverable [\[D5.1\]](#).

3.2.2.3 Final evaluation results

This subsection presents the measured baseline values for the O-RAN platform's throughput performance, power consumption, and their interdependency at different levels of network load. All tests are performed in an O-RAN deployment consisting of single near-RT RIC, CU, DU, and RU components as well as of a dedicated Open5GS 5GC instance. The hardware infrastructure consists of two Supermicro AS-1015A-MT servers, one hosting the CU, near-RT RIC, and 5GC software as containerised services and one running the DU protocol stack, as well as of a Benetel RAN650 RU. These values can be used as input for a simple threshold-based control logic in the data rate limiting xApp (rate_lim) that is responsible for making decisions on the increasing or decreasing the Session-AMBR value for individual active UEs or PRB quota for individual network slices. However, the implemented monitoring functionality at the O-RAN platform provides a variety of KPIs in addition to the ones utilized in this experiment, so it can also facilitate the integration of more complex algorithms, e.g., by 3rd party developers.

Figure 30 presents the achieved average DL throughput values for different levels of offered network load. All throughput values on the presented curves are the highest measured throughput after averaging over the whole duration of a five-minute test run. During the experiments, the maximum throughput value is defined by gradually increasing the number of emulated XR users in the contributing to the overall network load. All test traffic is generated with iPerf3 [\[iPerf2025\]](#) server located at the network edge, emulating the XR user with applications from 6G-XR use cases UC3 and UC4 with parameters shown in Table 3. All test traffic is using UDP as the transport layer protocol. As UDP does not include any congestion control mechanisms, it enables us to overload the E2 Node by using only DL traffic. As can be seen from the throughput curves presented in Figure 30, the larger data packet size of UC4 results in a higher maximum throughput, which saturates around 770 Mbps when the offered network load starts to approach 1 Gbps. For UC3, the smaller data packets size and, consequently, higher number of individual data packets to process results into a maximum throughput around 610 Mbps, which is achieved already before the offered network load reaches 800 Mbps. This indicates that the E2 Node is overloaded much faster by the larger amount of smaller packets, as expected.




Figure 30: Measured maximum DL throughputs for UC3 and UC4.

Figure 31 presents the measured average power consumption for the whole O-RAN infrastructure during the maximum throughput test described above. The higher amount of processing needed to handle the larger amount of smaller packets generated in UC3 can also be seen in the power consumption. The difference comes from the CU, RIC, and 5GC services, all of which are focusing on the processing related to the higher protocol layers and handling of the user plane data packets. Overall, the measured power consumption with UC3 data traffic is approximately 7-14 W higher than with UC4 data traffic at the same level of offered network load.




Figure 31. Measured power consumption of the O-RAN infrastructure for UC3 and UC4.

Taking into consideration the nearly linear relationship between the achieved throughput and offered network load under 500 Mbps for both UC3 and UC4 in Figure 32, a simple verification test for energy

consumption optimisation with the xApp-based dynamic control approach described in subsection 3.2.2.2 was performed. In the test setup, two UEs are receiving a varying amount of data from the iPerf3 edge server. UE1 is a high priority XR user receiving emulated traffic streams from UC4. UE2 is a low priority user receiving a continuous UDP data stream at 300 Mbps. In the beginning of the 26-minute test run depicted in Figure 32, UE1 starts to receive one UC4 data stream at 88 Mbps (dark blue dotted line) and UE2 starts to receive its full 300 Mbps UDP data stream (orange dotted line). The number of UC4 data streams received by UE1 are increasing by one every 5 minutes. When the third UC4 data stream is added for UE1 around 11 minutes into the test, the combined data rate of UE1 and UE2 (green solid line) crosses the example threshold of 500 Mbps and the maximum data rate of UE2 is limited to 200 Mbps to compensate. Again, as the fourth UC4 data stream is added for UE1 around 16 minutes into the test, the 500 Mbps threshold for the combined data rate is crossed and the maximum data rate of UE2 is further limited to 100 Mbps. Finally, UE1 stops receiving all four of its UC4 data streams around 21 minutes into the test and the data rate limitations of UE2 can be lifted so that it again receives the full 300 Mbps as in the beginning of the test.




Figure 32. Example of the changes in the overall network load and O-RAN infrastructure power consumption with high priority XR UE (UE1) and low priority UE (UE2).

The solid light blue line in Figure 32 shows the measured power consumption of the O-RAN infrastructure during the described test run. The power consumption is measured 1/s, but to smooth out the measurement noise, the curve shows the average consumption over each 5-minute step where the amount of UC4 data streams received by UE1 is changed. The Y-axis on the right-hand side in Figure 32 shows that by limiting the total network load to 500 Mbps the power consumption also saturates to 380 W level. Taking into consideration the power consumption curves in Figure 32, the potential savings achieved with this simple approach varies from few Watts up to ten Watts, which in terms of relative power consumption is few percents, but adds up when the network deployment consists of more than one cell side and there are also higher level power saving methods (discussed in subsection 2.2.4) in use.

3.2.2.4 Takeaways and conclusions

The performed measurements confirmed the two assumptions used as main motivation for the design of the dynamic energy consumption optimization functionality for O-RAN architectures. First, the measurements showed that the network load and power consumption are directly linked also in a disaggregated RAN deployment and the power consumption is not only dependent on the overall load, but also on the data traffic parameters causing the load. Second, the verification test with two UEs proved that the maximum power consumption of the overall O-RAN infrastructure can be controlled by limiting the allowed data rates of lower priority users when saving energy is made the most important objective at the system level. These early results encourage us to seek further dynamic energy consumption optimisations also with more complex algorithmic approaches in the future and show that there is also interesting opportunities for 3rd parties to experiment with these topics on top of suitable O-RAN platforms.

3.2.3 O-RAN enabled slicing to support XR services

3.2.3.1 Motivation

This enabler was described in [\[D4.2\]](#) with Split-8 for FR1. Plan during the creation of [\[D4.2\]](#) was to update the system with FR2 RU from LiteOn. That has now been done as the LiteOn FR RU has been integrated into the Allbesmart OAIBOX based O-RAN system. Otherwise, the system has been kept as it was for the FR1 implementation where Ettus USRP X410 was used as the Radio Unit (RU). System details are described in the document [\[D4.2\]](#) and they will not be repeated in this document but only the changes and upgrades to the system are explained and defined.

3.2.3.2 Final Solution design

As described in the previous chapter, the radio unit of the system defined in the [\[D4.2\]](#), was upgraded to support functional RAN Split-7.2a. Previously system supported Split-8 where the majority of the radio related tasks, even the lower physical layer (PHY), are performed in the DU. In Split-7.2a lower functions of the physical layer are executed in the RU, as found in most practical RU implementations [\[O-RAN\]](#).

In the OAIBOX setup, gNB connects to RIC via E2 interface. XApps is hosted on this Near-Real Time RIC. RAN slicing is being managed by xApps by defining SLAs and dynamically controlling MAC scheduler. Exclusive PRBs are assigned at the MAC layer for isolation. When maximum throughput is requested by multiple UEs, PRBs are divided dynamically among both slices. So xApps adapt the slice configuration in real time to meet specific UE requirements.




Figure 33: O-RAN-enabled slicing to support XR systems

Other components of the system are the same as in the FR1 composition. Slices are to be the same, except that the frequency will naturally be changed to be in band n257 instead of the n77. Test cases remain the same to guarantee the comparability of the FR2 test results with the ones from the previous FR1 tests.



Figure 31: O-RAN FR2 (mmWave) split 7.2a setup

In Figure 31, on the left side, there is a Liteon radio connected to a robot arm. On the right side, there are CU and DU components which are hosted on Allbesmart OAIBOX. The system operates on n257 (FR2 mmWave) band. The user equipment used for the measurements is Quectel RM530 module.

3.2.3.3 Final evaluation results

Figure 32 below shows the throughput values for both URLLC and eMBB slice in scenario 1. In this measurement scenario both UEs are connected to the network slices operating under normal traffic load conditions.




Figure 32: Throughput measurements for scenario 1 (normal load)




Figure 33: Delay measurements for scenario 1 (normal load)

Figure 34 below provides the visual comparison between data throughput and latencies for both URLLC and eMBB slices in scenario 1. URLLC is shown on the left side of the picture and eMBB on the right side. Downlink values are shown in green colour and uplink values are shown in blue colour.




Figure 34: Throughput and latencies for both slices in scenario 1 (normal load)

Figure 35 below shows the throughput values for both URLLC and eMBB slice in scenario 2. In this scenario eMBB slice was overloaded by generating the background traffic. One UE is connected to URLLC slice and other to overloaded eMBB slice.




Figure 35: Throughput measurements for scenario 2 (eMBB slice overloaded)




Figure 36: Delay measurements for scenario 2 (eMBB slice overloaded)

In above Figure 36, latency measurements are shown for scenario 2. Latency measurements are performed in both directions and visual comparison can also be seen in Figure 37 below. Downlink values are shown in green colour and uplink values are shown in blue colour. In the Figure 37, the right side represents the URLLC slice, while the left side corresponds to the eMBB slice.



Figure 37: Throughput and latencies for both slices in scenario 2 (eMBB slice overloaded)

3.2.3.4 Takeaways and conclusions

In the presented measurement results it is clearly visible that similar behaviour that was in FR1 measurements presented in [\[D4.2\]](#) can be observed. When eMBB slice is overloaded with data in downlink direction, the eMBB downlink latency increases as in the FR1 setup. Measurement results also have many latency spikes visible when the eMBB slice is overloaded with data as was done earlier also in the FR1 tests. These spikes impact the user VR glass Quality of Experience just as they did with FR1. With FR2, as in FR1, when inspecting the URLLC slice, it can be seen that the overloaded eMBB

slice is not impacting the URLLC resources. The downlink latency for URLLC is in the same range as in the test case 1 without extra latency spikes that are experienced in eMBB slice. So, the slice isolation in the FR2 test platform with split 7.2a is working as expected and in similar way as it was working when using the FR1 RU with split 8.

In this enabler O-RAN implementation presented in D4.2 was upgraded from FR1 RU at 3.8 GHz to support also FR2 RU at 25 GHz. Then the same tests executed for the FR1 case were repeated for the FR2 setup. As the LiteOn FR2 RU did not support split 8 but only split 7.2a so the resource share between RU and DU was changed to support split 7.2a. It can be concluded that implementation in both cases, FR1 with split 8 + USRP X410 RU and FR2 with split 7.2a + LiteOn RU, were working as expected and slice isolation is working as it should be in both test cases.

4 6G DISRUPTIVE XR ENABLERS

4.1 SUMMARY OF DISRUPTIVE ENABLERS PROPOSED IN D4.2

[D4.2] described the motivation and initial solution design and evaluation results for new disruptive technology enablers in the areas of Reconfigurable Intelligent Surfaces (RIS), THz and Sub-THz frequency bands, and Integrated Sensing and Communications (ISAC). Analysis was made in relation to their applicability for 6GXR use cases. In addition, some suggestions for integration of the baseband THz-RIS and ISAC components with the high-frequency transceivers in 140 GHz and 300 GHz were made.

This chapter completes the description by providing details of the final solution design and final evaluation results of the proposed disruptive enablers. To facilitate the integration, an offline mechanism based on files between the baseband and RF transceivers is described, and some useful results of the integration process are given. An additional enabler for channel estimation and equalization in doubly-dispersive channels is also proposed, and the corresponding evaluation results are provided in the form of numerical simulations.

4.2 6G DISRUPTIVE XR PROPOSED ENABLERS

4.2.1 High-Frequency Transceivers for THz-RIS and ISAC

4.2.1.1 Motivation

To stimulate the adoption and further development of XR applications, XR technologies should support wireless connectivity with extremely high data rates and ISAC capabilities for sensing. Current wireless connectivity solutions cannot satisfy the throughput and accurate sensing requirements of XR use cases. High throughput and accurate sensing can best be obtained using wide band spectrum, as currently available in the mm-wave and sub-THz spectrum. The so-called sub-THz regime involves frequencies between 100 GHz and 300 GHz and is considered the most promising range for ultra-high frequencies, where some regulation already exists for communications and active/passive sensing services. Above it, the THz frequency range further extends from 300 GHz up to 10 THz but is less explored due to limited availability of efficient transceiver components. The D-band (110-170 GHz), the G-band (140-220 GHz) and the H-band (220-330 GHz) are considered as the frequency bands where sufficient spectrum can be allocated for wireless communication with data rates of 100 Gbit/s and beyond. Utilizing these frequency bands requires disruptive research to overcome the challenges inherent to operating at such high frequencies.

4.2.1.2 140GHz Transceiver design

4.2.1.2.1 Final Solution design

In [D4.2], the 140 GHz transmitter and its evaluation system have been described. The reference receiver from commercial waveguide components is now replaced with the 140 GHz integrated receiver module. This full communication link set-up is depicted schematically in Figure 34.

At the core of this set-up are the 140 GHz transmitter module (Teradio TX in the figure) and the 140 GHz receiver module (Teradio RX in the figure).

They are both connected to a signal generator generating a 15 GHz local oscillator frequency signal. Both in the transmitter as in the receiver, this local oscillator input is needed to generate a carrier

frequency for the mixers in the receiver and transmitter front-ends a factor 9 higher in frequency, at 135 GHz.

At the side of the transmitter, an arbitrary waveform generator (AWG) is used to input a baseband signal to the transmitter front-ends. Commercial attenuators and bias-tees are used to give the baseband signal the swing and offset expected by the chip.

At the side of the receiver, the received baseband signal, with DC offset removed, is captured by a digitizing oscilloscope.

A lab power supply is only used for the offset voltage of the bias-tees. The teradio modules get their 4 power domains from a dedicated PCB with a 5V usb-C power input, that is mounted on the backside of the RF modules.

We the correct front-end settings (both for transmitter and receiver modules) through a Raspberry Pi with a dedicated interface board mounted on its GPIO pin header. Apart from calibration settings, the most noteworthy settings of the front-ends are the phase shifter settings: by adjusting the relative phase of the 4 front-ends, the resulting beam pattern of the module can be adjusted.

The full set-up can be controlled from a host PC over a network.




Figure 34: 140 GHz communication link set-up

4.2.1.2.2 Final evaluation results

We evaluated the performance of the D-band link by placing the set-up described in the previous paragraph in a chamber with walls absorbing electromagnetic waves, as shown in Figure 35. It is not a full anechoic chamber, but it is sufficient for a stable line-of-sight link. In the figure below, the unit on the right (in purple casing) is the receiver unit, the unit on the left (in white casing) is the transmitter unit. In the depicted set-up, the distance between both units is about 1.5m and the receiver is placed on a programmable turntable. By rotating the turntable, we can measure the link for given reception angle to obtain a beam-pattern measurement.

A set-up the modules positions switched, with the transmitter on the turntable and the receiver at a fixed position was used for the transmission angle measurements.




Figure 35: D-band link in absorber cage

Figure 36 and Figure 37 show the beam pattern for the transmitter and the receiver respectively. In both cases, the beam of the measured module (TX or RX) was steered to azimuth angles from -60° to $+60^\circ$ in steps of 10° . For each case, the beam pattern was measured in step of 5° from -70° to $+70^\circ$. For the readability of the figure, lines are only plotted for a limited range around the peak observed signal strength. From the figure we can conclude that both TX and RX modules allow beamsteering without significant signal level reduction in azimuth angles from -40° to $+40^\circ$.




Figure 36: TX Beam pattern for various beamsteering configurations




Figure 37: RX beam pattern for various beamsteering configurations

Next to the beam pattern measurements, we also performed EVM measurements of the D-Band link using an 802.11ad-like single carrier, frequency domain equalised waveform. The highest link capacity was observed using a 4 GHz Bandwidth QPSK modulated waveform: 5.7 Gbps. Received baseband spectrum and constellation diagram are shown in Figure 38.




Figure 38: Baseband spectrum and constellation diagram of highest throughput mode

The spectrally most efficient mode is using (up to) 1 GHz of bandwidth and a 16 QAM modulation, yielding a throughput of 2.6 Gbps, or a spectral efficiency of 2.6 bits/Hz. Received baseband spectrum and constellation diagram of this mode is show in Figure 39.




Figure 39: Baseband spectrum and constellation diagram of most spectrally efficient mode

4.2.1.2.3 Takeaways and conclusions


We evaluated a D-band link consisting of a transmitter and receiver module, both with 4 front-ends and antennas to support analog beam-steering in the azimuth plane. We have demonstrated a link capable of reliably delivering a physical layer throughput of up to 5.2 Gbps using a bandwidth of 4 GHz and up to 2.6 Gbps using a bandwidth of 1 GHz. Using the beamsteering functionality, this link can be adjusted in the azimuth plane for angles between -40° and $+40^\circ$, both for the transmitter and the receiver.

4.2.1.3 300GHz Transceiver design

4.2.1.3.1 Final Solution design

The 6G-XR 300 GHz Transceiver design is an ongoing activity during the project duration. The 300 GHz transceiver includes both receiver and transmitter functionalities and the receiver part has been now implemented and verified with the measurements. The final 300 GHz array chip during the 6G-XR project was taped out to IHP SiGe process in December 2024. The final RF chips will be available in June 2025 and those will be tested until the end of 2025 or to the end of the project. The final 300 GHz RFIC will be verified with over-the-air (OTA) measurements since the antennas of the antenna array have been integrated into the RFIC. The implemented RFIC will be mounted on specially designed printed circuit board (PCB) enabling OTA testing. The final RFIC's radio performance cannot be

The final 300 GHz chip is an array receiver which is based on the previously implemented RF architecture of the 300 GHz sliding intermediate frequency (IF) receiver chip and a photograph of the RF chip are shown in Figure 40.



[SNM+2023]

The local oscillator (LO) signal uses doubling frequency stages where a low-frequency reference signal, in this case 50 GHz, is raised to the required mixing frequency inside the chip. The LO signal at two-thirds and one-third of the carrier frequency is generated by the multiplication of the external LO signal of 50 GHz. First, the external LO input signal frequency is multiplied by 2 by a frequency doubler that generates the LO signal at 100 GHz. This signal is amplified and then divided into two paths. Path-I of the LO chain generates the LO signal for the RF mixer using one additional frequency doubler and LO signal amplifiers. Path-II generates quadrature LO signal for I/Q mixer stages by using signal amplifiers and a quadrature hybrid coupler (QHC) for 90-degree phases shift. The final RFIC's LO generation signal path has been modified so that it can be driven with 12.5 GHz reference signal instead of 50 GHz compared to be previous generation 300 GHz RFIC.

Since the final 300 GHz RFIC is not available at the time of writing, we discuss based on measurement results of the previous generation 300 GHz RFIC. The operation of the implemented 300 GHz receiver has been verified with single-tone sinewave (CW) measurements and with modulated signal measurements. First, the CW and S-parameter measurements of the receiver have been measured. The measurement setup and the gain curve of the receiver are shown in Figure 41. The measured 3-dB and 6-dB RF bandwidths of the receiver are 26 GHz and 36 GHz, respectively when keeping fixed baseband frequency at 2 GHz.




Figure 41: Sinewave signal measurement setup and 300 GHz receiver's gain curve

Modulated signal measurements have been done with the following measurement equipment setup with the developed 300 GHz receiver in below Figure 42.




Figure 42: Modulated signal measurement setup for 300 GHz receiver testing [Sin2024]

A modulated signal is generated from the Keysight M8195A arbitrary waveform generator (AWG). Using VDI's WR3.4 modulated frequency extender (mFEx) SAX module, the modulated signal is upconverted to sub-THz/THz frequencies. A VDI's power amplifier is used to boost the low output power of the mFEX. PA is followed by a mechanical waveguide attenuator to control the overall RF power going to the DUT. Two separate Keysight's E8267D PSG are used to provide the LO signal to both mFEX and DUT. The input LO signal at 50 GHz is provided with the help of PSG and an external Marki Markowave MMD-2060LU frequency doubler. The downconverted low-IF I/Q modulated signal is demodulated and analyzed with the help of Keysight's UXR 0404A oscilloscope and PathWave Vector Signal Analysis (VSA) 89600 software, respectively.

One major challenge of the previous conducted RF performance measurement was to perform input signal power sweeps with both CW and modulated signals with known power level in the input of the RFIC. The power calibration was improved by implementing measurement setup as shown in Figure 43.




Figure 43: Modulated signal measurement setup for 300 GHz receiver testing [SPR+2025]

The additional power amplifier PA2 enabled improved signal level for CW testing that the signal level can be measured accurately with RF power meter. The RF power meter provides absolute power level, but the dynamic range of those is limited and thus additional power amplification is needed for the accurate power calibration.

The final evaluations of the 300 GHz RFIC will be performed with OTA measurements and for this purpose we have developed a robotic arm-based OTA measurement system to be used in RF anechoic chamber. The block diagram of the developed robotic arm OTA system is shown in Figure 44.




Figure 44: Block diagram and photograph of 300 GHz OTA measurement setup at University of Oulu [HRM+2025]

4.2.1.3.3 Takeaways and conclusions

The 300 GHz RFIC development has been done during the project duration and the final chip will be available in June 2025. The final validations will be performed OTA manner at the second half of 2025.

The developments of measurement systems for 300 GHz radio verifications for conductive and OTA have been essential for the success of the project. Similarly, the CW and the modulated measurements require different measurement systems, so during the project we have developed multiple new measurement capabilities up to 330 GHz in UOulu.

Measured RF performances from the 300 GHz RFICs have been mostly aligning with simulations results but significant debugging of the chip and continuous development of measurement systems has been required to achieve these results.

We will continue the measurement system development for 300 GHz systems that when those are needed for 6G extreme data rate system validations then we will be ready for the purpose.

4.2.2 Baseband implementation for THz-RIS and ISAC based on SC-FDE

4.2.2.1 Motivation

XR Metaverse use cases like remote control, remote maintenance, enterprise metaverse, real-time holographic communications, and collaborative 3D digital twin-like environment demand ultra-high bandwidths and sensing capabilities. Most of the current approaches at very high frequencies to cope with these use cases generally exhibit very limited communication capabilities and lack ISAC capabilities. 6G-XR has bridged that gap by developing an ISAC-capable baseband design for Sub-THz communications integrated with the High-Frequency Transceivers for THz-RIS and ISAC.

In addition to that, further research is also needed to cope with scenarios characterized by highly doubly-dispersive channels, like those encountered at high mobilities or in the presence of significant phase noise. Both situations lead to delay and Doppler channel impairments that need to be compensated at the receive side prior to performing sensing detection or data demodulation. Most of the state-of-the-art approaches rely on novel waveforms (like OTFS, FM-OFDM, etc.) with better Doppler support, and are therefore less applicable to existing waveforms like CP-OFDM or SC-FDE. Further work is thus needed to accomplish this goal while keeping the characteristics of the waveform unchanged.

4.2.2.2 Solution design

The solution design of this enabler comprises two main components described in what follows.

4.2.2.2.1 Baseband Integration with the High-Frequency Transceivers

The ISAC-capable baseband for Sub-THz communications was designed with the goal of being able to integrate it with the High-Frequency Transceivers for THz-RIS and ISAC. Such integration allows for performing experimental measurements in a controlled lab environment that can be more valuable than any simulated results.

An offline integration mechanism was used to ease the integration process while overcoming some of the hardware constraints at these frequencies. Sub-THz transmission requires very large bandwidths and significantly higher processing capabilities at the baseband hardware. In order to circumvent the constraints that most of the current available programmable baseband platforms pose, an offline integration approach was leveraged based on text files containing IQ samples of a number of transmit and receive waveform symbols. These files fed, or were produced from, laboratory equipment capable of parsing the contents and generating a baseband waveform that is then upconverted to Sub-THz frequencies by the high-frequency transceivers in subsection 4.2.1. The contents include one I,Q sample per line in floating- or fixed-point format, with as many files as symbols are transmitted.

A simple transmission was first tested to check the main RF characteristics of the waveform signal, with the following configuration:

- SCS = 960 kHz
- BW = 1.152 GHz (100 RB = 1200 complex constellation values per SC-FDE symbol)
- With CP
- Sampling frequency = 1.966 Gsps
- FFT size = 2048
- 16QAM modulation
- RRC roll-off factor = 0.35

Further integration involved multiplexing DMRS-like and PTRS-like signals with data, so that channel estimation and phase noise compensation could be performed on detection, respectively.

At the receive side, suitable IQ samples were captured by down-converting and quantizing the analog signal received by the 140 GHz and 300 GHz High-Frequency Transceivers into the baseband domain. Any synchronization inaccuracies in the received signal were overcome by leveraging the cyclic prefix duration, and the useful part of the symbol was taken after leaving a guard time equal to one half of the CP duration. Further channel estimation, phase noise compensation, and equalization must be performed to yield the decoded data.

4.2.2.2 Channel estimation and equalization in doubly-dispersive channels

This technique aims to study the problem of ISAC detection in channels with high delay and Doppler spread, assuming a waveform with Frequency Domain Equalization (FDE) capabilities, e.g., SC-FDE, CP/ZP-OFDM, DFT-s-OFDM, FM-OFDM, CE-OFDM, etc. Although this is not implemented in the baseband testbed, this is a relevant problem to study that can bring multiple insights for ISAC in scenarios with high mobility and/or high phase noise impairments, as often encountered in ISAC.

A reference signal structure like in Figure 45 is proposed to be added on the time-frequency grid of the waveform. It comprises two components in yellow denoted as WB-RS and NB-RS, respectively, that can coexist with other data and control information in the time-frequency grid defined for the waveform. The presence of WB-RS and NB-RS enables separate estimation of the long-term and short-term channel components, respectively. The long-term channel components are the multipath delays, whose changes are several orders of magnitudes slower than the short-term channel components represented by the multipath complex amplitudes. Both long-term and short-term channel can be collectively referred to as incremental CSI. The goal is to estimate incremental CSI in channels characterized by significant delay spread and Doppler spread.




Figure 45: Reference Signal structure based on incremental CSI estimation

The solution, as detailed in the Annex 8.2, reduces to solving the optimization problem

$$\tilde{b}_{i,l} = \arg \min \left\{ \|\bar{\mathbf{R}} - \mathbf{S}\tilde{\mathbf{B}}\|_F^2 + \lambda^2 \|\mathbf{B} - \mathbf{B}_0\|_F^2 \right\},$$

where $\tilde{b}_{i,l}$ are the complex multipath amplitudes that characterize the time-domain channel impulse response. This vector, together with the multipath delays estimated via WB-RS, completes the channel characterization for any doubly-dispersive channel with arbitrary variation statistics.

4.2.2.3 Final evaluation results

4.2.2.3.1 Results on the Baseband Integration with the High-Frequency Transceivers

In this section, the results of a first integration of the TX path with the 140 GHz and 300 GHz High-Frequency Transceivers are shown.

Figure 46 provides the RF output spectrum of the 300 GHz measurement system measured by UOulu after down conversion to baseband. The operational bandwidth is ~800 MHz, aligned with the theoretical one ($1.152/2 \times 1.35 = 777.6$ GHz) and smaller than the nominal 1.152 GHz due to the smoothing effect of the RRC filter. The EVM plot is not showing a correct constellation because the lab equipment did not incorporate the appropriate SC-FDE decoding algorithms for this waveform.

Similarly, Figure 47 provides the RF spectrum of the 140 GHz measurement system obtained by IMEC. The same operational bandwidth close to 800 MHz is visible in the spectrum.

Plots related to the RX integration will be provided later.




Figure 46: RF output of the test signal at 300 GHz measured by UOulu after down conversion frequency extender.




Figure 47: RF spectrum of the test signal at 140 GHz as measured by IMEC.

4.2.2.3.2 Results on the Channel estimation and equalization in doubly-dispersive channels

In this section, the bit error rate (BER) performance with piecewise channel estimation and equalization was assessed for ZP-OFDM and FM-OFDM in two different high-speed channel models. These waveforms were considered given their single-carrier nature (similar to SC-FDE) and the fact that they can naturally incorporate a ZP field, as needed for highly time-varying channels. Results can easily be extended to SC-FDE if the CP field is replaced by ZP. The results depicted are a summary of the provided results in the paper submitted to IEEE Transactions on Communications.

A WB-RS transmission with 4 symbols spanning full bandwidth (11.1 MHz, for a 1024-point FFT size and a subcarrier spacing of 15 kHz) was set to estimate the multipath delays of the channel, with a periodicity of 99 slots. A NB-RS transmission based on $M = 32$ or 64 was set in every symbol. Piecewise MMSE equalization was considered with $N_L = 4$. Results for a ETSI Highway NLOS channel model at 300 km/h, a carrier frequency of 4 GHz, 384 active subcarriers, $M = 32$, and QPSK and 16QAM are plotted in Figure 48. The NB-RS overhead is 25%, while the WB-RS overhead is negligible. It is apparent that FM-OFDM and ZP-OFDM behave much better than CP-OFDM with both ideal or LS channel estimation, with a degradation of 3-4 dB at $\text{BER} = 10^{-1}$ with respect to ideal piecewise estimation.




Figure 48: bit error rate performance of piecewise estimation and equalization in a Highway NLOS channel model at 300 km/h and $M = 32$.

Similar results are obtained in Figure 49 with a 3GPP TDL-C (300) channel model at 300 km/h, with $M = 64$ and 256 active subcarriers in this case to account for the more complex multipath profile (24 taps). In this case the degradation is 5-6 dB at $\text{BER} = 10^{-1}$ despite the higher NB-RS overhead (50%).




Figure 49: bit error rate performance of piecewise estimation and equalization in a 3GPP TDL-C (300) channel model at 300 km/h and $M = 64$.

The NMSE for both channel models is plotted in Figure 50, where the value of the squared bound of the piecewise approximation error ε_{LTI}^2 (explained in the Annex 8.2) is also shown. It is apparent that most of the NMSE is dominated by imperfections in the channel estimation process, rather than the piecewise approximation error, which justifies the applicability of the piecewise approach.




Figure 50: NMSE and squared bound for Highway NLOS and TDL-C (300) channel models at 120 km/h and 300 km/h.

Both the overheads and the degradation incurred in the proposed realistic channel estimation and equalization techniques are in line with other state-of-the-art, like 5G NR, where channel estimation typically incurs overheads between 33% and 50% at the OFDM symbols carrying DMRS, and even 100% overhead (full symbol devoted to DMRS) when using DFT-s-OFDM in the UL. Channels at very high speeds generally require multiple DMRS symbols per slot. In the extreme case of doubly-dispersive

channels with significant variations within the symbol, a DMRS-like signal needs to be mapped in every symbol to properly track the channel variations.

4.2.2.4 Takeaways and conclusions

The proposed baseband implementation for THz-RIS and ISAC with SC-FDE aims to enable very high throughputs (as foreseen in THz and Sub-THz technologies) and high sensing accuracy (as allowed by the very high bandwidths and small carrier wavelengths in these bands) for XR applications. Given the challenging amount of baseband processing required with real-time equipment, the processing functions were developed in the form of Matlab blocks for easier post-processing of the received samples after passing through the High-Frequency transceivers and a real channel in the lab, without suffering any constraints on the bandwidths. The results obtained demonstrate the possibility to build a full disruptive system by splitting it into baseband and frontend components that are offline integrated in the lab.

Future lines of work could aim to implement the blocks in specialized hardware (e.g., based on USRPs or FPGAs) to enable real-time experiments, and to develop ISAC algorithms leveraging the fine resolution of the testbed.

ISAC applications often require tracking high-speed targets in challenging environments with strong multipath, Doppler and phase noise. Channel estimation and equalization in these conditions sometimes rely on assumptions a-priori made in the channel, e.g., to mimic its variations via BEM models. We have proposed and evaluated a technique to estimate and equalize any type of doubly-dispersive channel via a Reference Signal inserted, or superimposed, to subcarriers of the signal waveform, applicable to waveforms with FDE capabilities under any type of channel with no a-priori assumptions on their variations. Further works in this field could aim to apply it to different waveform types and compare their resilience in the presence of channel impairments.

4.2.3 Deep Reinforcement Learning for THz-RIS

4.2.3.1 Motivation

The 6G disruptive XR proposed enablers, as described in reports [\[D4.1\]](#) and [\[D4.2\]](#), highlight Reconfigurable Intelligent surfaces (RIS) and Deep reinforcement Learning (DRL) as essential solutions to overcome the inherent challenges faced by wireless networks at mmWave and THz frequencies. RIS technology, combined with DRL algorithms like Deep Deterministic Policy Gradient (DDPG), dynamically configures phase shifts in RIS elements, significantly improving spectral efficiency and adaptability [\[Huang2020\]](#). Report D4.3 extends this foundational work by introducing substantial advancements, notably refining RIS channel modelling for more accurate simulations at THz frequencies and fine-tuning machine learning algorithms such as DRL algorithms to enhance learning efficiency and convergence, thus optimizing network performance further [\[Sun2019\]](#).

4.2.3.2 Final Solution design

The final solution design is basically divided into two parts:

Part 1: The THz-RIS performance evaluation: Initially, comprehensive evaluations of THz-RIS performance were conducted at **28 GHz**. These evaluations systematically assessed RIS configurations with varying numbers of RIS elements (N). Specifically, we examined configurations across two main scales: low N (ranging from 4 to 64 RIS elements) and high N (ranging from 64 up to 250 RIS elements).

These studies clearly demonstrated significant spectral efficiency (SE) gains as the number of RIS elements increased. This enhancement is attributed primarily to increased beamforming gains and improved RIS optimization techniques [Vulchi2024]. Building upon these results, we extend our simulations to 'Hyper N' where extremely large values of N (the number of RIS elements) are considered. The goal here is to push the boundaries of THz-RIS technology, understanding its upper limits in terms of spectral efficiency, beam forming and steering capabilities, and real-time optimization adaptability under realistic conditions. Additionally, simulations and evaluations are being extended to a higher frequency band of **140 GHz**. Next, simulations and evaluations have been extended to higher frequency bands of **140 GHz and 300 GHz**. This transition involves critical modifications, specifically in the positioning and distances among the transmitter, RIS, and receivers to accurately reflect realistic THz scenarios in conjunction with an XR indoor scenario. Evaluations at 140 GHz similarly cover three categories of RIS elements: low N, high N, and Hyper-N scales.

The proposed system model represents an indoor XR meeting use case scenario operating at a frequency of **140 GHz**. It comprises a transmitter (TX) positioned at coordinates (3, 0, 3) m and a Reconfigurable Intelligent Surface (RIS) at coordinates (0, 0, 3) m, strategically placed opposite each other. Four receivers (RX1 at (5, 5, 1)m, RX2 at (6, 4, 1)m, RX3 at (4, -4, 1) m, and RX4 at (5, -5, 1)m) form two distinct user clusters located outside the direct coverage area of the TX. The RIS utilizes directional beamforming to steer signals effectively toward these user clusters. Different receiver positions are explored to analyse the impact of directional beamforming on system performance.

Part 2: DRL simulations In parallel to evaluating RIS performance at both **28 GHz** and **140 GHz**, we further extended our efforts on the **Deep Reinforcement Learning (DRL)** front, emphasizing algorithmic refinement and accuracy enhancements. The theoretical and mathematical foundations of the algorithm and process is well explained as D4.1 and D4.2.

Specifically, we conducted detailed fine-tuning of our **Deep Deterministic Policy Gradient (DDPG)** algorithm. We varied critical hyperparameters, notably increasing the number of training episodes to achieve more robust learning convergence, especially when dealing with a higher number of RIS elements (high N). This fine-tuning effort allowed us to carefully observe the algorithm's behaviour over extended training periods, the inspiration of which was taken from [Saglam2023] facilitating the identification of optimal configurations to maximize spectral efficiency gains. Furthermore, to ensure precise, realistic, and reliable DRL training, we integrated our refined DRL framework directly with accurate channel data generated by MATLAB's detailed THz-RIS channel modeling. This integration ensured that the DRL agent learned and optimized the RIS configurations based on realistic channel characteristics, closely reflecting practical deployment scenarios. Consequently, this led to improved adaptability and spectral efficiency performance of the DRL-driven RIS system. This refined and carefully integrated DRL approach is now being consistently applied and validated across both frequency bands, allowing us to accurately assess, compare, and confirm the effectiveness and scalability of our optimized DRL solutions for a wide range of RIS element configurations (low, high N) and deployment scenarios. The advanced results are provided in the Figure 51 below. As our next step, we extended our investigation to **140GHz** scenario to explore the feasibility and potential benefits at higher THz frequencies Figure 51.

4.2.3.3 Final evaluation results




Figure 51: First advanced improved DRL results showing the efficacy of DRL

Analysis of the Improved DRL Results at 28 GHz:

The graphs clearly illustrate significant advancements and improvements in the performance of the Deep Reinforcement Learning (DRL) algorithm compared to previously employed methods, namely the Cosine Similarity algorithm and the Random Phase method.

Compared to the earlier simulations (left-hand graphs), the updated DRL simulations (right-hand graphs) show remarkable improvement in the overall spectral efficiency and scalability with increasing N. The previous DRL implementation (left graphs) had modest performance gains and sometimes saturation at higher values of N. However, the refined DRL method ("DRL E10") on the right clearly resolves this limitation, showcasing significantly better results, especially noticeable at moderate-to-high RIS element counts (N=100 to 250 and up to 400). The fine-tuned DRL algorithm achieved through more precise training episodes and accurate integration of MATLAB's THz-RIS channel data enabled precise channel modeling, thus leading to more realistic learning and better optimization decisions.

THz-RIS performance evaluation at 140 and 300 GHz cover three categories of RIS elements: low N, high N, and Hyper-N and are shown in Figure 52.




Figure 52: Average spectral efficiency (SE) versus number of RIS elements N for MU-MIMO and MU-MISO systems across low, high, and hyper N regimes. Results are shown for random and cosine phase configurations at 140 GHz and 300 GHz.

Insights from the results:

The DRL approach consistently outperforms the other two methods across all tested values of RIS elements (N). At lower RIS element counts (up to about $N=50$), DRL demonstrates rapid improvement in spectral efficiency (**Average SE**), clearly exceeding both the cosine and random methods. At higher RIS counts ($N > 100$), DRL shows a robust upward trend, achieving higher spectral efficiencies than previously recorded, highlighting the effectiveness of the refined DRL training and channel modeling integration. The previous DRL implementation as in [D4.2] also the left graphs in Figure 51 had modest performance gains and sometimes saturation at higher values of N . However, the refined DRL method ("DRL E10") on the right clearly resolves this limitation, showcasing significantly better results,

especially noticeable at moderate-to-high RIS element counts ($N=100$ to 250 and up to 400) One of the main reasons for these improved results is credited to the fine-tuned DRL algorithm achieved through more precise training episodes and accurate integration of MATLAB's THz-RIS channel data thus leading to more realistic learning and better optimization decisions. This precise modeling directly resulted in higher achievable spectral efficiency in practical scenarios.

Additionally, from the simulations in Figure 52, we draw the following key observations:

Spectral efficiency (SE) increases consistently with the number of RIS elements N across all scenarios, confirming the benefit of larger surfaces for both MU-MIMO and MU-MISO systems.

Random phase configurations outperform cosine-based designs in both MIMO and MISO setups, especially at lower N , suggesting higher diversity or constructive scattering with randomized phases.

MU-MIMO consistently achieves higher SE than MU-MISO, indicating that spatial multiplexing gains are better exploited when both transmitter and receiver are equipped with multiple antennas.

Performance gains saturate in the hyper N regime, with diminishing returns as N grows beyond a certain point, particularly visible for cosine phase configurations.

Low N regimes show severely limited SE, with values remaining well below 1 bit/s/Hz in many cases, highlighting the inefficiency of small RIS deployments without proper phase optimization.

Importance of Maximum Spectral Efficiency (SE)

The theoretical importance of achieving maximum spectral efficiency (SE) in RIS-based MIMO systems is clearly highlighted by [\[Yigit2020\]](#), who emphasize optimal RIS configurations as essential for ideal communication performance. The provided graph at 140 GHz shows the theoretical maximum SE (maxSE) represented by the black line, which serves as an upper bound achievable with perfect channel state information. Notably, our Deep Reinforcement Learning (DRL) method closely approaches this theoretical limit, significantly outperforming conventional methods such as the cosine similarity and random phase approaches, especially at higher numbers of RIS elements (N). The narrow performance gap between the DRL and maxSE demonstrates the DRL's robust optimization capabilities, reinforcing the accuracy of channel modeling and algorithm refinement in our approach.

The **black line** labeled "maxSE" clearly represents this theoretical maximum SE. It illustrates the idealized scenario of perfect RIS configuration. The **brown dotted line** labeled "DRL" closely approaches the max SE curve, particularly noticeable at moderate and high RIS element counts ($N > 100$) Figure 53. The small gap between DRL and max SE highlights how effectively your DRL-based method approaches the theoretically optimal scenario. Conversely, the Cosine Algorithm and Random Phase methods remain significantly below the max SE across all RIS sizes, emphasizing their suboptimal nature in practical settings.

The improvements made in the DRL algorithm with Large N for 28 GHz and 140 GHz is shown in Figure 53 One of the main reasons for these improved results is credited to the fine-tuned DRL algorithm achieved through more precise training episodes and accurate integration of MATLAB's THz-RIS channel data thus leading to more realistic learning and better optimization decisions. This precise

modeling directly resulted in higher achievable spectral efficiency in practical scenarios for both **28 GHz** and **140 GHz** see Figure 53.




Figure 53: Spectral efficiency vs RIS elements calculated with different phase optimization approaches such as DRL, cosine algorithm and random phases.

4.2.3.4 Take aways and Conclusions

The presented results strongly encourage further exploration and confidence in extending DRL simulations to very high RIS element counts **Hyper N** scenarios as indicated in the first part, predicting even higher gains in spectral efficiency.

Given the promising nature of the refined DRL performance at **28 GHz** and **140GHz**, we are optimistic about extending these evaluations to greater RIS elements at **140 GHz** along with new explorations at 300GHz. At higher frequencies, where path losses and environmental complexities are even more pronounced, DRL's adaptability will likely yield even more compelling performance improvements. We also intend to explore multi-RIS scenarios with our further developments.

5 THE 6G-XR TRIAL CONTROLLER

5.1 FINAL TRIAL CONTROLLER IMPLEMENTATION

The architecture of the 6G-XR trial controller is structured to facilitate seamless experiment management across different geographical locations. It consists of a Unified Web Portal, which acts as the central point for user management, trial management, and facility selection.

This portal connects to two region-specific web portals:

- South Web Portal (linked to the 6G-XR Facility in Madrid - Barcelona, Spain)
- North Web Portal (linked to the 6G-XR Facility in Oulu, Finland)

Each of these portals serves distinct roles in managing experiments, NST (Network Slice Template) handling, transportation, and data visualization. The architecture enables secure remote access for experimenters, ensuring effective control over trials and facility interactions. The current implementation does not enforce HTTPS or multi-factor authentication (MFA) directly at the portal level. However, access to the portals is restricted through a secure internal VPN, which acts as a gateway to the system. This VPN ensures that only authorized users can reach the login interface. The VPN supports various authentication mechanisms, including username and password, One-Time Passwords (OTP), and Passkey/WebAuthn-based authentication. These methods are available to users but are not enforced by default. Users who voluntarily enroll in any multi-factor or passwordless authentication method will be required to use it during login. All users are required to use MFA only if they have previously enrolled in one. Adding TOTP, Passkeys, or other WebAuthn-based methods is voluntary unless enforced by access policy.

The inclusion of APIs and node adapters further facilitates interoperability between different facilities.

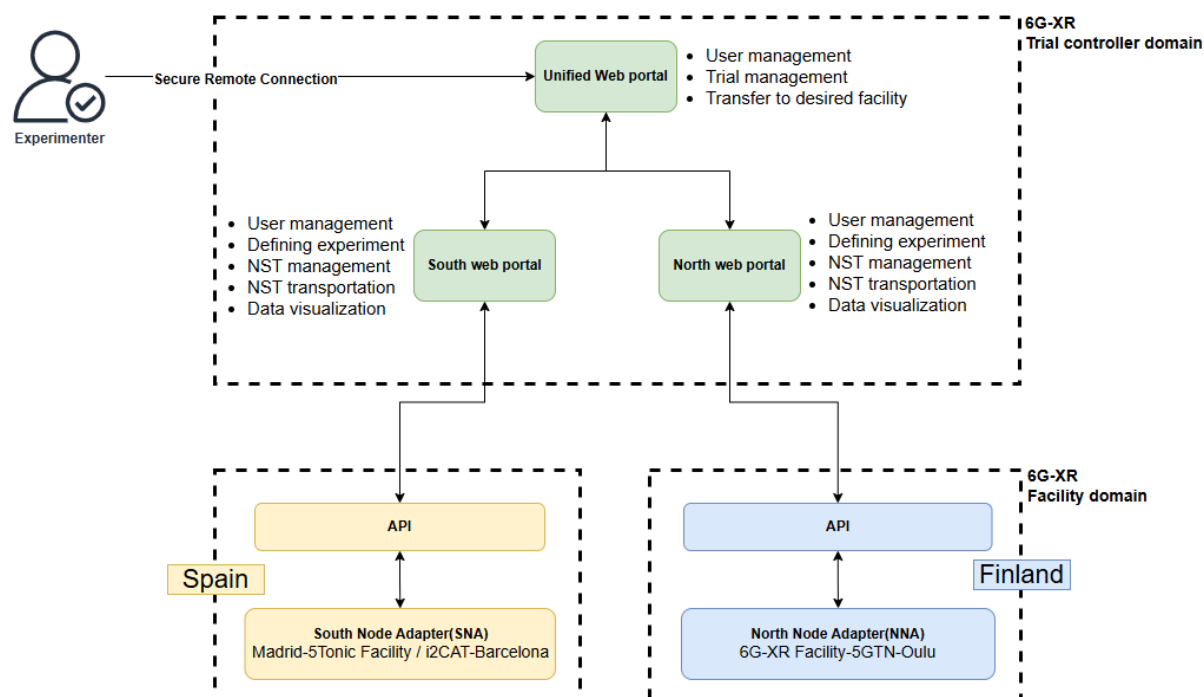


Figure 54. Trial Controller Architecture

This section provides a comprehensive explanation of the final implementation of the Trial Controller, detailing its core functionalities and system architecture. It covers the integration of different components, including the Unified Web Portal, facility portals, API interactions, and experiment management processes.

The trial controller has three main components. The unified web portal, the north web portal, and the south web portal. By authenticating and entering the unified web portal, the user can define the trial and be transferred to the desired portal. Then, in the north and south web portals, the user defines the Network Slice Template (NST), creates an experiment, send NST to the node adapters (NNA/SNA) and finally controls the experiments, including edit, start, stop, delete and status of the experiments. The following subsections will break down each module, explaining its role in ensuring efficient trial execution and seamless experiment handling.

5.1.1 Unified Web Portal

The unified web portal is designed with the aim of managing the trial on both the north and south web portals. The unified web portal is implemented using the Django framework, including two Django applications to handle user authentication, registration, permissions, and trial management. As a frontend, it uses React.js and uses PostgreSQL as a database. By logging into this system, the user can create a trial and then, according to the user's needs, the user is redirected to the desired portal (North or South). The main key of communication between the unified web portal and the north and south portals is the trial ID. By creating a trial, a unique number and a trial facility are assigned to it. Each of the North and South web portals can only access and view trials that were assigned to their own facility during the trial generation process in the unified portal. Some features of the Unified web portal include:

- User management: including registration, logging in and out, forgetting password, and editing user information.
- Trial management: create, delete, and edit the trial.

- Transfer to facility: By clicking on the facility, the user is redirected to the desired web portal facility (North or South).

The following Figure 55 to Figure 59 illustrate the main sections of the Unified web portal.



6G-XR Home Page 

Welcome to Unified Portal

Trial Management



Sign in to your account

Username

Password

6+ characters 

[Forgot password?](#)

Sign In

[Don't Have an Account? Sign Up Now](#)

Figure 55. User login page in the Unified web portal



6G-XR Home Page 

Welcome to Unified Portal

Trial Management



Let's Get Started

Already have an account? [Sign in](#)

First name

Last name

Email address

phone number

Username

Password

6+ characters 

Create account

By signing up, I agree to [Terms of service](#) and [Privacy policy](#).

Figure 56. User signup page in the Unified web portal

UNIFIED WEB PORTAL

Trial List

About Unified

in X 6G_{XR}

Logout

Trial List

Add new Trial

Trial ID	Trial Name	Start Time	End Time	Facility	description	Action
58	test1	2025-03-13 00:00:00	2025-03-16 00:00:00	Oulu	test1	
59	test 2	2025-03-13 02:00:00	2025-03-14 01:00:00	South	test2	
60	test 3	2025-03-13 03:00:00	2025-03-14 03:00:00	South	test3	
61	test4	2025-03-15 00:00:00	2025-03-17 10:00:00	Oulu	test4	
62	test5	2025-03-17 02:00:00	2025-03-18 02:00:00	Oulu	test5	
63	test6	2025-03-19 00:00:00	2025-03-21 00:00:00	South	test6	

1-6 of 6 < >

Figure 57. Unified web portal

UNIFIED WEB PORTAL

Trial List

About Unified

in X 6G_{XR}

Logout

Add Trial

Trial Name

Start Time

04/17/2025 01:16 PM

End Time

04/18/2025 03:00 AM

Facility Option

Oulu

Description

test

Create Trial

Figure 58. Unified web portal- trial definition




Figure 59. User profile in the Unified web portal

In Figure 60 to Figure 63 the sequence diagrams related to the Trial controller are presented, which include Trial Management and User Management. Since the architecture and design of applications differ between the North and South portals, the sequence diagrams have been provided separately for each portal.

5.1.1.1 Trial Management

The Trial Management sequence diagram illustrates the process of handling trials within the system. It begins with a user creating a trial, where the system stores the start and end times in UTC format. Then, the user can select a facility and is redirected to the facility portal login page, carrying the Trial ID managed by the frontend application. After logging in, the user can proceed with experiment creation. Trials can be edited only before they reach their end time. All the trial management process runs similarly, both when the Unified Web Portal interacts with the North Node and South Node test facilities.

The only difference is the APIs to delete the experiment when the trial ID is deleted. If a user decides to delete a trial, the Unified Web Portal ensures that all associated experiments in the facility portals are also deleted. This is executed by calling an API endpoint to remove related experiments, including those running in North Node Adapter (NNA) and South Node Adapter (SNA), ensuring a complete cleanup of trial-related data. In the case of the North Node, the Unified Web Portal calls the North web portal APIs, which subsequently call the NNA, while in the case of the South Node, it directly calls the SNA.




Figure 60. Trial Management– Northern Portal Connection




Figure 61. Trial Management–Southern Portal Connection

5.1.1.2 User Management

The User Management sequence diagram details the registration, authentication, password recovery, and deletion processes within the system between the Unified Portal, the North Portal, and the South Portal.

In the case of the North Node, the Unified and facility portals have the same credentials. It means that the user only registers in the Unified Web Portal, which triggers an API request to create the user in the North Portal with the same credentials. Users can then edit their profiles, and the system ensures synchronization between portals through API calls. During login, the system authenticates users and provides access to the necessary functionalities. If a user chooses to delete their account, the system first removes all trials and associated experiments, followed by deleting the user from both the Unified and North portals. The system ensures proper data integrity by handling deletions systematically through API requests to remove user data and related experiments.




Figure 62. User Management Across Unified and Northern Portal

In the case of the South Node, Unified and facility portals do not use the same credentials. It means that the User Management in the South Node facility works differently than the North Node one, as the user credentials at the facility portal are not the same of the Unified portal. In the South Portal, the admins of the MEC Orchestrator in South Node assign the experimenter an appProviderId, which will be used for the login. With this, together with the user and password, a record in the MongoDB database will be saved.

When performing login with user and password, a session record will be saved on MongoDB with the sessionId (which is the unique identifier MongoDB has for the record) and the generated OAuth token from MEC Orchestrator API.




Figure 63. User Management Across Unified and Southern Portal

5.1.2 North Node Web Portal and Adapter

The North web portal is designed to manage the experiments. The North web portal is implemented using the Django framework, including two Django applications to handle users and experiments. As a frontend, it uses React.js and uses PostgreSQL as a database. The North web portal handles experiment actions and fetches KPIs from Qosium to visualize results in the dashboard. By logging into this system, the user can create an experiment including trial name, target facility, start and end time, application name, and selecting slice type.

When a user creates an experiment, an NST (Network Slice Template) is generated, which can then be sent to the NNA via an API call. The user is able to control the experiment, including editing, starting, stopping, deleting, and checking its status. Additionally, the North Portal includes a dashboard to visualize data retrieved from Qosium. This is achieved using APIs and the Grafana plugin for real-time data display.

The following Figure 64 to Figure 68 illustrate the main sections of the North web portal.




Figure 64. User login page in the North web portal

My Experiments

Trial ID	Experiment	Start Time	End Time	Node	Slice Option	Slice 1	Slice 2	App1	App2	Status	Actions
172	EXP1	2025-05-31 00:00:00	2025-06-02 00:00:00	North	no_slices			app1	app2		▶ ✎ 👁 🗑
172	EXP2	2025-05-31 00:00:00	2025-06-04 00:00:00	North	one_slice	uRLLC		app3	app4		▶ ✎ 👁 🗑
172	EXP3	2025-06-03 00:00:00	2025-06-04 00:00:00	North	two_slices	uRLLC	eMBB	app5	app6		▶ ✎ 👁 🗑

page: 1 | 3-3

Figure 65. North web portal

Create a New Experiment

Trial Name: test

Experiment Name: EXP4

Start Date: 05/31/2025 02:17 AM

End Date: 06/02/2025 03:19 AM

Node: North

Facility: Oulu

Application 1 (optional): app5

Application 2 (optional): app8

Slice Option: Two_slices

slice 1: ☒ uRLLC ☐ eMBB

slice 2: ☒ uRLLC ☐ eMBB

Create Experiment

Figure 66: North web portal-experiment definition




Figure 67: User profile in the North web portal

In Figure 68 the sequence diagram related to Experiment Management in the Northern web portal is presented. The sequence diagram outlines the process of creating, editing, and deleting experiments within a trial. The process begins when a user selects the Oulu facility, which redirects them to the North Portal login page. The user must log in using the same credentials from the Unified portal to proceed. Once logged in, the user can create an experiment, which is linked to a Trial ID. The Trial ID is stored on the frontend application. Before allowing experiment creation, the system checks if the trial has not expired by calling an API in the Unified portal. If the trial is still active, the experiment can be created. Users can also edit experiments, but before making any changes, the system again verifies if the trial is still valid by checking its end date. If the trial is active, the modifications are allowed. When a user deletes an experiment, the system ensures that all data associated with that experiment will also be deleted from the North Portal. Additionally, an API call is made to the NNA (North Node Adapter) to delete any running experiments related to that experiment. The deletion process ensures that no residual experiments remain in the system, maintaining data consistency.




Figure 68: Experiment Management in the Northern web portal

The Unified and North portals are tightly integrated to enable seamless user and experiment management across the system. This integration is achieved through a set of API communications that allow data synchronization and coordinated workflows between the two platforms. When a user registers in the Unified portal, their information is simultaneously sent to the North portal via a dedicated API to create a matching user account. This ensures that users can log in to both systems using the same credentials, maintaining a consistent identity across the environment.

Experiment management in the North portal is directly dependent on the trials created in the Unified portal. When a user attempts to create or modify an experiment in the North portal, the system first checks the validity of the corresponding trial by querying the Unified portal through an API call. Only if the trial is still active (not expired), the operation is permitted. This ensures data integrity and prevents the creation of experiments based on invalid trials.

Additionally, when a user deletes a trial or account from the Unified portal, the system initiates a cascading deletion process. It sends API calls to the North portal to delete all related experiments and further communicates with the North Node Adapter (NNA) to terminate any running experiment instances. This approach ensures that no orphaned resources remain in the system and that user-related data is completely removed from all subsystems.

Overall, this integration design provides a reliable, automated, and secure connection between the portals, enabling users to manage their trials and experiments efficiently from a unified interface.

North Node adapter (NNA) is responsible for configuring the North Node facility components according to the Network Slice Template (NST) and starting/stopping trials. The NNA manages the setup and management of network slices and KPI collection and communicates with the underlying resources and management technologies. In other words, NNA acts as a middleware to hide the implementation details of the trial background processes the user can manage via North Node web portal.

NNA provides a REST API towards the North Node web portal, by which the portal can create and start trials. Through the API the web portal can also stop and delete existing trials. The North Node web portal can check the status of trials running in NNA by sending a status request through the REST API. NNA replies with the state NNA is at (“initialized”, “executing”, “terminating”, “done”) or with an error text if something has gone wrong with the running trial.

5.1.3 South Node Web Portal and Adapter

The South Node Web Portal is developed to manage the experimental processes at the South Node test facilities (5Tonic and i2CAT) through the interconnections with the Unified Web Portal, where the trial IDs are managed, and the South Node Adapter (SNA), in charge of deploying experiments and managing the monitoring components based on Qosium.

After accessing the Unified Web Portal and selecting South Node (see Section 5.1.1), the user is redirected to the South Node Web Portal, whose starting page is shown in Figure 69. This page is employed by the experimenter to register a new account during the first access to the portal (left figure) or to login for all the following accesses (right figure). Only experimenter previously registered at 5Tonic and provided with a unique identifier (“APP Provider ID” in the left figure) are allowed to register and login through the web portal. The login information is not the same as the Unified and North Node Web Portals.




Figure 69: Register form (left) and login form (right) of the South Node Web Portal

Upon logging in, the experimenter can interact with two different sections of the South Node Web Portal: *Dashboard* and *Network Slice Template (NST)*. The *Dashboard* section, shown in Figure 70,

provides a link to access the Grafana dashboard where the monitoring information collected by Qosium is shown. The KPIs collection and visualization is explained in Section 5.2.2.





Figure 70: Dashboard section of the South Node Web Portal

The NST section is instead shown in Figure 71 and allows the experiments to access the functionalities related to the experiment management, including listing of already configured experiments and the creation of a new one. In the experiment list, only basic information is shown, such as *Experiment Name*, *Node*, *Start Time*, *End Time* and *Status*. Additional information can be found by clicking on *Application* and *GPIs*, thus triggering the pop-ups presented in Figure 72.



#	Experiment Name	Node	Applications	GPSIs	Start Time	End Time	Status	Actions
1	Experiment-3322	SOUTH			25/03/2025 04:51 PM	30/03/2025 05:51 PM	CREATED	
2	Experiment-1234	SOUTH			25/03/2025 04:51 PM	30/03/2025 05:51 PM	CREATED	
3	Experiment-182	SOUTH			25/03/2025 04:51 PM	30/03/2025 05:51 PM	CREATED	

Create new NST

6G-XR User
6gxuser@6gxr.com
Logout

Figure 71: NST section of the South Node Web Portal.



Figure 72. Pop-ups to show further information of the experiment at the South Node

Figure 73 presents the form employed to create a new experiment. Among the information to be introduced, Trial ID, Start Time and End Time are provided by the Unified Web Portal, while all the rest are chosen by the experimenter. Within the experiment form, *Application & Slice* section allows to select the applications previously onboarded at South Node Edge infrastructure, the location where to deploy the, the slice where to attach them, and the processing resources needed. *GPSIs & Slice* section is instead meant for selecting the User Equipments (UEs) and connect them to specific edge location by means of the appropriate slices.




Figure 73. Form to create a new experiment in the South Node

Finally, once all the information is inserted, the Network Slice Template (NST), associated with the experiment, is generated and sent to the South Node Adapter (SNA) via its REST API. The information is stored, and the experiment created by not executed yet. After this, the experimenter keeps the control over the experiment, as he is enabled with several actions including start, stop and delete the experiment.

In Figure 74, the complete workflow for experiment management at the South Node test facilities is presented. It includes the interaction of the experimenter with Unified Web Portal, South Node Web Portal and South Node Adapter.

The process begins when the user accesses the Unified Web Portal and selects the South Node to execute the experiment with a specific Trail ID. At the South Node Web Portal, the user specifies key parameters such as the experiment name, applications and the slices. Here, the South Node Web Portal is integrated with the Unified such to retrieve Trail ID information, including start and end times of the trial. Once the experiment is defined, the system automatically generates the associated NST, encapsulating the experiment configuration in a JSON format suitable for the deployment.

Following this, the user through the South Node Web Portal submits the NST to the South Node Adapter (SNA) using an API call. This submission acts as a registration mechanism, effectively storing the information of the experiment. Once submitted, the experiment becomes manageable via the portal. Finally, the user can choose to run, stop or delete the experiment, as well as to access the monitored metrics (see Section 5.2.2).




Figure 74. Experiment management with Unified Web Portal, South Node Web Portal and South Node Adapter

The SNA is responsible for configuring the components of the South Node facility in accordance with the specifications defined in the NST and for managing the initiation and termination of experiments. It oversees the setup, operation, and administration of network slices, as well as launching the KPIs monitoring components, by interfacing with the underlying resources and technologies, such as IEAP, NEF and Qosium. In essence, the SNA functions as middleware, abstracting the implementation details of the background processes, thereby simplifying the infrastructure management. The SNA exposes a REST API to the South Node Web Portal, through which the portal is able to create and initiate experiments. Furthermore, the API allows the web portal to monitor the operational status of each experiment, namely, "initialized," "running," "terminated," or "failed".

Since [\[D4.2\]](#) there have not been significant changes on the Trial Controller flow regarding interaction between the different components of the South Node platform, e.g South Node Web Portal and South Node Adapter. The following features have been implemented:

- Decoupled experiment creation from experiment launch, and experiment termination from experiment deletion in the API and the flow.
- Only one experiment can be run at a time. This is controlled by a database table with a single registry where the lock boolean is stored. In case of an experiment being run, terminate instance method will set the lock to false regardless of the success of the termination to prevent the application from getting stuck.

- Implemented a delete experiment by trialId so the Unified Web Portal will be able to perform cascade deletion when the trial is deleted from the Unified Web Portal.
- Performed necessary routing for the experimenter to access South Web Portal and South Node Adapter from North Node VPN (5GTN), which will be the only VPN the experimenter will need access to.

Integration between Unified Web Portal and South Portal

South Web Portal uses NGINX routed URLs to be functional inside North Node VPN, so deployment of South Node frontend and backend will be only available on South Node VPN, while testing South Web Portal will only be available in North Node VPN. Frontend local instance testing will be only available in North Node VPN, and backend local instance testing will be only available on South Node VPN.




Figure 75 Integration between Unified Web Portal and South Node Portal

5.2 KPIS COLLECTION AND VISUALIZATION

Trial controller uses Qosium as the KPI collection engine. Qosium is a bundle of software tools that allows measuring QoS in a network technology agnostic way. Qosium is used to measure and obtain QoS KPIs such as latency, throughput, jitter, and packet loss ratio during trial runs.

In our trial controller solution, KPI measurements are controlled by a measurement component called Qosium Storage. It is a server and database software which provides a REST API for starting/stopping measurements and retrieving results for programmatical use in North and South node adapters and web portals.

Measured KPIs are visualized in both North and South node portals using Grafana dashboard. The measurement logs (measured KPI values at one second interval) are stored in Qosium Storage at both nodes. From Qosium Storage the logs are available for download for further study with Qosium Scope, for example. The KPI data stored in Qosium Storage are persistent (held in database indefinitely).

5.2.1 North Node monitoring

The North web portal manages the full lifecycle of an experiment, from initiation to result visualization. The process begins when the user clicks the “Play” button to send the Network Slice Template (NST). This triggers two POST requests to the endpoint <Server IP>:8080/home/api/submit/ and 172.29.6.16:5000/experiment, which sends the NST to the North Node Adapter (NNA). Upon receiving the NST, the NNA initiates the experiment by starting per-slice measurements using Qosium. A function is configured to get measurement data every 50 seconds.

These measurements are requested through a POST call to 127.0.0.1:9090/home/api/sendquery/. Once received, the North portal stores the data in the local database for future access and visualization purposes. After the experiment starts running, the user can select four KPIs from the dashboard. The selected KPIs are saved into the SelectedFields table using a POST request to <Server IP>:8080/home/api/grafana/. The system then extracts the corresponding results from the database via a GET request to the same endpoint. These metrics are visualized in real-time using Grafana dashboards, which are embedded into the North portal interface. Users also have the ability to download the results. This is done by calling the API to <Server IP>:8080/home/api/downloadcsvview/ using the POST method. The backend extracts the selected KPI data from the database and returns it to the user's browser as a downloadable file, typically in CSV format.

Finally, the user can terminate the experiment at any time. The termination request sends DELETE commands to both the North portal API <Server IP>:8080/home/api/terminate/<id>/ and the NNA API http://172.16.226.56:5000/experiment/<id>, ensuring the experiment is completely stopped and The user can only access and use the results that were available up to the point when the experiment was stopped. To start a new experiment, the user must first delete the stopped one and then create a new experiment. This guarantees consistency and prevents leftover data or active sessions from persisting beyond their intended duration.




Figure 76. Real-Time Visualization via Qosium and Grafana




Figure 77. Grafana dashboard in the North web portal

5.2.2 South Node monitoring

We have deployed at the servers of 5TONIC an instance of Qosium Probe, by using a .deb installation at OS level. The Qosium Probe acts as a client that will provide this data to the main Qosium Storage on I2CAT infrastructure, which will provide the metrics for the monitoring app.

On API level, when an experiment is run, the South Node Adapter will call the Qosium Storage API, which will start the monitoring on the machine where Qosium Probe was installed. On experiment termination, the SN Adapter will store the id to later use it to call the Qosium API again to stop the monitoring.

At the level of SN Web Portal, it was decided that we will show a permanent link to Qosium monitoring interface that will list all instances currently run in the machine. With the change on SN adapter in 5.1.3 explained above, only one experiment will be run at a time, so the experimenters will be able to tell them apart by filtering by date ranges that match the startTime and stopTime in their trial's NST.

The SN metric collection system is designed considering that experiments will consist of application flows that will run between an edge node, namely MAD or BCN, and a CPE. The number of CPEs available to experiments is limited to two (CPE1 and CPE2), and the experimenter will be informed about the CPE that has been assigned to run the experimenter. Figure 78 represents the interface that allows the experimenter to retrieve the experiment results. First, at the top-right corner the experimenter can select the time interval during which its experiment was run. Second, on the left top corner, the experimenter selects the experiment ID, which is the same that appears in the dashboard. Third, the experimenter selects the CPEs that have been used in the experiment. Finally, the experimenter clicks on the "zoom to measurement" field to retrieve its experiment results. The metrics displayed are all the ones made available by Qosium. Figure 79 depicts an example of metrics collected in the SN for an experiment transmitting a ping application, showing throughput, dropped packets and malformed packets.

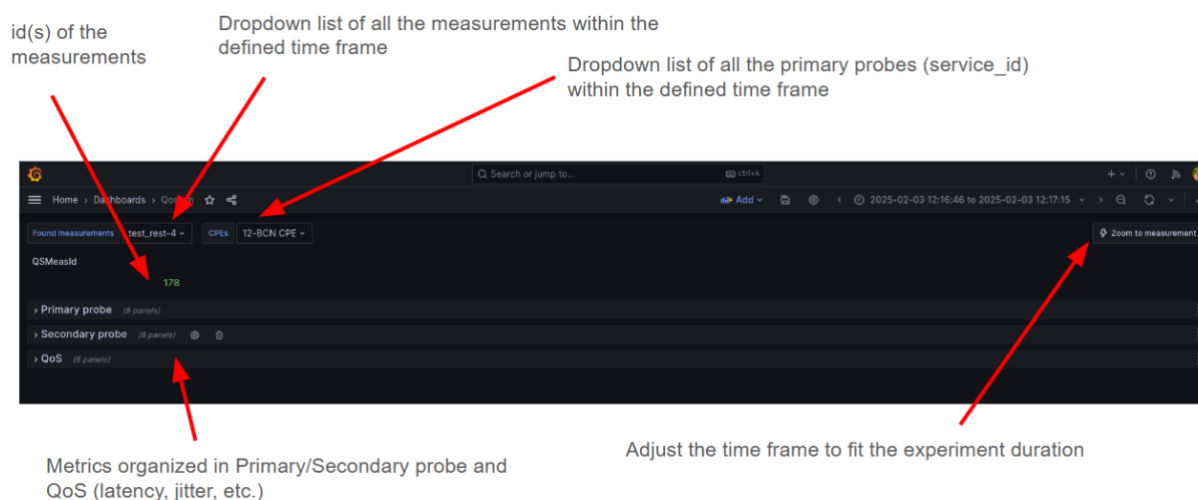


Figure 78. Configuration interface to enable the SN experimenter to retrieve results




Figure 79. Example of metrics retrieved through the SN metric collection system

6 SUMMARY

This deliverable completes the description of the XR enabler technologies proposed in the project, which are verified on the 6G-XR distributed nodes – North Node (UOULU 5GTN and VTT 5GTN) and the South Node (5GBarcelona and 5TONIC). In addition, it completes the necessary activities towards successful deployment of trial controlled used for managing experimentation activities.

The XR enablers are categorized based on three main pillars: (1) 3GPP XR Enablers, (2) O-RAN XR Enablers and (3) 6G disruptive XR Enablers. For each enabler, a motivation is given leveraging the previous deliverables [\[D4.1\]](#) and [\[D4.2\]](#), and the final solution design and main takeaways are provided. It also describes the trial-controller used for managing and orchestrating 5G infrastructure resources during experimentation. Enablers are validated using above-mentioned experimentation nodes, which are therefore evolving to demonstrate the performance of key Beyond 5G/6G candidate technologies, components, and architectures. Used experimentation nodes will remain valid not only in the short-term but also in the mid and long-term.

This report completes the final deployment and validation results achieved in this project for the new enablers and the trial controller. The final trial controller implementation was completed via Unified Web Portal, North Node Web Portal and Adapter, and South Node Web Portal and Adapter. The corresponding KPIs were also collected and visualized in both the North Node and South Node.

The following table summarizes the main outcomes, achievements and synergies with other WPs for each the enablers studied in WP4.

Table 10: Summary of main characteristics and achievements of the XR enablers and Trial Controller.

Category	XR Enabler	Use Case	TRL	Main achievements in WP4 and synergies with other WPs
3GPP XR Enablers	Network-Controlled Repeaters in indoor environment [IT]	UC1 – Holographic communications (XR gaming in the Vienna Simulator)	3	NCR performance was evaluated in indoor scenarios using both simulator-based and experimental XR traffic models. Baseline implementation and integration feasibility assessed in WP4.
	Network assisted Rate Control API [I2CAT]	UC1 – Holographic communications	4	This enabler will be used in WP6 as part of UC1, but integration will be limited to a laboratory environment, since the counters required for an accurate rate recommendation are not available in the commercial Ericsson gNBs deployed in the SN.

	ATSSS based capacity enhancement [I2CAT]	UC1 – Holographic communications	4	This enabler has been designed considering the service requirements of the HoloMIT application developed in WP3 (i.e. holograms identified as TCP flows). However, in this deliverable the enabler has been validated in laboratory conditions. Integration with the HoloMIT service is left as future work
	Energy and Application aware management of 3GPP infrastructure [VTT]	UC1-3 – Holographic communications UC4 – Collaborative 3D Digital Twin-like Environment UC5 – Energy Measurement Framework for Energy Sustainability	4	The impact of the different energy saving states to XR application scalability assessed with end-to-end measurements in a live network. XR applications from UC1-4 (WP3 and WP6) emulated with traffic generators for the tests. RAN energy saving states from UC5 (WP5) exploited for the measurements.
	Upgrade and evaluation of experimental RAN infrastructure in SN and NN [ERI, OULU, NOK]	SN: UC1-3 – Holographic communications	7	The campus deployments carried out both in SN and NN will be used to support the experiments proposed by the OC3 winners.
O-RAN XR Enablers	Congestion aware load balancing [I2CAT]	UC1 – Holographic communications	5	The CDF component of this enabler will be demonstrated together with the HoloMIT service developed in WP6. An initial demonstration will be performed at EuCNC using a portable network setup, and a final demonstration will be showcased as part of UC1, using the campus network in the SN.
	Energy-aware end-to-end resource management [VTT]	UC1-3 – Holographic communications UC4 – Collaborative 3D Digital Twin-like Environment UC5 – Energy Measurement	3	The achievable gains in the energy consumption of the O-RAN platform verified with end-to-end measurements in a live network. XR applications from UC3-4 (WP3 and WP6) emulated with traffic

		Framework for Energy Sustainability		generators for the tests. Energy measurement framework from UC5 (WP5) exploited for the measurements.
	O-RAN enabled slicing to support XR services [OULU]	UC4 – Collaborative 3D Digital Twin-like Environment	FR1: 8 FR2: 7	OAIBOX O-RAN Implemented with both FR1 and FR2 radio units. Slice isolation tested for both cases in UC4 environment.
Disruptive XR Enablers	High-Frequency Transceivers for THz-RIS and ISAC [IMEC, OULU]	UC1 – Holoportation UC4 – Collaborative 3D Digital Twin-like Environment	3	The 140 GHz radio modules tested over the air with relevant digital modulations. The 300 GHz chip development done and tested. The 6G signal candidate integrated to 300 GHz test system.
	Baseband implementation for THz-RIS and ISAC based on SC-FDE [IDE]	UC1 – Holoportation UC4 – Collaborative 3D Digital Twin-like Environment	3	Baseband implementation based on Matlab blocks tested via simulations. WP4 integration with High-Frequency Transceivers via IQ files in the TX path.
	Deep Reinforcement Learning for THz-RIS [IT]	ML based THz-RIS for XR traffic adaptation & beam optimization	3	DRL algorithm integrated into system-level simulator to optimize beamforming and RIS phase shifts. Preliminary results show improved SINR and throughput under XR scenarios
6G-XR Trial Controller	Unified Web Portal [OULU]	UC4 – Collaborative 3D Digital Twin-like Environment	6	A web-based GUI (Unified Web Portal) developed for users to create and manage their trials, and then access the North or South web portals through it to execute their experiments and view measured KPIs in real-time.
	North Node Web Portal and Adapter [OULU]	UC4 – Collaborative 3D Digital Twin-like Environment	6	A web-based GUI (North Node Web Portal) developed for users to run experiments at the North node and view measured KPIs in real-time.

				North Node Adapter, which is responsible for fetching 2-slice KPI data from Qosium and using the data with AI to dynamically adjust bandwidth allocation.
	South Node Web Portal and Adapter [VICOM / CGE]	Not employed by 6G-XR use cases, but for the Open call 3 vertical use cases.	6	A web-based GUI (South Node Web Portal) available for experimenters to run trials at the South node infrastructure, including 5Tonic and i2CAT. South Node Adapter to interpret inputs from the web-based GUI, deploy experiments and collect metrics through Qosium integration.

7 REFERENCES

- [3GPPTS24501] 3GPP TS 24.501 V19.1.1, "Non-Access Stratum (NAS) protocol for 5G System (5GS)," December 2024. Available [here](#).
- [D4.1] Deliverable D4.1, "State-of-the-art analysis and initial design of beyond 5G RAN, core, and open-source networks, disruptive RAN technologies and trial controller", [Online]. Available [here](#).
- [D4.2] Deliverable D4.2, "Intermediate deployment of beyond 5G RAN, core, and open-source networks, disruptive RAN technologies and trial controller", [Online]. Available [here](#).
- [D5.1] Deliverable D5.1, "Description of sustainability experimentation framework", [Online]. Available [here](#).
- [Goyal2025] S. Goyal, M. Kapuruhamy, B. Golebiowski, K. Wanuga, E. Visotsky, J. Korhonen, A. Lappalainen, and A. Talukdar, "Understanding the 3GPP Standardization Aspects of Network-Controlled Repeaters," IEEE Communications Magazine, vol. 9, pp. 36–43, 2025. Available [here](#).
- [Hlawatsch2011] F. Hlawatsch and G. Matz, Wireless Communications Over Rapidly Time-Varying Channels, 1st ed. USA: Academic Press, Inc., 2011. Available [here](#).
- [HRM+2025] A. Husein, K. Rasilainen, J.-P. Mäkelä, A. Pärssinen, and M. E. Leinonen, "6G OTA Measurements at Sub-THz Band Using a Compact Robotic System," in EUCAP2025, 30.3.-4.4.2025, Stockholm, Sweden, p. 1-4. Available [here](#).
- [Huang2020] C. Huang, R. Mo, and C. Yuen, "Reconfigurable Intelligent Surface Assisted Multiuser MISO Systems Exploiting Deep Reinforcement Learning," IEEE Journal on Selected Areas in Communications, vol. 38, no. 8, pp. 1839–1850, 2020. Available [here](#).
- [iPerf2025] iPerf, [Online]. Available: <https://iperf.fr/> [Accessed 14 June 2025].
- [LinuxPTP2025] The Linux PTP Project, [Online]. Available: <https://linuxptp.sourceforge.net/> [Accessed 15 June 2025].
- [Lorca2023] J. L. Hernando and A. G. Armada, "Piecewise Equalization of Zero Padding OFDM and FM-OFDM in Doubly-Dispersive Channels," 2023 IEEE Globecom Workshops (GC Wkshps), Kuala Lumpur, Malaysia, 2023, pp. 811-817. Available [here](#).
- [Lorca+2025] J. Lorca Hernando, L. Méndez Monsanto and A. García Armada, "Channel Estimation and Equalization of Zero-Padded Waveforms in Doubly-Dispersive Channels," accepted for publication in IEEE Transactions on Communications, 2025. Available [here](#).
- [O-RAN] Open RAN functional splits, explained, <https://www.5gtechnologyworld.com/open-ran-functional-splits-explained/> (accessed 26 June 2025)
- [Paymard+2023] P. Paymard, A. Amiri, T. E. Kolding and K. I. Pedersen, "Optimizing Mixed Capacity of Extended Reality and Mobile Broadband Services in 5G-Advanced Networks," IEEE Access, vol. 11, pp. 113324-11338, 2023. Available [here](#).
- [Qosium2025] Kaitotek Qosium, [Online]. Available: <https://www.kaitotek.com/qosium> [Accessed 15 June 2025].

[Saglam2023] B. Saglam, D. Gurgunoglu, and S. S. Kozat, "Deep Reinforcement Learning Based Joint Downlink Beamforming and RIS Configuration in RIS-aided MU-MISO Systems Under Hardware Impairments and Imperfect CSI," Proc. of IEEE, 2023. Available [here](#).

[Sin2024] S. P. Singh, Radio receiver front-end ICs at the sub-THz/THz frequency range in silicon technology, Doctoral thesis, University of Oulu, Finland, p. 123, 2024. Available [here](#).

[SNM+2023] S. P. Singh, M. J. Nokandi, M. H. Montaseri, T. Rahkonen, M. E. Leinonen and A. Pärssinen, "A 300-320 GHz Sliding-IF I/Q Receiver Front-End in 130 nm SiGe Technology," 2023 IEEE Radio Frequency Integrated Circuits Symposium (RFIC), San Diego, CA, USA, 2023, pp. 37-40. Available [here](#).

[Sousa2024] D. Sousa, T. Maciel, F. Lima, M. Åström, V. Monteiro, B. Makki, D. Moreira, F. Carvalho, R. Paiva, and L. Bao, "Network-Controlled Repeater – An Introduction," arXiv preprint, arXiv:2403.09601, 2024. Available [here](#).

[SPR+2025] S. P. Singh, P. Phukphan, T. Rahkonen, A. Pärssinen and M. E. Leinonen, "Power calibration methods for frequency extenders aided modulated measurements at sub-THz/THz," in Gemic2025, Mar. 17–19, 2025, p. 1 -4. Available [here](#).

[Sun2019] Y. Sun, M. Peng, Y. Zhou, Y. Huang, and S. Mao, "Application of Machine Learning in Wireless Networks: Key Techniques and Open Issues," IEEE Communications Surveys & Tutorials, vol. 21, no. 4, pp. 3072–3108, 2019. Available [here](#).

[Vulchi2024] H. Vulchi, S. A. Busari, J. Rodriguez, and A. Garcia Armada, "Phase shift configuration for RIS-based 6G networks using DRL technique," Proceedings of the 2024 IEEE International Workshop on Computer-Aided Modeling and Design of Communication Links and Networks (CAMAD), Athens, Greece, 2024. Available [here](#).

[Xu2022] S. Xu, H. Xu, S. Xiong, H. Zhang, and J. Xin, "A Survey on Network Controlled Repeater Technology," 2022 IEEE 8th International Conference on Computer and Communications (ICCC), pp. 1097–1101, 2022. Available [here](#).

[Yigit2020] Z. Yigit, E. Basar, and I. Altunbas, "Low Complexity Adaptation for Reconfigurable Intelligent Surface-Based MIMO Systems," IEEE Communications Letters, 2020. Available [here](#).

8 ANNEXES

8.1 DESCRIPTION OF XAPPS AND CONTROL SIGNALLING FOR THE ENERGY-AWARE END-TO-END RESOURCE MANAGEMENT O-RAN ENABLER

Table 11 lists the implemented xApps and describes their functionality (see Figure 29 in section 3.2.2.2).

Table 11: xApps and their functionality

xApp name	Interactions	Functionality
pwr_mon	Receives energy monitoring data from energy measurement database. Sends control data to rate_lim xApp.	The power monitoring xApp (pwr_mon) is responsible for fetching the relevant power consumption data (pwr_cons) from the energy measurement database and assessing if the current consumption is below or above the set allowed maximum consumption value. This threshold value can be defined based on the measured relationship between the O-RAN infrastructure power consumption and network load. In case the consumption is above the target, it sends a power threshold indication message (pwr_thold) to the data rate limiter xApp (rate_lim) asking it to lower the load in the cell by limiting the maximum data rate of the active UEs or the PRB quota of the active network slices. When the consumption is again below the target, it sends a new pwr_thold message to the rate_lim xApp allowing for more load.
kpm_mon	Receives KPI monitoring data from CU and DU. Sends control data to rate_lim xApp.	The Key Performance Measurements (KPM) monitoring xApp (kpm_mon) is responsible for fetching the relevant cell load data from the E2 Node components. For the assessment of the cell load, it can use metrics such as the aggregated GTP tunnel throughputs (gtp_tputs) from the DU and number of active UEs (ue_count) from the CU. Based on these parameters, it assesses if the overall load of the cell is below or above the set allowed maximum value. This threshold value can be defined based on the measured relationship between the O-RAN infrastructure power consumption and network load. In case the load is above the target, it sends a load threshold indication message (load_thold) to the data rate limiter (rate_lim) xApp asking it to lower the load in the cell by limiting the maximum data rate of the active UEs or the PRB quota of the active network slices. When the load is again below the target, it sends a new load_thold message to the rate_lim xApp allowing for more load.
rate_lim	Receives power and load threshold values from pwr_mon and	Based on the power (pwr_thold) and load (load_thold) threshold indication messages received from the power (pwr_mon) and KPM monitoring (kpm_mon) xApps, respectively, the data rate limiter xApp (rate_lim) decides of the maximum data rate of the active UEs or the PRB quota of the active network slices needs to

	kpm_mon xApps. Sends control data to UE.	be decreased or if it can be increased. In case both pwr_thold and load_thold messages request for lowering the load, it sends a data rate modification control message (rate_mod) to the E2 Node/UE(s) lowering the network load. This process can be based on the modification of the Session-AMBR value of the UEs using the 3GPP specified Network-requested PDU session modification procedure defined in section 6.3.2 of [3GPPTS24501], or alternatively, on the modification of the slice-specific PRB quota. When both pwr_thold and load_thold allow for higher cell load again, a new rate_mod message is sent to the UE(s) to allow higher network load.
--	---	--

Table 12 lists the messages exchanged between the xApps, O-RAN architecture components, and energy measurement framework (see Figure 29 in section 3.2.2.2).

Table 12: Monitoring and control messages

Message name	Source and destination	Content
ru_meas	From RU energy meter to energy measurement database	Raw 1/s measurement data for the RU hardware: timestamp, current [A], frequency [Hz], power factor [PF], voltage [V], power [W], and cumulative power consumption [kWH].
du_meas	From DU energy meter to energy measurement database	Raw 1/s measurement data for the DU server hardware: timestamp, current [A], frequency [Hz], power factor [PF], voltage [V], power [W], and cumulative power consumption [kWH].
cu_meas	From CU energy meter to energy measurement database	Raw 1/s measurement data for the CU/RIC/5GC server hardware: timestamp, current [A], frequency [Hz], power factor [PF], voltage [V], power [W], and cumulative power consumption [kWH].
pwr_cons	From energy measurement database to pwr_mon xApp	Aggregated measurement data from the energy measurement database: timestamp, power [W], and cumulative power consumption [kWH].
gtp_thputs	From DU to kpm_mon xApp	Aggregated 1/s measurement data from the E2 Node: timestamp, DL GTP throughput (bps), and UL GTP throughput (bps).
ue_count	From CU to kpm_mon xApp	Aggregated measurement data for the E2 Node: timestamp, number of active UEs.
pwr_thold	From pwr_mon xApp to rate_lim xApp	Request to modify network load in the O-RAN infrastructure: increase/decrease load.

load_thold	From kpm_mon xApp to rate_lim xApp	Request to modify network load in the O-RAN infrastructure: increase/decrease load.
rate_mod	From rate_lim xApp to UE	Request to modify UE specific Session-AMBR value or slice specific PRB quota.

8.2 CHANNEL ESTIMATION AND EQUALIZATION IN DOUBLY-DISPERSIVE CHANNELS

This section describes the details of the piecewise channel estimation and equalization technique for doubly-dispersive channel conditions introduced in Section 4.2.2.2.2.

With reference to Figure 45, WB-RS can span up to the system bandwidth with a period equal or smaller than T_d . NB-RS spans MN_L subcarriers that can be split in two (or more) parts around DC, at frequency locations determined by the parameters k_1 and k_2 . Splitting in two parts ensures that the corresponding time-domain signal undergoes wider variations from the higher-order subcarriers involved, which reduces the estimation uncertainty from the higher value of the second-order derivatives in the Cramer-Rao Bound. N_L is the number of piecewise intervals that the symbol is divided into to perform incremental CSI estimation. N_L can be dimensioned from the rms error incurred by the piecewise approximation, given by the root of the universal bound for ICI power in a duration $\frac{T_s}{N_L} + \frac{1}{B_c}$ [Lorca2023],

$$\varepsilon_{LTI} \leq \frac{\pi f_D}{\sqrt{3}} \left(\frac{T_s}{N_L} + \frac{1}{B_c} \right),$$

where f_D is the doppler spread and B_c is the channel's coherence bandwidth.

The periodicity of WB-RS (T_d) and NB-RS (T_c) are related by the expression $T_d \lesssim \frac{f_c}{f_s} T_c$, where T_c is the channel coherence time, f_c is the carrier frequency, and f_s is the sampling rate. When T_c is equal or lower than the symbol duration, NB-RS must be allocated in every symbol, as in Figure 45. The guard subcarriers are aimed to isolate NB-RS from any inter-carrier interference (ICI) from an adjacent data/control subcarrier allocated in the same symbol, with

$$N_g \gtrsim \left\lceil \frac{f_D + W_{PN}}{SCS} \right\rceil$$

where SCS is the subcarrier spacing and W_{PN} is the phase noise spectral width.

A piecewise channel estimation technique can be applied to infer a series of linear time-invariant (LTI) channel responses whose combination represents the incremental CSI that best approaches the time-varying channel response, as in Figure 80 [Lorca2023]:




Figure 80: Illustration of the decomposition of the time-variant channel taps $b_{i,j}$ into a set of piecewise time-invariant channels whose channel responses are estimated by sampling the time-varying taps at each of the defined N_L symbol intervals.

Let us assume a received signal vector $\mathbf{R} \in \mathbb{C}^{N_T \times 1}$, where $N_T = N + N_{CP/ZP}$, with N as the DFT size and $N_{CP/ZP}$ as the length of the cyclic prefix (CP) or zero-padding (ZP) field. The corresponding symbol contains NB-RS and other data or control subcarriers that are multiplexed in frequency.

The symbol contents are first filtered to discard the subcarriers not carrying NB-RS. Doing so depends on whether the time-domain symbol structure contains a CP (e.g., as in CP-OFDM or DFT-s-OFDM) or a ZP (e.g., as in ZP-OFDM or FM-OFDM).

- If the symbol contains a CP, filtering out the subcarriers not devoted to NB-RS can be done in the frequency domain by nulling them after an N -point DFT:

$$\bar{\mathbf{R}} = \mathbf{F}_N^H \mathbf{\Pi} \mathbf{F}_N \mathbf{R} \quad (4)$$

where $\mathbf{F}_N$ is the N -size DFT matrix, $\mathbf{F}_N^H$ is its Hermitian transpose, and the diagonal matrix $\mathbf{\Pi} = \{\pi_{jk}\} \in \mathbb{C}^{N \times N}$ filters out the subcarriers not containing NB-RS as follows:

$$\pi_{jk} = \begin{cases} 1, & (j, k) \in \mathcal{K} \\ 0, & \text{otherwise} \end{cases}$$

with $\mathcal{K}$ as the set of subcarrier indices allocated to NB-RS.

- If the symbol contains a ZP, an overlap-and-add (OLA) operation is first applied to recover the circularity of the symbol by adding the contents of the ZP field to the beginning of the symbol,

$$\mathbf{R}_{OLA} = \mathbf{R}_u + [\mathbf{R}_l \mathbf{0}_{1 \times (N - N_{ZP})}]^T \in \mathbb{C}^{N \times 1}, \quad (5)$$

where $\mathbf{R}_u$ and $\mathbf{R}_l$ are column vectors containing the upper N rows and the lower N_{ZP} rows of $\mathbf{R}$, respectively. OLA yields an N -point circular vector whose control/data subcarriers can be nulled prior to channel estimation as in the CP case, to yield a vector

$$\bar{\mathbf{R}} = \mathbf{F}_N^H \mathbf{\Pi} \mathbf{F}_N \mathbf{R}_{OLA}. \quad (6)$$

If the receiver knows the multipath delays and the length of the Channel Impulse Response L (in samples), a matrix equation can be written

$$\bar{\mathbf{R}} = \mathbf{S}_{OLA} \mathbf{B} \quad (7)$$

where $\mathbf{B} \in \mathbb{C}^{N_L \times 1}$ is the unknown L -point vector containing the complex tap coefficients at the positions of the a-priori known delays and 0 elsewhere, and the following matrices may be defined as follows with $\Delta = N/N_L$:

$$\mathbf{S}_{OLA} = \begin{bmatrix} \mathbf{V}_0 & \mathbf{0} & \cdots & \mathbf{U}_0 \\ \mathbf{U}_1 & \mathbf{V}_1 & \ddots & \vdots \\ \vdots & \vdots & \ddots & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{U}_{N_L-1} & \mathbf{V}_{N_L-1} \end{bmatrix} \quad \mathbf{V}_i = \begin{bmatrix} s_{(i\Delta)} & 0 & \cdots & 0 \\ s_{(i\Delta+1)} & s_{(i\Delta)} & \vdots & \vdots \\ \vdots & \vdots & \ddots & 0 \\ s_{(i\Delta+L-1)} & \cdots & s_{(i\Delta+1)} & s_{(i\Delta)} \\ s_{(i\Delta+L)} & \ddots & s_{(i\Delta+2)} & s_{(i\Delta+1)} \\ \vdots & \vdots & \vdots & \vdots \\ s_{((i+1)\Delta-1)} & \cdots & s_{((i+1)\Delta-L+1)} & s_{((i+1)\Delta-L)} \end{bmatrix} \quad \mathbf{U}_i = \begin{bmatrix} 0 & s_{(i\Delta-1) \bmod N} & \cdots & s_{(i\Delta-L+1) \bmod N} \\ 0 & 0 & \ddots & \vdots \\ \vdots & \vdots & \ddots & s_{(i\Delta-1) \bmod N} \\ 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & 0 & 0 \end{bmatrix}$$

The terms $s_{(n)}$ in $\mathbf{V}_i, \mathbf{U}_i \in \mathbb{C}^{\Delta \times L}$ are the elements of the N -point time domain signal vector $\mathbf{C}_N = [s_{(0)} \dots s_{(N-1)}]^T$ constructed from a complex sequence of symbols $\mathbf{C}$ in the frequency domain that comprise NB-RS, by means of an interpolation

$$\mathbf{C}_N = \mathbf{F}_N^H \text{vec}(\mathbf{C}, \mathbf{0}_{(N-MN_L) \times 1}). \quad (8)$$

The solution to (7) is the complex vector $\tilde{\mathbf{B}} = \{\tilde{b}_{i,l}\}$ that satisfies

$$\tilde{\mathbf{B}} = (\mathbf{S})^+ \bar{\mathbf{R}}, \quad (9)$$

where $(\mathbf{S})^+$ is the Moore-Penrose pseudo-inverse of $\mathbf{S}$. Calculation of the pseudo-inverse can be done, e.g., by solving the optimization problem

$$\tilde{b}_{i,l} = \arg \min \left\{ \|\bar{\mathbf{R}} - \mathbf{S} \tilde{\mathbf{B}}\|_F^2 \right\}. \quad (10)$$

This can be challenging if the problem is ill-conditioned, as happens when the ratio MN_L/N is very small. An alternative is to use a Tikhonov regularization method given by

$$\tilde{b}_{i,l} = \arg \min \left\{ \|\bar{\mathbf{R}} - \mathbf{S} \tilde{\mathbf{B}}\|_F^2 + \lambda^2 \|\mathbf{B} - \mathbf{B}_0\|_F^2 \right\}, \quad (11)$$

where $\mathbf{B}_0$ is a pre-defined constant and λ is the ridge parameter. This solution is equivalent to the one in (10) with an improved condition number depending on the value of λ .

The Channel Frequency Responses of the N_L time-invariant channels that approximate the LTV channel are obtained over a number of points that depends on whether CP or ZP is present, as follows:

- CP: $\tilde{\mathbf{H}}_i = \mathbf{F}_N \text{vec}(\tilde{\mathbf{B}}_i, \mathbf{0}_{(N-L) \times 1})$.
- ZP: $\tilde{\mathbf{H}}_i = \mathbf{F}_{N_T} \text{vec}(\tilde{\mathbf{B}}_i, \mathbf{0}_{(N_T-L) \times 1})$.

The normalized mean square error from channel estimation can be given by

$$\text{NMSE} = \frac{\|\bar{\mathbf{R}} - \mathbf{S}\tilde{\mathbf{B}}\|_F^2}{\|\bar{\mathbf{R}}\|_F^2}. \quad (12)$$

NMSE includes the combined impact of thermal noise, imperfect channel estimation, and the error incurred by the piecewise approximation of an LTV channel into N_L time-invariant channels. NMSE can be used in the design of NB-RS to determine an optimum value of M .

With the above channel estimates, a piecewise equalization technique can readily undo the effect of the channels, as, e.g., described in [\[Hlawatsch2011\]](#).

The described approach has been accepted for publication in [\[Lorca+2025\]](#).

8.3 UNIFIED WEB PORTAL USER GUIDE

This annex is intended for experimenters who wish to execute their experiments and view the results through the designated North and South portals. It provides a detailed overview of the Unified web portal, where experimenters can log in, create new trials, and then proceed to the corresponding North and South portals to launch experiments and monitor results based on each specific trial.

System Overview

The system is a web-based application that, along with its database and graphical front-end interface, allows experimenters to create and manage their trials.

The application consists of several components, including:

- Tabs for “Trials” and “About Unified”
- User authentication features such as Login, Logout, Registration, and Password Recovery
- User management: View, edit, and delete user accounts
- Trial management: Create, edit, and delete trials
- Redirection of experimenters to the appropriate North or South portal based on the selected trial.

User Functionality

User Registration:

- The experimenter is required to provide their personal information, including full name, phone number, email address, username, and a password.
- The password must meet complexity requirements to ensure account security.
- All credentials and experimenter information are securely stored in the database using encryption.

Password Recovery:

- If an experimenter forgets their password, they can enter their registered email address to receive a verification code.
- This code allows them to reset their password securely.

User Deletion:

- Experimenters have the ability to delete their accounts.
- Once an account is deleted, all associated data is permanently removed from the database and the application.

Post-Registration:

After successful registration, the experimenter is redirected to the unified portal dashboard.

Creating a Trial

Upon logging into the unified portal dashboard, the experimenter is presented with a list of existing trials and also has the option to create a new one.

- All fields are mandatory; if any field is left empty, the trial will not be created.
- After successfully creating a trial, the experimenter is redirected to the trial list page, where they can view and manage all their created trials.
- If the trial's end time is set to a time earlier than the current system time, the trial will appear as inactive in the trial list page.

To create a trial, the experimenter must fill in all required fields, including:

- Trial Name
- Start Time
- End Time
- Facility
- Description

Trial List

On this page, the experimenter can view a list of all created trials, including both active and inactive ones.

If a trial has not yet expired, the experimenter can click on the desired facility to be redirected to either the North or South portal. Additionally, this page allows experimenters to manage their trials, including options to edit, view, and delete each trial.

Additional Interface Components

In the top-left corner of the application, the experimenter can find the main navigation tabs. One of these tabs is the "About Unified" section, which provides general information and guidance on how to use the system. In the top-right corner, there is an experimenter menu. By clicking on it, experimenter can view and edit their personal information.

Note: The email field is not editable.

From this same experimenter menu, experimenters also have the option to delete their accounts. Next to the profile icon is a settings gear icon that allows experimenters to customize the interface's appearance.

Options include:

- Switching between light and dark themes
- Adjusting page layout sizes
- Aligning the interface left-to-right or right-to-left

Finally, at the bottom left of the interface, experimenters can log out of the system or access external links related to the 6G-XR project.

Data Deletion Notes

- When an experimenter's account is deleted, all associated data will be permanently removed. This includes:
 - The experimenter's personal information
 - All trials created by the experimenter
 - All experiments related to those trials in both the North and South portals
- When a trial is deleted, all experiments linked to that trial in both the North and South portals will also be permanently deleted.

Contact Info / Support

If you encounter any issues or conflicts while using the system, please contact Mahdi Salmani, the person responsible for the Unified and North portals:
mohammad.salmani@oulu.fi

8.4 NORTH WEB PORTAL USER GUIDE

System Overview

In the North Portal, the experimenter can create their experiments based on selected slice types such as eMBB and URLLC. After initiating the experiment, experimenters can monitor the resulting data through the integrated dashboard and perform evaluations based on the observed outcomes within the North and South portals. The primary purpose of this portal is to enable the creation of experiments and the visualization of related performance results in a clear and structured manner.

User Functionality in the North Portal

- Once an experimenter registers in the Unified Portal, their account is automatically created in the North Portal as well.
- In the North portal, the experimenter is permitted to view only their profile.
- All user-related actions, such as registration, account editing, deletion, and password recovery, are handled exclusively through the Unified Portal. Experimenters must return to the Unified Portal to perform any of these operations.

Create an experiment:

Upon accessing the dashboard of the North Portal, the experimenter can view a list of existing experiments or choose to create a new one. If the trial period has expired, the user will no longer be able to create new experiments. Additionally, any experiment must be created within the valid duration of the associated trial.

To create an experiment, the experimenter must provide the following information:

- Experiment Name (more than one character)
- Start Time
- End Time
- Application name
- Slice Type: either URLLC or eMBB

Specific fields, such as Trial Name, Facility, and Node, are predefined and fixed. The Trial Name is automatically imported from the Unified Portal, establishing a clear link between each experiment and its corresponding trial.

After successfully creating an experiment, the experimenter is redirected to the experiment list page, where all created experiments can be viewed and managed.

Experiment List Page

Each experiment entry provides options to view, edit, delete, and send the experiment to the NNA (North Node Adapter).

Important:

After creating an experiment, the experimenter must first send it to NNA. Only then can the results of the experiment be viewed.

Experiment Results Visualization

A details page opens by clicking the "View" (eye icon). This includes:

- Experiment information
- A button that redirects the user to the Grafana dashboard for monitoring experiment results

Access Restrictions:

- The experimenter cannot enter the Grafana dashboard if the experiment has already ended.

- If the current system time is later than the experiment's start time, and the experiment has not yet been sent to NNA, the experimenter will be blocked from accessing the dashboard.
- In such cases, the experimenter must create a new experiment with a valid future start time.

Selecting KPIs

- The experimenter can select 4 KPIs from the provided list.
- After submitting the selected KPIs, the system will wait for the experiment to start.
- If the experiment has already started, results will be shown immediately. Otherwise, the experimenter must wait until the scheduled start time.

Grafana Dashboard

- The experimenter is redirected to the Grafana dashboard to view real-time results.
- Grafana provides options such as time range adjustments and panel settings.
- Experimenters can download results anytime or wait until the experiment ends to access more complete data.

Terminating an Experiment

- The experimenter can terminate an experiment at any time, whether it is currently running or not.
- The experimenter can also reselect KPIs and resubmit them to see updated results.

Concurrent Experiment Rule

- Only one experiment can be active at a time.
- If there is already an active experiment, a warning message will appear at the top of the experiment list.
- The experimenter must either wait for it to finish or delete the current active experiment (if they created it) to start a new one.
- After the experiment is completed, it will remain in the system for only 10 minutes before being deleted. Users must review the information and download their results within this 10-minute window.

Summary Flow:

1. Create an experiment with a valid time.
2. Send the experiment to NNA.
3. Go to the Grafana dashboard
4. Select and submit 4 KPIs.
5. Wait for the start time (or view results if already started).
6. Monitor on the Grafana dashboard.
7. Download results or terminate the experiment as needed.

Additional Interface Components

In the top-left corner of the application, the experimenter can find the main navigation tabs. One of these tabs is the "User Guide" section, which provides general information and guidance on how to use the system. In the top-right corner, there is an experimenter menu. By clicking on it, experimenters can view their personal information. (Editing is not available in the North Portal)

Next to the profile icon, there is a settings gear icon that allows experimenters to customize the appearance of the interface.

Options include:

- Switching between light and dark themes
- Adjusting page layout sizes
- Aligning the interface left-to-right or right-to-left

Finally, at the bottom left of the interface, experimenters can log out of the system or access external links related to the 6G-XR project.

Contact Info / Support

If you encounter any issues or conflicts while using the system, please contact your open call mentor.

8.5 SOUTH WEB PORTAL USER GUIDE

System Overview

In the South Portal, the experimenter can create his own experiment based on selected applications and slice types. After initiating the experiment, he can also monitor the resulting data through the provided dashboard. The primary purpose of this portal is to enable the creation of experiments and the visualization of related performance results in a clear and structured manner.

As preliminary step before using the South Portal, the experimenter must contact their corresponding mentor and provide the applications to be onboarded on the edge infrastructure, clarifying the node where they should be executed (5Tonic, i2CAT or both). After the onboard, Capgemini provides the experimenter with an “AppProviderID” to be employed for registering his account on the South Portal.

User Functionality in the South Portal

- The experimenter registers his account, selecting username, password and inserting the provided “AppProviderID”.
- Once the account is created, the experimenter logs in through the selected username and password (“AppProviderID” is no more required).

Dashboard Page

Upon accessing the South Portal, the experimenter has access to the dashboard page, where the information on how to access the Grafana dashboard is provided.

NST Page

In this page, the experimenter can view a list of existing experiments or choose to create a new one. Each experiment entry provides options to view, delete or run the experiment. The execution is performed through the SNA (South Node Adapter), while metrics are collected through Qosium and shown in a Grafana dashboard.

Creating an Experiment

The experimenter can create a new experiment, providing the following information:

- Experiment Name
- Applications
 - Application Name
 - Slice -> This represents where it will be deployed
 - Resource Consumption -> This is the policy the MEC Orchestrator will use. This should be given beforehand by the mentor/administrator.
 - Flavour -> The dedicated resources option will only be inserted if specified by the MEC Orchestrator admin. It represents different configuration templates for the image regarding assigned resources.
- GPSI & Slices (if they apply)
 - GPSI -> The identifier of pre-registered UEs for testing
 - Slice -> The network slice the UE will subscribe to

Certain fields—such as Trial Name, Facility, and Node—are predefined and fixed. The Trial Name is automatically imported from the Unified Portal, establishing a clear link between each experiment and its corresponding trial.

After successfully creating an experiment, the user is redirected to the experiment list page, where all created experiments can be viewed and managed.

Concurrent Experiment Rule

- Only one experiment can be active at a time.
- If there is already an active experiment, a warning message will appear at the top of the experiment list.
- The experimenter must either wait for it to finish or terminate the current active experiment (if they created it) to start a new one.

Grafana Dashboard

- Once deployed, the experimenter will have a menu option in Actions that links to the Grafana URL. The Grafana installed in South Node monitors activity of the MEC Orchestrator as a whole. Using the Concurrent Experiment Rule specified above, time filters can be used to locate the launched experiment and monitor it.

Terminating an Experiment

- The experimenter can terminate an experiment at any time once it is launched.

Summary Flow:

1. Enter South Node Web Portal with user and password, register if necessary.
2. NST Page / create new NST button, it will fill with the trial data from UWP, specify applications and slices previously specified by the corresponding mentor/administrator
3. Launch the saved NST
4. Monitor on the Grafana dashboard.
5. Terminate the experiment as needed, not later than the end of trial date, so the lock is released for the next experimenters.