Horizon Europe MSCA Project

(Grant Agreement Number 101086219)

Intelligent and Proactive Optimisation for Service-centric Wireless Networks

IPOSEE

D3.1

Service-centric RAN Optimisation Applications for O-RAN RIC

UPPSALA UNIVERSITET

Authors	Xinyu Gao, Jiming Chen
Authors' Affiliation	Ranplan Wireless Network Design Ltd.
Editors:	Xinyu Gao, Jiming Chen, Xiaoli Chu, Yiouli Kritikou, Panagiotis Demestichas, Kan Lin, Weijie Qi, Di Yuan, Lingyou Zhou, and Xin Dong,
Status-Version:	Final
Project Number:	101086219
Project Title:	Intelligent and Proactive Optimisation for Service-centric Wireless Networks
Project Acronym:	IPOSEE
Work Package Number	3

Table of Contents

List of Figures

Executive summary

Wireless networks have become a critical part of modern infrastructure, supporting industries such as transport, logistics, utilities, and manufacturing. As service types continue to grow and diversify, the demands placed on wireless networks have become increasingly complex and difficult to manage with traditional approaches. In particular, understanding and forecasting complex, high-dimensional traffic patterns — whether for individual services or aggregated groups — remains a significant challenge for future network planning and optimisation.

The Open Radio Access Network (O-RAN) architecture is designed to introduce intelligence, flexibility, and openness into the management and optimisation of radio access networks (RANs)[1]. RAN Intelligent Controller (RIC) is a central component of the O-RAN architecture. There are two types of RIC defined by the O-RAN Alliance: Non-Real-Time RIC and Near-Real-Time RIC. Service-centric RAN optimisation applications for the O-RAN RIC are designed to enhance the performance of specific services (e.g., voice, video streaming, URLLC, eMBB, IoT) by dynamically adapting radio network behaviour to meet diverse service-level requirements[2]. These applications typically run as xApps in the Near-Real-Time RIC for timescales of 10ms–1s, or as rApps in the Non-Real-Time RIC for longer-term optimisation and policy control.

In this deliverable, we build upon the initial work completed in D1.1, where methods for spatial-temporal traffic forecasting and mobility prediction were developed using Artificial Intelligence (AI) techniques. Our goal in D3.1 is to extend these methods into practical evaluation: summarizing the main findings from D1.1, applying them to selected simulation scenarios, and laying the groundwork for network optimisation based on forecasted traffic patterns.

We focus not only on identifying representative traffic profiles but also on understanding how these patterns can influence typical RAN deployment strategies. All simulations and preliminary optimisation studies are conducted using the Ranplan Professional software platform, ensuring a consistent and realistic evaluation environment.

This report marks a step forward in the IPOSEE project's aim to develop intelligent, proactive network optimisation strategies, by moving from theoretical development towards service-centric RAN optimisation applications for O-RAN RIC and their practical validation.

[1] M. Polese, *et al.*, "Understanding O-RAN: Architecture, interfaces, algorithms, security, and research challenges," *IEEE Communications Surveys & Tutorials*, vol. 25, iss. 2, pp. 1376-1411, 2023.
[2] P. Sroka, *et al.*, "Policy-based traffic steering and load balancing in O-RAN-based vehicle-to-network communications," *IEEE Transactions on Vehicular Technology*, vol. 73, iss. 7, pp. 9356-9369, 2024.

1. Introduction

This deliverable, D3.1, is produced under Work Package 3 (WP3) of the IPOSEE project. It builds directly on the work completed in earlier stages, particularly the datasets, models, and methodologies developed in Deliverable D1.1, which focused on high-resolution wireless traffic and mobility forecasting with uncertainty quantification. Service-centric RAN optimisation in the O-RAN architecture ensures that the diverse and stringent requirements of next-generation applications (5G/6G) are met efficiently[3]. By combining real-time control via xApps and policy/AI-driven oversight via rApps, O-RAN RIC will enable programmable, intelligent, and service-aware RANs.

The purpose of D3.1 is to extend these developments towards practical evaluation. It summarizes the key methods and findings reported in D1.1, applies them to selected simulation cases designed to reflect realistic deployment scenarios, and explores preliminary approaches to network optimisation based on forecasted traffic patterns. Through this approach, D3.1 continues the technical progression initiated in WP1 and supports the broader goals of WP3 in enabling proactive, service-centric wireless network operations.

This document is structured into three main parts:

- A summary of the key methodologies, datasets, and results developed in D1.1, with a focus on their applicability to simulation and optimisation tasks.
- A presentation and analysis of selected use cases, demonstrating how the forecasting and modelling techniques can be applied to evaluate network performance in practice.
- A preliminary discussion of optimisation approaches based on forecasted traffic patterns, with further extensions to be finalized as the project progresses.

By linking the initial research activities with application-driven evaluation and optimisation, D3.1 plays a key role in advancing the IPOSEE project towards achieving its overall objectives.

1.1 Purpose of this document

The purpose of this deliverable, D3.1, is to consolidate and extend the outcomes achieved during the initial phases of the IPOSEE project, particularly the work carried out under Work Package 1 (WP1) and documented in Deliverable D1.1. While D1.1 focused on developing data-driven methods for forecasting wireless traffic and mobility patterns, D3.1 moves the work forward by applying these methods in practical simulation environments and preparing for optimisation activities under Work Package 3 (WP3).

[3] A. Lacava, *et al*., "Programmable and customized intelligence for traffic steering in 5G networks using open RAN architectures," *IEEE Transactions on Mobile Computing*, vol. 23, iss. 4, pp. 2882-2897, 2024.

Specifically, this document serves to:

- Summarize and highlight the key methodologies, datasets, and results established in D1.1, ensuring a consistent technical foundation for subsequent simulation and evaluation tasks.
- Demonstrate the practical applicability of the developed forecasting and modelling techniques through the analysis of selected use cases, providing insights into network behaviour under different service and traffic conditions.
- Set the groundwork for necessary further developments, particularly the optimisation strategies and decision-support tools that will be finalized in later stages of WP3.

By achieving these objectives, D3.1 ensures continuity between early-stage research and the application-driven phases of the project, helping to bridge the gap between theoretical modelling and practical network optimisation.

1.2 Document Structure

This document is organized into the following main sections:

- **Section 2 – Summary of D1.1 Methodologies and Findings:** This section provides a concise overview of the key methodologies, datasets, and results developed during D1.1. It focuses on their relevance to simulation and optimisation tasks, covering dataset preparation, environment identification, spatial-temporal traffic forecasting, and integration into network simulation tools.
- **Section 3 – Use Case Analysis:** This section presents two selected simulation use cases. For each case, the network setup, simulation configuration, and performance evaluation are discussed, along with key insights into the practical application of the forecasting and modelling techniques.
- **Section 4 – Preliminary Optimisation Approach:** This section introduces initial optimisation strategies informed by the forecasting and simulation results. It discusses how predictive traffic information can guide network deployment and optimisation decisions, providing a basis for further development of optimisation tools within WP3.
- **Section 5 – Further Work:** This section outlines areas for future work, including potential extensions of the optimisation strategies and additional simulation studies that will be completed as the project progresses.

2. Summary of D1.1 Methodologies and Findings

This section provides a concise yet structured summary of the key methodologies, datasets, and findings presented in Deliverable D1.1, with an emphasis on their relevance to the simulation and optimisation tasks addressed in the later stages of the IPOSEE project.

The primary objective of D1.1 was to develop data-driven techniques for high-resolution forecasting of wireless traffic and user mobility, with a particular focus on incorporating uncertainty quantification. This emphasis on uncertainty is critical to ensuring that the predictions are not only accurate in expectation, but also robust to outlier behaviours, context shifts, and sparse measurement conditions—all of which are common in real-world network environments.

The outputs of D1.1 form the analytical foundation for the predictive simulation platform developed in WP3. The models, training pipelines, and data representations introduced in D1.1 are directly integrated into simulation workflows, enabling forecast-informed network planning, dynamic spectrum management, and resource optimisation strategies to be explored in Deliverables D3.1 and beyond.

2.1 Methodologies Developed in D1.1

Deliverable D1.1 focused on the development of data-driven methodologies designed to predict wireless traffic and mobility patterns with high spatial and temporal granularity. The overarching goal was to support proactive network optimisation by generating not only accurate traffic forecasts but also quantified uncertainty estimates, thereby enabling more informed and risk-aware decision-making processes in simulation and planning environments.

The technical work in D1.1 was structured around four core components, which collectively support the pipeline from raw data ingestion to predictive inference and uncertainty quantification. These components are described in the following subsections.

2.1.1 Data Collection and Preparation

The first stage of methodology development focused on building a representative and reliable dataset ecosystem to support model training and evaluation. Multiple heterogeneous datasets were curated to reflect a wide range of urban environments, user mobility patterns, and network usage behaviours. These included:

- **Historical wireless traffic records**, collected from diverse public available;
- **Mobility traces**, capturing aggregate and fine-grained user movement patterns;
- **Environmental context data**, such as building footprints, heights, land use classifications, and transportation infrastructure.

To ensure consistency and generalizability across datasets, several standardized preprocessing techniques were applied:

- **Normalization** was used to align feature scales across time and location dimensions, ensuring comparability between heterogeneous traffic indicators;
- **Data augmentation** techniques—such as random perturbations, temporal resampling, and geographic jittering—were employed to increase dataset

diversity and reduce model overfitting to specific city layouts or population densities.

This preprocessing pipeline resulted in a robust, high-variance dataset foundation suitable for training complex forecasting models, and ensured that the trained models could generalize across different network topologies and traffic regimes.

2.1.2 Environment Identification and Classification

Understanding the relationship between physical environments and traffic patterns was critical for enabling accurate and context-aware forecasting. To achieve this, high-resolution satellite imagery was analysed using semantic segmentation techniques. These methods allowed the model to classify urban regions into functional categories such as:

- **Commercial zones** (e.g., business districts, shopping centres),
- **Residential neighbourhoods**,
- **Transportation hubs** (e.g., train stations, airports, major intersections).

This environmental classification layer was integrated into the prediction pipeline to inform the learning model of the semantic function of different urban zones, which is highly correlated with both traffic volume and temporal activity patterns. **Figure 1** shows an example of environment classification results, where different types of areas are visually distinguished to reflect their traffic generation characteristics.

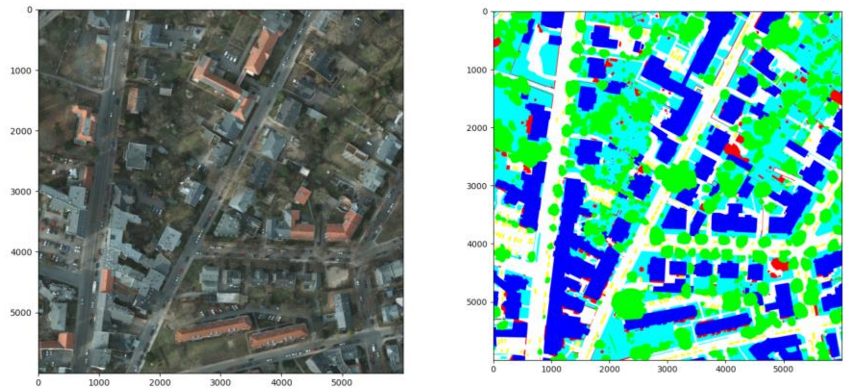

Figure 1. Environment type identification.

By explicitly embedding environmental context into the forecasting process, the model could better learn and generalize spatial correlations between urban structure and observed network behaviour—laying the groundwork for spatially adaptive optimisation in later work packages.

2.1.3 Spatial-Temporal Traffic Forecasting

Building on the environment classification, models were developed to forecast traffic variations across both space and time. The modelling approach addressed spatial dependencies—how traffic in one area influences surrounding regions—and temporal dynamics such as daily and weekly usage cycles. The spatial dimension of traffic patterns was captured through clustering techniques applied to traffic distribution data, as illustrated in **Figure 2**. This clustering helped identify areas with similar traffic characteristics, forming the basis for localized forecasting.

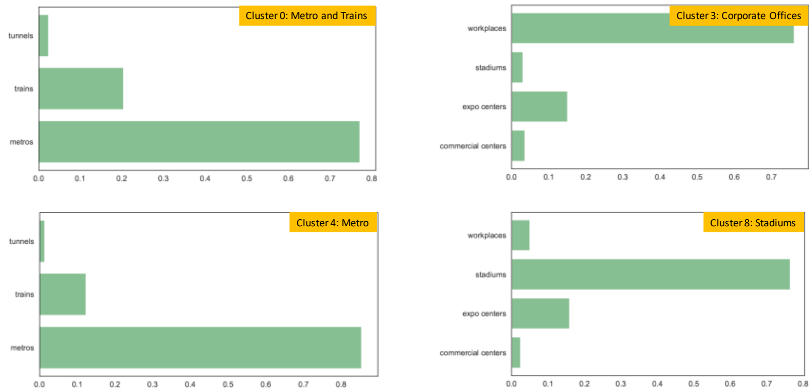

Figure 2. Spatial traffic classification.

Additionally, traffic clustering was refined using unsupervised learning methods to further group traffic behaviours into meaningful categories, as shown in **Figure 3**.

Figure 3. Traffic clustering.

Temporal analysis was also a critical component. **Figure 4** presents typical daily traffic variation patterns observed in different environmental types, providing insights into peak and off-peak usage cycles.

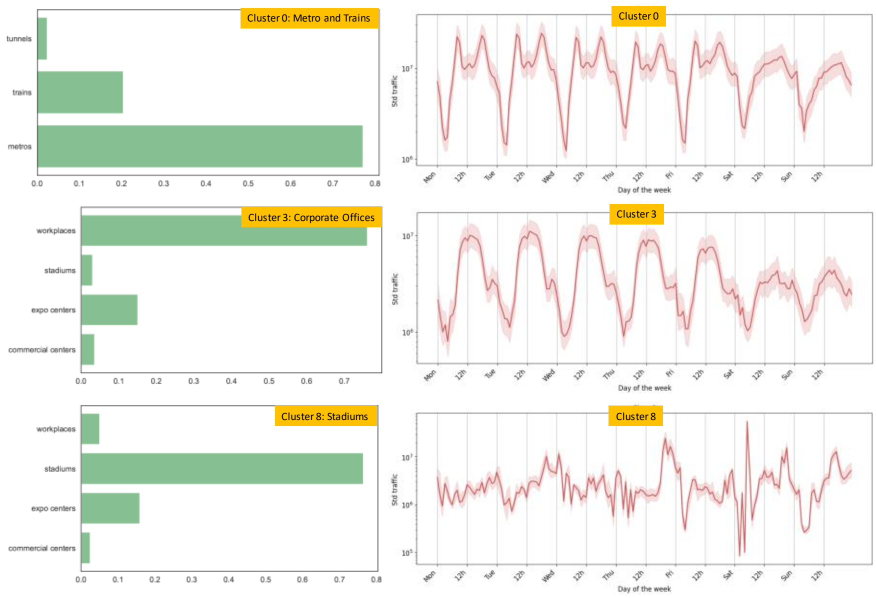

Figure 4. Temporal traffic analysis.

The forecasting models leveraged both historical traffic patterns and environmental context to anticipate future traffic states. **Figure 5** shows an example of predicted traffic distributions across urban areas.

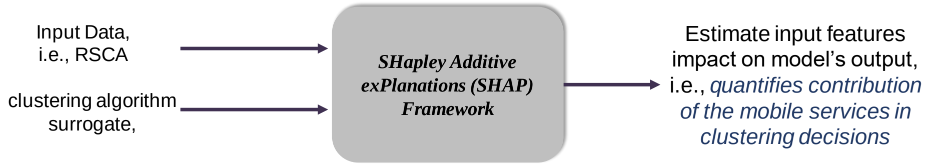

Figure 5. Traffic predictions.

Finally, **Figure 6** provides a consolidated view of spatial-temporal traffic dynamics, highlighting how traffic intensity evolves throughout the day across different zones.

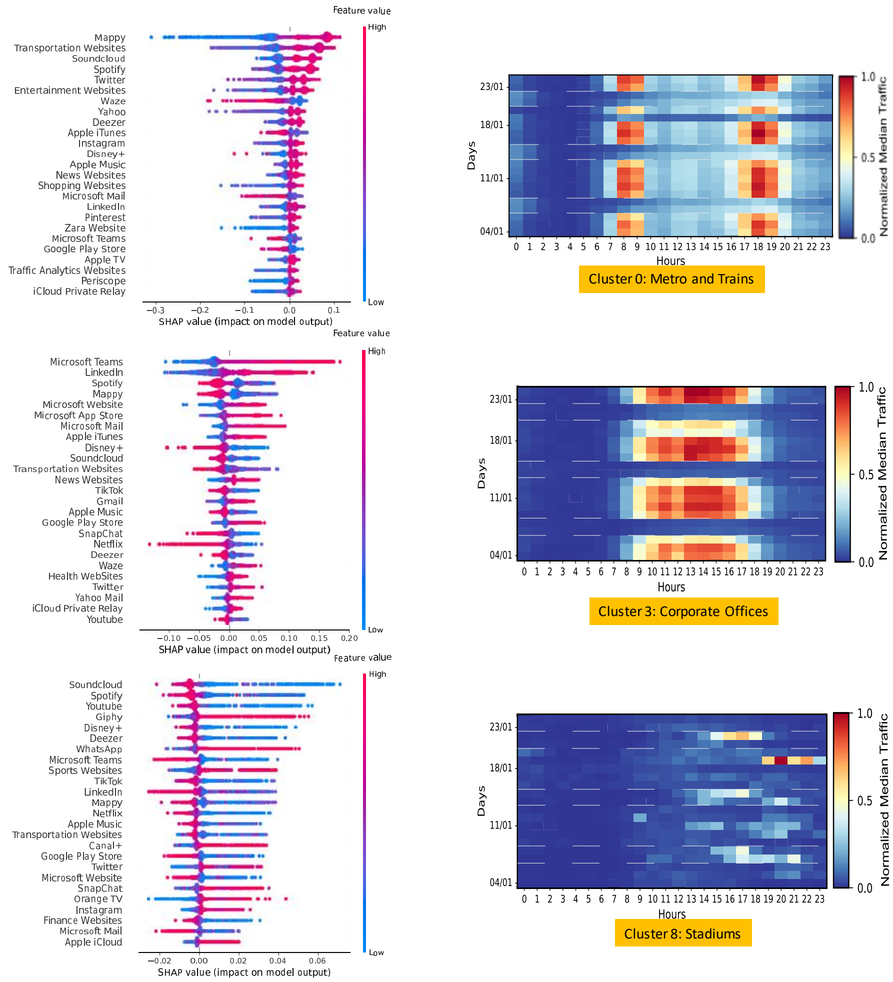

Figure 6. Spatial-temporal traffic patterns.

2.1.4 Uncertainty Quantification

Given the inherent variability and unpredictability of wireless traffic—particularly in dynamic urban environments—it was essential to complement deterministic forecasting with explicit uncertainty quantification. This ensures that predictions not only indicate likely traffic levels, but also convey the degree of confidence associated with each estimate. To achieve this, D1.1 adopted a combination of Bayesian inference techniques and Monte Carlo dropout methods, both of which are widely used in probabilistic deep learning. These approaches allowed the forecasting models to generate confidence intervals around predicted values, rather than point estimates alone. The resulting uncertainty bounds provided valuable insights into:

- The expected range of traffic values at a given time and location;
- Areas with high prediction volatility, which may require conservative planning strategies;
- Time windows with low variance, offering opportunities for aggressive spectrum reuse or optimisation.

By integrating uncertainty into the forecasting pipeline, the model outputs became more actionable for downstream use in simulation, resource provisioning, and risk-aware network design. This aligns with the broader IPOSEE objective of enabling proactive and adaptive network planning, particularly under conditions of incomplete or evolving information.

In summary, uncertainty-aware predictions offer a more robust foundation for decision-making, empowering network designers to better anticipate and respond to fluctuations in traffic demand across heterogeneous deployment scenarios.

2.2 Key Findings in D1.1

The work carried out in D1.1 produced several key insights that have directly shaped the design and priorities of subsequent simulation and optimisation activities in WP3. These findings confirmed the importance of spatial context, temporal regularity, and predictive robustness in building practical tools for intelligent network management. The most influential findings are outlined below.

2.2.1 Environment Classification Brings Practical Benefits for Traffic Modeling

One of the first and most impactful conclusions from D1.1 was that semantic environment classification significantly enhances the accuracy and realism of traffic prediction models. By applying semantic segmentation to satellite imagery and other geospatial data, urban areas were automatically divided into functional categories such as:

- Commercial districts,
- Residential neighbourhoods,
- Transportation hubs.

Each of these categories is associated with distinct traffic behaviours, usage patterns, and peak-load signatures. Rather than treating the urban landscape as a homogeneous prediction surface, incorporating these distinctions enabled the forecasting models to adopt a context-aware formulation. For example, traffic patterns in commercial zones exhibited sharp, cyclical peaks aligned with business hours, while residential areas showed a more gradual increase in activity during the evening hours. This environment-aware modelling approach allowed spatial forecasting methods to move beyond generic assumptions, yielding stronger alignment with real-world usage and enabling more reliable resource planning.

2.2.2 Spatial-Temporal Forecasting Models Capture Real Usage Patterns

Another major finding from D1.1 was that spatial-temporal forecasting models, when properly trained and calibrated, are capable of reliably capturing realistic traffic dynamics. These models account for:

- **Spatial dependencies** (e.g., how neighbouring areas influence local demand),
- **Temporal evolution** (e.g., daily peaks, weekly cycles, seasonal changes).

Testing results showed that the forecasting models successfully anticipated traffic surges linked to recurring behavioural patterns—such as morning and evening commutes, lunchtime congestion, and weekend retail activity. In addition, they could identify off-peak periods, which are particularly useful for guiding energy-saving or low-priority resource allocation strategies.

While some prediction noise remained in highly dynamic or mixed-use zones, the models proved robust enough to support proactive and adaptive optimisation, a central objective of the IPOSEE framework.

These findings reinforce the notion that forecasting is not just a descriptive tool, but a key enabler of real-time decision-making in modern wireless networks.

2.2.3 Uncertainty Estimation Adds Meaningful Insights

Incorporating uncertainty estimation into the forecasting process proved to be far more than a theoretical enhancement—it delivered practical planning value. By producing not only point forecasts but also confidence intervals, the models provided insight into the reliability of their own predictions. This enabled network planners to make more informed decisions about where forecasts could be confidently trusted, and where more flexible, resilient designs might be required.

This was particularly important in environments such as transport hubs, where traffic volumes are subject to sharp, irregular fluctuations driven by external factors like weather, flight schedules, or public events. In these cases, the ability to anticipate not just the average demand but also its expected variability made a direct impact on the risk profile of deployment strategies.

By quantifying uncertainty, the forecasting framework transitioned from a static predictive tool into a strategic planning instrument, enabling better alignment between capacity provisioning and demand volatility. This capability supports a more risk-aware, adaptive approach to network design—central to the IPOSEE project's vision of proactive service-centric RAN optimisation.

2.2.4 Environment-Specific Traffic Behaviour Supports Tailored Deployment

Another critical insight from D1.1 was that environment-specific traffic behaviours are not merely academic observations—they have direct operational implications for wireless network deployment. Traffic patterns were found to vary significantly across different types of environments:

- **Commercial zones** exhibited sharp daytime peaks and quiet evenings;
- **Residential areas** followed slower build-up and evening-dominant usage curves;

- **Transportation nodes** showed high burstiness, often disconnected from time-of-day norms.

Recognizing these distinct patterns early in the planning process allowed for far more efficient network deployment. Rather than applying a uniform planning strategy across the entire urban area, network resources could be allocated adaptively, based on the actual usage rhythms and service needs of each environment type.

This observation supports the broader goal of service-centric RAN optimisation applications, which moves away from rigid, one-size-fits-all rules toward intelligent, environment-aware optimisation in O-RAN. It also laid the groundwork for the functional zoning approach later applied in WP3 simulation studies, where deployment and configuration decisions are directly linked to the spatial semantics of traffic behaviour.

3. Evaluation of Forecast-Driven Simulation Environment in an Urban Scenario

Building upon the forecasting methodologies and findings established in Deliverable D1.1, this section presents the design, calibration, and evaluation of a simulation environment tailored to support forecast-informed wireless network planning. The evaluation is conducted within a realistic urban scenario, designed to emulate the structural and mobility complexities that are typical in high-density European cities.

The overarching goal is to validate whether the simulation platform developed in WP3 can meaningfully incorporate traffic and mobility predictions from WP1, and whether it can serve as a reliable digital testbed for forecast-guided optimisation strategies. By grounding simulation behaviour in both empirical measurements and environmental structure, we aim to reduce the performance gap between theoretical planning and real-world deployment outcomes.

3.1 Bath Urban Scenario: Calibration and Validation of Simulation Environment

As a first step toward building a trustworthy simulation framework, we conducted a scenario-based validation exercise in the city centre of Bath, United Kingdom. This location was chosen for its representative urban characteristics, including narrow historical streets, dense mid-rise buildings, and non-uniform elevation profiles—all of which contribute to complex wireless propagation conditions.

The purpose of this case study is twofold:

1. To verify the **accuracy of the propagation models** embedded in the WP3 simulation platform against real-world signal measurements;
2. To establish a **calibrated, high-fidelity baseline environment** that can later be used to test forecast-driven resource planning and network optimisation.

This section details the measurement campaign, calibration steps, and statistical validation of the radio model used. The resulting **propagation engine**, once calibrated, enables simulations that faithfully reflect the structural and environmental features of realistic urban settings, thereby supporting meaningful evaluation of the predictive strategies developed in the IPOSEE project

3.1.1 Motivation and Role

One of the key objectives of the IPOSEE project is to establish an intelligent, prediction-aware framework for wireless network optimisation. At the heart of this framework lies a suite of high-resolution traffic forecasting and uncertainty quantification models developed under Work Package 1 (WP1), as detailed in Deliverable D1.1. These models encompass environment-aware traffic density estimation, spatial-temporal mobility clustering, and probabilistic demand modelling.

To ensure that these predictive models can be effectively leveraged in practical network planning and decision-making, Work Package 3 (WP3) is tasked with building a simulation and optimisation platform capable of responding to forecast inputs. However, the reliability of any such simulation platform critically depends on the accuracy of its underlying radio propagation models. If the predicted path loss is inaccurate, then—even when driven by precise traffic forecasts—the simulation outputs may be misleading or untrustworthy.

To address this challenge, we carried out a calibration and validation exercise in a realistic urban environment. The objectives of this exercise were to:

- Assess the accuracy of simulation models against real-world measurement data;
- Calibrate ray-tracing parameters to better reflect actual urban propagation characteristics;
- Ensure the simulation platform accurately represents the structural and environmental complexity of dense urban areas;
- Establish a reliable foundation for the forecast-driven optimisation strategies to be developed in subsequent IPOSEE deliverables.

3.1.2 Measurement Campaign and Scenario Setup

The selected validation area is located in the city centre of Bath, United Kingdom, which presents a complex and diverse urban propagation environment. This area was chosen due to its rich combination of physical characteristics, including:

- Narrow historical streets and open plazas;
- Dense mid-rise building clusters constructed with a variety of materials;
- Elevation variations caused by the city's naturally hilly terrain.

To replicate realistic small cell deployment conditions, 16 pico base stations were installed on street-level lampposts at a height of 4 meters. Each base station was equipped with omnidirectional antennas configured for 4×4 MIMO operation. The system operated in the 3.7 GHz band (n77) with a channel bandwidth of 100 MHz. The spatial configuration of the base stations, along with the coverage area for simulation and measurement, is illustrated in **Figure 7**.

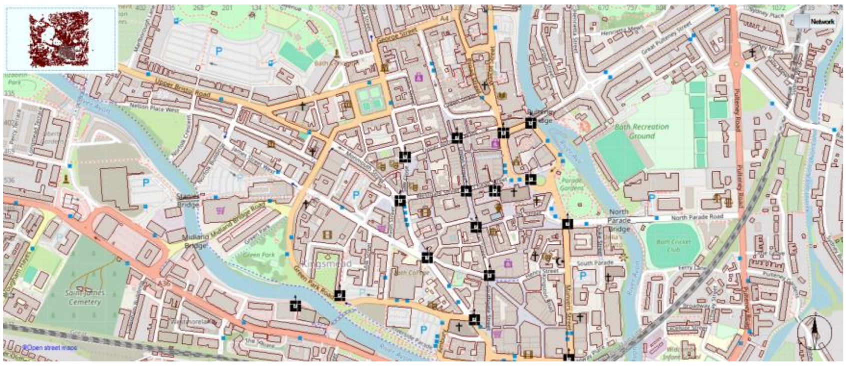

(a) Network deployment

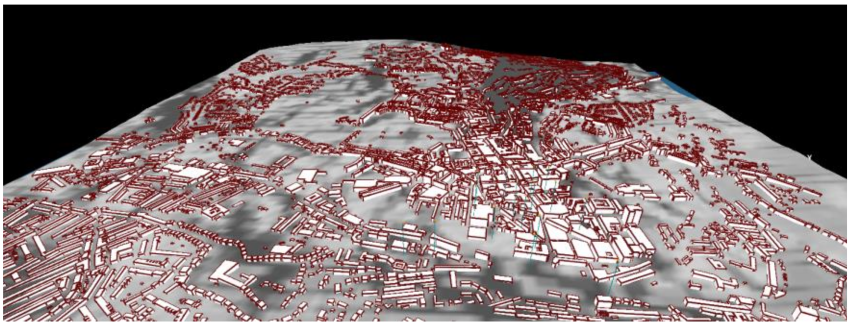

(b) 3D view

Figure 7. Bath urban measurement scenario with 16 picocell deployments using omnidirectional antennas on 4-meter poles.

Receiver locations were carefully distributed throughout the measurement area, ensuring broad coverage of both line-of-sight (LOS) and non-line-of-sight (NLOS) regions. These receiver points were selected to capture the full range of propagation conditions, from open, unobstructed streets to deep urban canyons and heavily shadowed corners. Signal strength data were collected through on-site measurements and subsequently compared to predicted values generated by the default (uncalibrated) ray-tracing propagation model within the WP3 simulation platform. This comparative analysis served as the basis for the model calibration process, which is described in the following subsection.

3.1.3 Performance Across Deployment Scenarios

Using the default ray-tracing configuration, the simulation platform produced path loss predictions which were then systematically compared against the measured signal strength data. As illustrated in **Figure 8**, the discrepancies between simulated and empirical values were significant—especially in NLOS environments, such as alleyways, narrow side streets, and building-shielded corners. The mean absolute error exceeded 9 dB, and the standard deviation was unacceptably high, indicating that the default model failed to accurately capture the propagation dynamics of the Bath urban scenario.

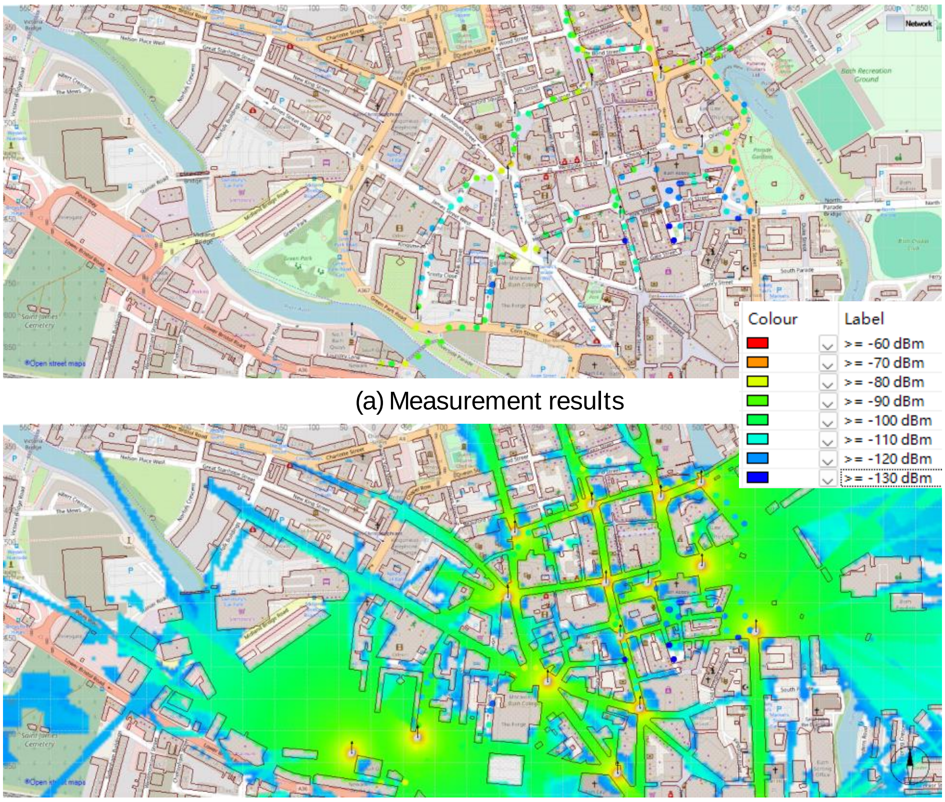

(a) Measurement results

(b) Simulation results

Figure 8. Comparison between simulated and measured path loss using default parameters, showing high variance and location-dependent prediction errors.

This outcome underscored the limitations of generic propagation assumptions when applied to structurally diverse environments. As emphasized in Deliverable D1.1, accurate wireless simulation requires the integration of environment-specific knowledge, particularly when historical or non-standard construction materials are present. In this case, the default wall penetration losses, diffraction losses, and surface reflection coefficients were not suitable for the stone façades, arched passages, and mixed-material building envelopes that characterize the Bath city centre. These findings clearly justified the need for a dedicated calibration process, which aimed to align simulation predictions with actual urban signal behaviour. The calibration methodology and resulting improvements are described in the next subsection.

3.1.4 Propagation Model Calibration Procedure

To improve prediction accuracy, a targeted calibration of the ray-tracing engine was carried out. The following key adjustments were implemented:

- Increased wall penetration loss to reflect the presence of thick stone structures commonly found in Bath's historical buildings;
- Reduced reflection loss in close-spaced building facades to better model multipath persistence in narrow streets;
- Adjusted diffraction thresholds to capture the effects of street-canyon propagation more accurately;
- Refined receiver sensitivity settings to better reflect practical detection limits in typical user equipment.

These modifications were informed by WP1's environmental classification outputs, which were generated through supervised learning models trained on datasets of building footprints, elevations, and material attributes. This approach enabled the simulation platform to contextualize propagation behaviour according to locally identified environmental classes, resulting in a more precise calibration per zone type. The updated simulation output is presented in **Figure 9**, where the predicted coverage map exhibits significantly improved alignment with real-world measurements across diverse urban conditions.

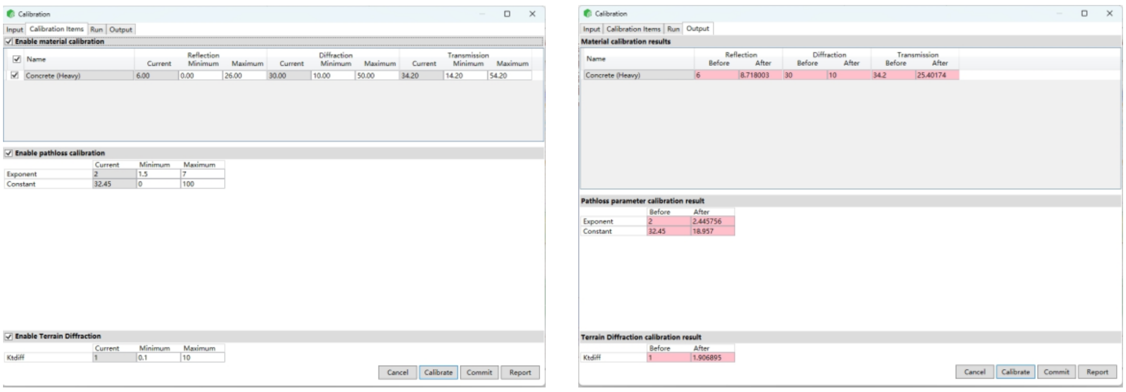

(a) Path loss parameters before calibration

(b) Path loss parameters after calibration

Figure 9. Calibrated propagation engine output showing improved alignment with real-world measurements in urban Bath.

This calibration phase substantially reduced both over- and under-prediction of path loss, particularly in NLOS conditions and transition areas between open squares and dense building corridors. The result is a more reliable and consistent baseline for simulation-based evaluation and forecast-driven network planning activities in WP3.

3.1.5 Statistical Validation and Integration into WP3 Platform

A statistical evaluation of simulation accuracy before and after calibration is presented in **Figure 10**. The results reveal a significant improvement in model performance: the average prediction error was reduced from +9.44 dB to −1.85 dB, while the standard deviation dropped from 64.10 dB to 13.05 dB. These reductions reflect a substantial gain in both accuracy and stability, directly supporting the quality assurance of the simulation engine employed in WP3.

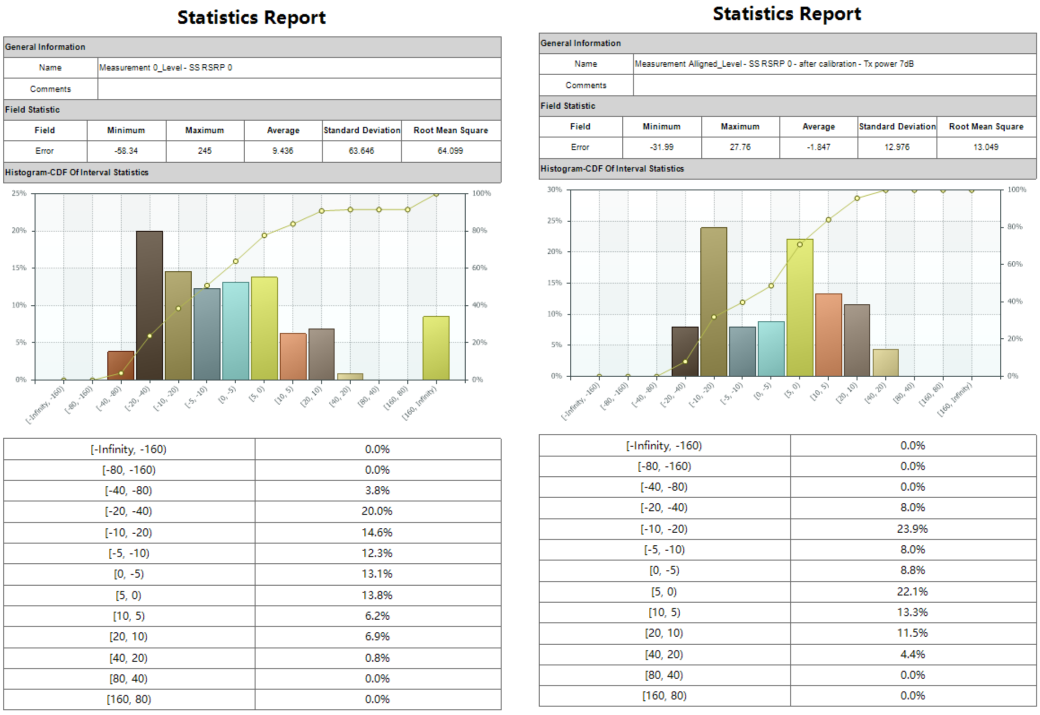

Statistics Report (a)

General Information	
Name	Measurement 0_Level - SS RSRP 0
Comments	

Field	Minimum	Maximum	Average	Standard Deviation	Root Mean Square
Error	-58.34	245	9.436	63.646	64.099

Histogram-CDF Of Interval Statistics

[-Infinity, -160)	0.0%
[-80, -160)	0.0%
[-40, -80)	3.8%
[-20, -40)	20.0%
[-10, -20)	14.6%
[-5, -10)	12.3%
[0, -5)	13.1%
[5, 0)	13.8%
[10, 5)	6.2%
[20, 10)	6.9%
[40, 20)	0.8%
[80, 40)	0.0%
[160, 80)	0.0%

(a) Error statistic before calibration

Statistics Report (b)

General Information	
Name	Measurement Alligned_Level - SS RSRP 0 - after calibration - Tx power 7dB
Comments	

Field	Minimum	Maximum	Average	Standard Deviation	Root Mean Square
Error	-31.99	27.76	-1.847	12.976	13.049

Histogram-CDF Of Interval Statistics

[-Infinity, -160)	0.0%
[-80, -160)	0.0%
[-40, -80)	0.0%
[-20, -40)	8.0%
[-10, -20)	23.9%
[-5, -10)	8.0%
[0, -5)	8.8%
[5, 0)	22.1%
[10, 5)	13.3%
[20, 10)	11.5%
[40, 20)	4.4%
[80, 40)	0.0%
[160, 80)	0.0%

(b) Error statistic after calibration

Figure 10. Statistical comparison of prediction error before and after propagation model calibration.

The validated model resulting from this calibration process will now serve as the radio propagation foundation for upcoming simulation scenarios under WP3. These scenarios will explore forecast-informed resource allocation, spectrum reuse strategies, and deployment optimisation frameworks. While the current deliverable does not implement full optimisation loops, it provides a critical enabling step toward that goal by ensuring the simulation platform can meaningfully respond to high-resolution input data. Moreover, this process exemplifies how WP1 outputs—particularly the environmental classification models and probabilistic traffic forecasts—can be systematically used to inform and refine the simulation environment, thereby enhancing the fidelity and credibility of subsequent prediction-driven evaluations.

3.2 Airport Scenario: UWB and Mobile Network Coexistence at Heathrow

While Section 3.1 focused on validating the simulation platform in a conventional urban deployment setting, this section demonstrates how the forecasting capabilities developed in WP1 can be extended beyond network deployment to address more complex challenges—specifically, spectrum coexistence between heterogeneous wireless systems.

With increasing pressure on mid- and high-band spectrum resources, technologies such as 5G New Radio (NR) and Ultra-Wideband (UWB) radar systems are beginning to share frequency ranges, especially around the upper 6 GHz band. This creates practical risks of mutual interference, particularly in high-density environments like airports, where multiple critical systems operate in parallel.

Set at London Heathrow Airport, this scenario illustrates how traffic prediction, environmental awareness, and calibrated propagation models can be combined to:

- Proactively identify interference-prone regions before deployment;
- Define dynamic protection zones for sensitive UWB systems based on forecasted user flows;
- And inform scheduling and coordination strategies to enable safe and efficient spectrum sharing.

By integrating this scenario into the IPOSEE framework, we demonstrate that the project's forecasting capabilities are not limited to traditional mobile network planning, but also serve as a powerful tool for predictive interference management and multi-system coordination. This further highlights the practical value of the IPOSEE methodology in supporting future wireless ecosystems, where cross-system coexistence will be increasingly essential.

3.2.1 Scenario Description and Setup

This scenario evaluates spectrum coexistence between UWB radar systems and 5G mobile networks at London Heathrow Airport, one of the most complex and high-density wireless environments in the United Kingdom. In this context, UWB systems represent law enforcement, infrastructure monitoring, or airport security sensors, operating in the upper 6 GHz band, while the mobile network comprises a 5G NR deployment delivering continuous coverage across passenger terminals, aprons, and access routes.

The simulation models incorporate the following environmental and deployment elements:

- A flat airport terrain with large open surfaces, metallic structures, and strong structural obstructions;
- Three UWB radar stations positioned near critical buildings or high-traffic transit zones;
- Co-channel coexistence with 5G NR base stations operating in partially overlapping frequency bands;
- Signal interactions governed by spectrum sensing-based coexistence rules and strict interference protection thresholds.

The simulation layout is illustrated in **Figure 11**, which highlights the placement of UWB radar systems, 5G base stations, and key structural elements within the airport environment.

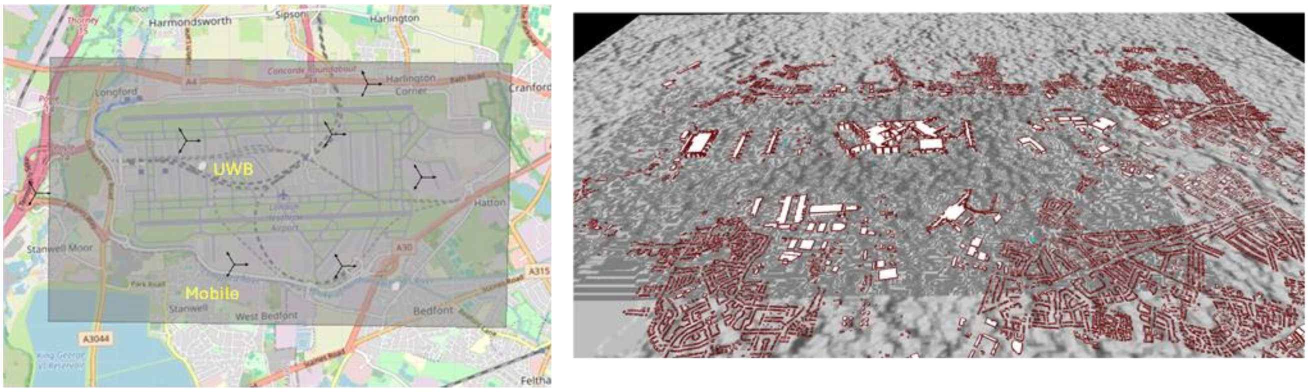

(a) UWB case

(b) 3D view

Figure 11. Simulation layout of Heathrow Airport with UWB radar stations and mobile infrastructure for coexistence analysis.

In this scenario, the UWB systems are modelled according to standard compliance specifications: they operate over a 500 MHz bandwidth with very low transmit power spectral density (PSD). The maximum PSD is limited to −41.3 dBm/MHz, aligning with international regulatory frameworks for short-range device operation (e.g., ITU-R and ETSI standards). These conditions create a challenging coexistence environment, particularly because UWB signals—despite their low power—occupy wide spectral bands that may overlap with mobile uplink or downlink subcarriers. The aim of the simulation is to analyse how well spectrum sensing mechanisms and interference management thresholds protect both system types under realistic deployment constraints.

3.2.2 Forecasting-Based Protection Modelling

To evaluate realistic interference risks and support proactive spectrum planning, we employed a forecast-driven modelling approach similar to that described in Section 3.1. This approach leverages both environmental and user behaviour predictions to determine context-aware protection zones around sensitive UWB devices. Specifically, we applied the following steps:

- Environmental classification was used to distinguish between obstructed and open-space areas within the airport environment, based on structural layouts and material datasets;
- Forecasted traffic density maps were integrated to simulate typical user movement flows across terminals, boarding areas, and service corridors;
- Reciprocal path loss estimation—following the methodology outlined in Deliverable D1.1—was used to calculate dynamic protection radii around UWB stations, accounting for varying propagation and occupancy conditions.

Using the known EIRP of 5G NR base stations (58.7 dBm) and the interference threshold for UWB reception (−116.6 dBm), we derived a required protection path loss threshold of 175.3 dB. This means that, under typical propagation assumptions, mobile transmitters must be kept at least 150 meters away from UWB receivers to prevent harmful interference in most cases. The spatial pattern of these constraints is visualized in **Figure 12**, which presents a UWB path loss heatmap overlaid on the simulated airport layout. The results clearly show how architectural features—such as terminal walls, hangars, and vehicle corridors—introduce strong variations in propagation, effectively reshaping the protection zones from simple circles into complex, structure-aware contours.

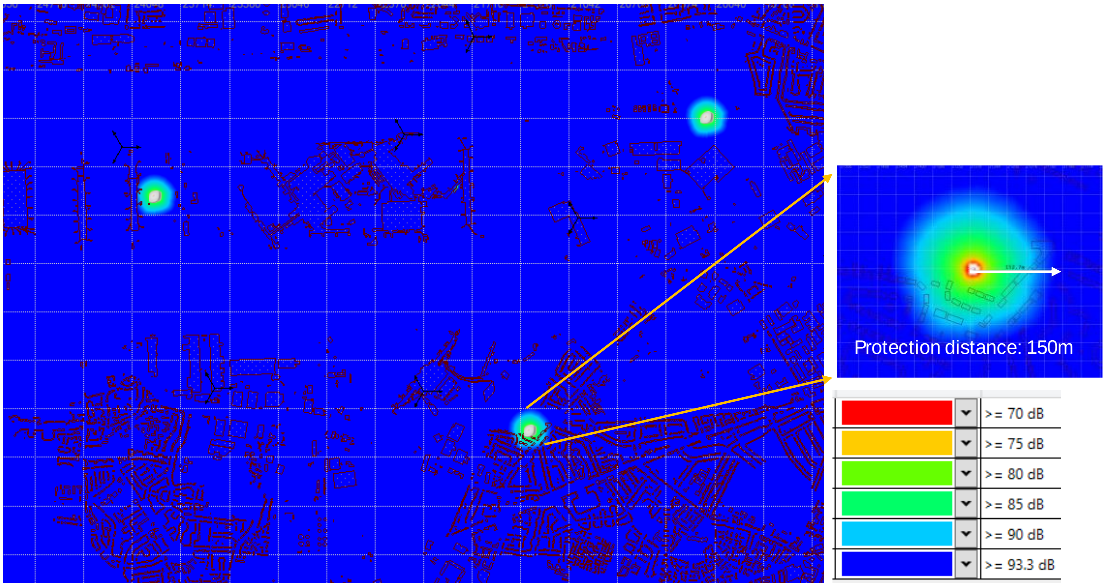

Figure 12. UWB path loss heatmap at Heathrow, illustrating protection zones based on interference threshold modelling.

This forecasting-based approach enables the proactive definition of exclusion zones, spectrum reuse windows, and dynamic guard-band enforcement strategies. By aligning protection modeling with real-world environmental structure and traffic flows, the system can move beyond static constraints and toward context-aware, predictive coexistence management.

3.2.3 Network Simulation and Results

The simulation assessed coexistence performance under dynamic spectrum sharing rules, where UWB systems and 5G mobile networks operate in overlapping frequency bands but rely on RSSI-based sensing mechanisms to avoid harmful interference. The goal was to validate whether real-time environmental awareness and traffic sensitivity can enable safe and efficient spectrum reuse in a high-density setting such as an international airport.

Key findings from the simulation include:

- Elevated interference risk in areas near UWB radar stations, particularly when user hotspots dynamically shift into overlap zones;
- Significant uplink performance degradation in 5G NR cells when UWB beam paths align with mobile base station coverage sectors, due to receiver-side desensitization;
- Strong spatial variation in coexistence behaviour, depending on local terminal geometry, construction materials, and signal reflection characteristics.

To quantify the impact of shared-band coordination, **Figure 13** compares system-level throughput under two configurations:

i) Single-band operation using only the conventional 3.5 GHz mid-band spectrum;

ii) Dual-band operation that adds the upper 6 GHz band under sensing-based coordination constraints.

In the second case, the simulation shows a throughput increase of up to 970.5 Mbps per cell, confirming that opportunistic access to high-frequency spectrum—when governed by real-time sensing—can provide substantial capacity gains without violating UWB protection thresholds.

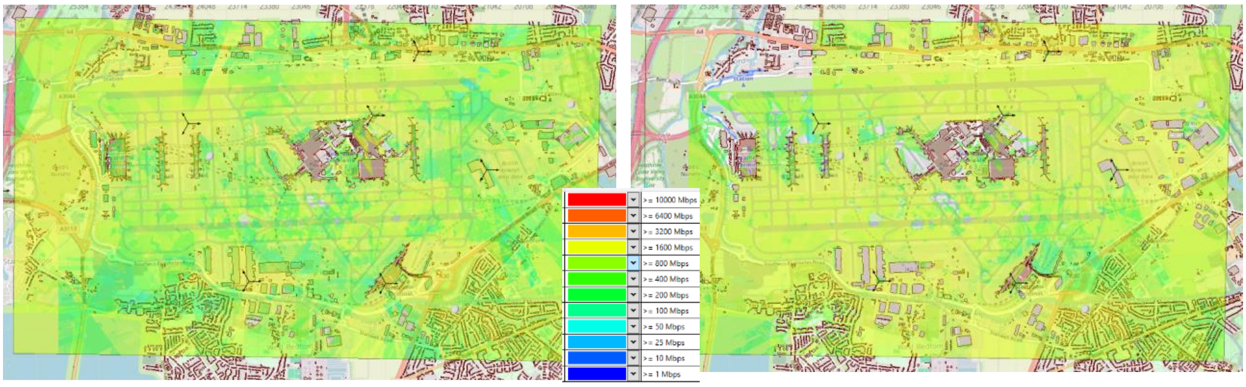

(a) throughput of 3.5 GHz   (b) throughput of upper 6 GHz

Figure 13. Throughput improvement with spectrum sharing between 5G and UWB at Heathrow Airport. Sensing-based exclusion maintains UWB protection while increasing network capacity.

These results also provide strong empirical support for the proposition in Deliverable D1.1: That traffic-aware and interference-aware scheduling mechanisms enable fine-grained coexistence, even in dense, mission-critical environments such as airports. The combination of forecast-driven traffic modelling and sensing-guided coordination offers a promising path forward for intelligent spectrum reuse.

3.2.4 Insights and Observations

The use case demonstrates that spectrum sharing between UWB and 5G mobile networks is technically feasible, provided that coexistence conditions are intelligently managed and context-aware adjustments are applied. The application of forecasting tools—particularly those capable of modelling environmental structure and dynamic user mobility—plays a pivotal role in defining adaptive protection zones that safeguard UWB system performance while enabling mobile networks to reclaim underutilized spectrum for high-capacity service delivery.

One of the core challenges in this scenario stems from broad-area, multipoint interference, driven by fluctuating user density and mobility within a dense and operationally complex environment like an airport. Such dynamics require continuous spatial and temporal adaptation, which can only be effectively achieved through forecast-informed modelling and simulation.

The results highlight several key insights:

- The importance of mobility-aware interference mapping, particularly in environments with high user churn such as airports;
- The need for proactive, site-wide spectrum reuse policies, guided by accurate traffic prediction and calibrated simulation;

- The practical viability of prediction-guided spectrum sensing as a scalable and low-overhead approach to maintaining performance in multi-service environments.

This evaluation strengthens the case for applying forecast-driven intelligence to future wireless system coordination and sets the stage for the optimisation strategies that will be developed in the subsequent sections of this deliverable.

4. Forecast-Guided Optimisation Framework

With a validated simulation environment in place and high-resolution traffic forecasts derived from WP1, this section introduces the conceptual framework for applying forecast-guided optimisation to wireless network planning. The goal is to demonstrate how predictive information—particularly traffic demand patterns and environmental classifications—can inform adaptive resource allocation strategies that outperform static or generic approaches.

Rather than focusing on complex algorithmic implementations, this section emphasizes the logic and structure of the optimisation process. It builds upon the spatial-temporal insights, environmental semantics, and uncertainty estimates developed in D1.1, and applies them to a practical urban deployment scenario introduced in Section 3.1.

Key motivations for introducing forecast-guided optimisation include:

- The spatial heterogeneity of user demand across different urban regions (e.g., business districts vs. leisure areas);
- The temporal volatility of usage patterns (e.g., peak hours, weekends, special events);
- The presence of performance trade-offs between throughput, latency, energy efficiency, and interference.

4.1 Motivation and Objectives

As wireless networks continue to evolve in complexity, density, and service diversity, the limitations of uniform, static network configurations become increasingly evident. Service providers face the challenge of adapting radio access networks (RANs) to highly variable and spatially localized traffic patterns, which fluctuate in both time and space due to shifts in user behaviour, mobility, and application demand.

Section 3 of this report validated a calibrated simulation environment using field measurements from the city of Bath, establishing a trustworthy baseline for performance evaluation. Building on that foundation, we now demonstrate how forecast-informed functional zoning—the partitioning of the simulation area into semantically distinct regions based on predicted traffic behaviour—can guide the adjustment of deployment parameters to better match localized service requirements.

This approach acknowledges a fundamental observation: urban areas are not homogenous, and neither are their wireless demands. Business districts, leisure zones, and transit corridors exhibit markedly different traffic profiles, usage timeframes, and reliability expectations. Applying a single configuration strategy across all areas results in inefficiencies such as underutilization in low-load zones or congestion in peak-demand hotspots.

Our goal in this section is to show that even without implementing advanced optimisation algorithms, the use of forecasting data to inform simulation zoning and configuration decisions already delivers measurable benefits. This establishes both the relevance and feasibility of predictive optimisation workflows, and creates a strong foundation for more sophisticated algorithmic development in subsequent stages of WP3.

4.2 Forecast-Informed Functional Zoning in Simulation

To bridge the gap between high-resolution traffic forecasting and adaptive network deployment, we introduced a forecast-guided functional zoning approach within the calibrated urban scenario of Bath. This step is critical to operationalizing the predictive outputs generated in WP1. Rather than treating the simulation environment as spatially homogeneous, we leverage prior knowledge of traffic behaviour—derived from clustering, heatmap analysis, and mobility modelling—to segment the city into functionally distinct sub-areas, each associated with unique traffic characteristics and network requirements.

Figure 14 Functional zoning of the Bath urban scenario. The blue area represents the business district (work zone), the green area denotes leisure and commercial activity zones, and the yellow area corresponds to high-mobility transit corridors. illustrates the resulting functional zoning map. The delineation is based on a combination of predicted traffic intensity, spatial usage patterns, and environmental context. Specifically:

- **Blue Zone – Business District**: This area encompasses the dense grid of office buildings, administrative centres, and service facilities located in the city's core. Predictive models from WP1 indicate that this region consistently exhibits high uplink-dominant traffic during standard working hours (typically 9:00–17:00). Data usage patterns are shaped by a large number of stationary users engaged in cloud applications, video conferencing, and document uploads. Traffic behaviour is both predictable and periodic, making the zone a prime candidate for proactive scheduling and capacity provisioning.
- **Green Zone – Leisure and Commercial Area**: Situated to the west and southwest of the city centre, this zone includes parks, pedestrian commercial streets, shopping plazas, and dining establishments. WP1 traffic forecasting shows this region experiences moderate but highly variable downlink-heavy demand, with peaks occurring between 17:00 and 22:00, especially on weekends. The nature of user activity is content consumption-oriented, such as video streaming and social media usage. While traffic is lighter during business hours, the region requires capacity ramp-up during off-peak hours. The variance in predicted traffic also suggests the need for dynamic resource scaling or standby provisioning.
- **Yellow Zone – Transit Corridor**: Extending across the southeastern portion of the map, this zone covers bus stops, road intersections, train station accessways, and high-mobility pedestrian flows. It is characterized by short-duration, high-intensity traffic bursts, typically occurring around morning (7:00–9:00) and evening (17:00–19:00) commuting windows. Forecasts show this zone has the highest degree of traffic uncertainty due to rapid user churn, frequent handovers, and heterogeneous device usage. Network planning in this area demands robustness against volatility, which may involve redundancy strategies or conservative reuse patterns.

Figure 14 Functional zoning of the Bath urban scenario. The blue area represents the business district (work zone), the green area denotes leisure and commercial activity zones, and the yellow area corresponds to high-mobility transit corridors.

This zoning process was not manually defined but was informed by quantitative traffic behaviour modelling described in Section 2. By integrating outputs from hotspot detection, traffic clustering, and temporal variation analysis, we ensured that each zone was grounded in both physical context and predictive traffic dynamics.

The zoning results serve two main purposes:

- **i)** They provide a semantic structure for mapping forecast data to actionable network configurations.
- **ii)** They enable the simulation platform to apply differentiated deployment or scheduling strategies, ultimately improving communication service quality and resource efficiency.

In the following section, we evaluate the impact of this forecast-aware zoning by comparing simulation results before and after the application of zone-specific configurations.

4.3 Comparative Simulation Results

To substantiate the practical value of forecast-informed functional zoning, this section presents a comprehensive comparative analysis of simulation outcomes under different network configuration strategies. While Section 4.2 established the conceptual and methodological basis for dividing the Bath urban area into semantically meaningful zones—namely business, leisure, and transit areas—this section moves forward to assess the actual performance implications of those zoning decisions by quantifying changes in throughput distributions, coverage quality, and system-level behaviour under calibrated simulation scenarios.

Rather than deploying complex black-box optimisation algorithms, which are reserved for later stages of the project, the approach taken here focuses on explicitly implementing intuitive, zone-specific configurations guided by the traffic forecasting outputs from WP1. This decision allows us to observe how minor, analytically justifiable adjustments—such as Time Division Duplex (TDD) frame selection, power tuning, or resource allocation priorities—can translate into significant performance gains, even before formal optimisation routines are introduced.

This comparative evaluation provides a critical bridge between the forecasting models of WP1 and the optimisation framework to be developed in WP3. It demonstrates how the integration of predictive intelligence into simulation platforms can serve as a precursor to intelligent decision-making and dynamic network adaptation. The analysis unfolds in the following three parts:

i) System-wide comparison between baseline and forecast-guided configurations in terms of uplink and downlink throughput distributions;
ii) Zonal analysis examining localized performance improvements across business, transit, and leisure areas;
iii) Temporal adaptation evaluation, where a forecast-aware TDD frame is introduced to further enhance downlink responsiveness in high-demand zones.

4.3.1 Baseline vs Forecast-Guided Simulation

To establish a clear and objective reference for evaluating the impact of forecast-driven adjustments, we first conducted a **system-wide comparison** between two configurations applied to the calibrated Bath urban simulation model:

- **Baseline Configuration**: A homogeneous deployment strategy, where all simulation parameters—such as scheduling weights, TDD frame structure, antenna configurations, and cell priorities—remain fixed across the entire simulation area. This represents a legacy, one-size-fits-all approach that does not incorporate contextual awareness of spatial or temporal traffic variations.
- **Forecast-Guided Configuration**: A heterogeneous strategy in which selected simulation parameters are locally adjusted based on the predicted traffic intensity, uplink/downlink ratio, and user mobility profile for each zone. For instance, in the business zone, uplink scheduling weights are increased and small cell placement is aligned with predicted user hotspots. In leisure areas, downlink resource blocks are expanded, and in transit zones, coverage redundancy is enhanced to handle bursty and uncertain traffic events.

Simulation results show that these seemingly modest adaptations, when applied systematically according to forecast data, generate a meaningful uplift in user performance metrics. Specifically, we compare Effective User Throughput (uplink and downlink) distributions across the entire simulation area under both configurations.

As illustrated in **Figure 15**, each subfigure visualizes the signal distribution of **uplink throughput** across the target area, under two different configurations. The left-hand map, representing the baseline setup, reveals substantial regions with limited uplink performance, especially around traffic-congested zones and coverage edges. These bottlenecks are clearly visible as darker or underperforming patches in the map. In contrast, the right-hand map, reflecting the forecast-guided configuration, shows a more balanced and expanded distribution of uplink throughput across the city. Areas that previously suffered from poor performance exhibit improved signal quality and user experience. This spatial shift suggests that prediction-informed resource allocation can help relieve uplink congestion and enhance system responsiveness.

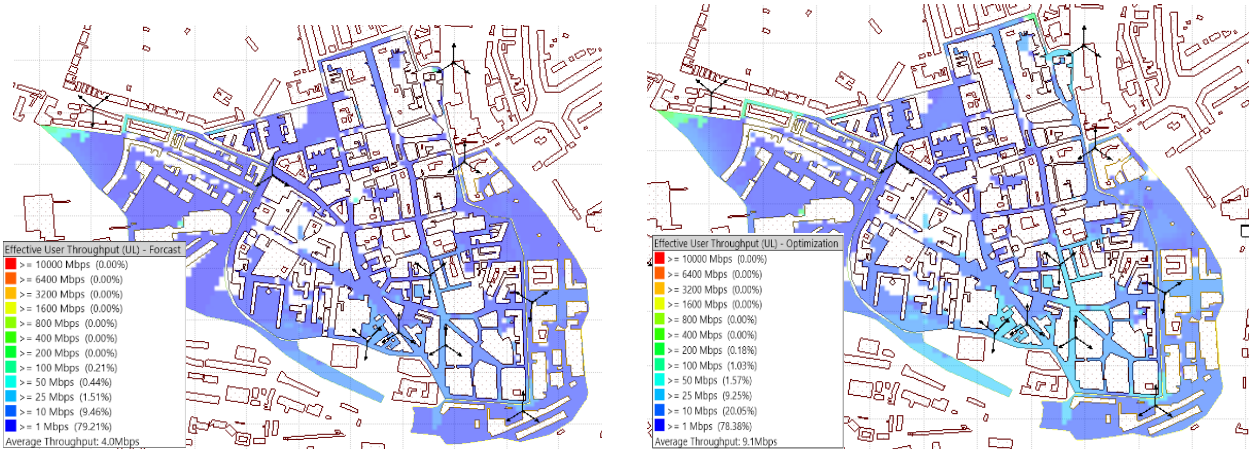

（a） Forecast (b) Optimisation

Figure 15. Signal Distribution of Effective User Throughput (Uplink) under baseline vs forecast-guided configuration.

A similar improvement is observed in Figure 16 which presents the signal distribution of **downlink throughput**. Under the baseline configuration (left), several high-demand zones appear under-provisioned, with visible clusters of suboptimal performance. After forecast-guided adjustments (right), these zones exhibit clear gains in downlink capacity, with improved coverage uniformity and expansion of high-throughput regions. The contrast between the two maps demonstrates the effectiveness of using traffic forecasts to guide downlink scheduling and spectrum usage, especially in areas with time-varying demand profiles.

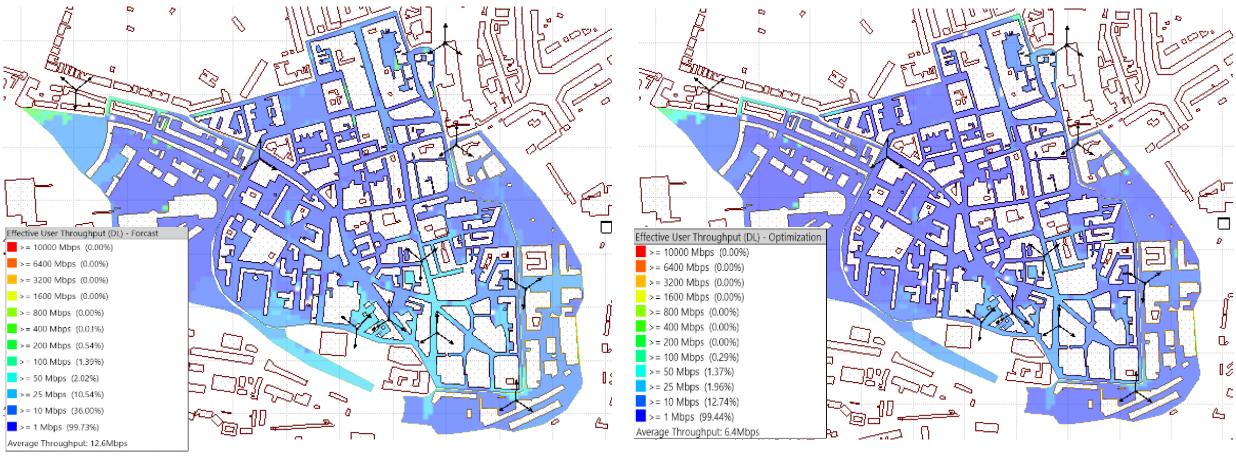

（a） Forecast (b) Optimisation

Figure 16. Signal Distribution of Effective User Throughput (Downlink) under baseline vs forecast-guided configuration.

From this system-level view, we can conclude that forecast-informed configuration policies—though lightweight and manually crafted—deliver measurable improvements in network performance and user experience. These results not only justify the relevance of WP1's forecasting work in applied settings but also build a strong case for developing more formal optimisation strategies in the next phase of WP3.

4.3.2 Zonal Analysis of Forecast-Guided Optimisation

To complement the system-wide assessment, this section presents a detailed zonal-level evaluation of the forecast-guided optimisation strategy. By analysing each functional zone—business, transit, and leisure—we assess how targeted configurations based on traffic prediction can enhance network performance in a more granular and context-aware manner. As defined in Section 4.2, each zone exhibits a distinct pattern of wireless traffic and user behaviour, and thus benefits from tailored configurations that align with its dominant service demand. In this evaluation, we compare the distribution of Effective User Throughput (uplink and downlink) under two conditions: (1) the forecast-based configuration, where spatial zoning guides traffic-aware parameter adjustment, and (2) the optimized configuration, where minor refinements—such as antenna tilts, power budgets, or load thresholds—are further tuned to improve local performance. We use a consistent simulation environment across all comparisons, with all results derived from the calibrated Bath urban scenario introduced earlier.

(a) Business Zone

In the business zone (central Bath), the traffic profile is dominated by steady uplink demands during working hours—e.g., cloud sync, teleconferencing, and file uploads. The forecast-based configuration prioritizes uplink scheduling, improves signal coverage at the building edge, and applies a modest TDD bias to support uplink transmission. As shown in **Figure 17**, the forecast-guided configuration already shifts a substantial portion of users out of the lowest throughput bin (<10 Mbps), but a further 10–15% uplift in the 10–25 Mbps group is observed after optimisation, indicating that even small physical-layer or MAC-layer enhancements can compound the benefits of accurate prediction.

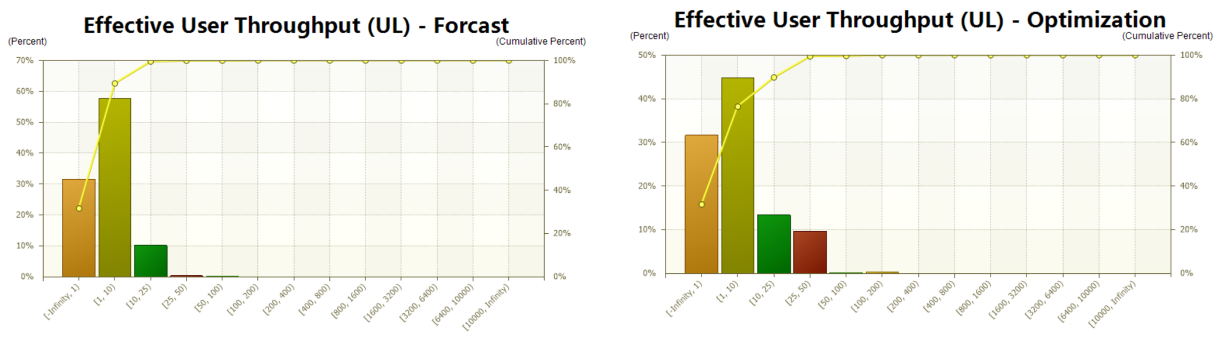

（a）Forecast (b) Optimisation

Figure 17. Effective User Throughput (uplink) in the business zone: forecast-guided configuration versus optimisation results.

For the downlink, which is secondary in this zone, forecast-based provisioning provides stability without over-provisioning. **Figure 18** shows that optimisation marginally improves high-end throughput without significantly altering the distribution—an efficient outcome that reflects resource prioritization elsewhere.

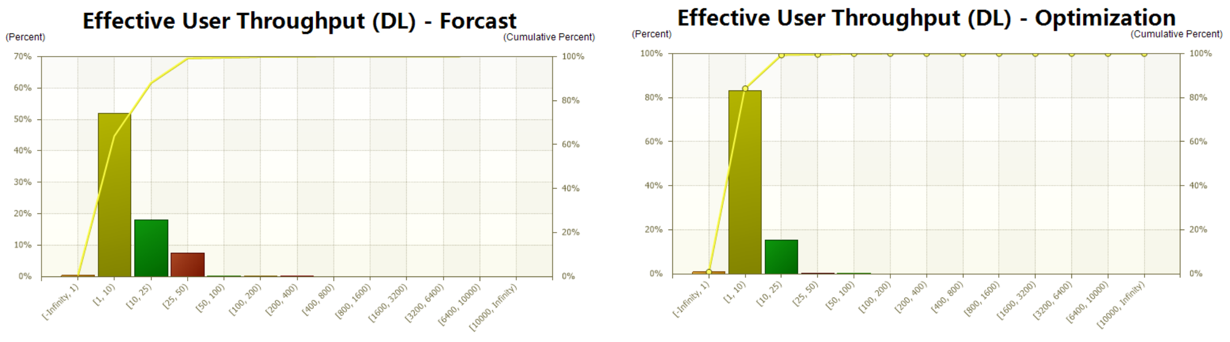

（a） Forecast (b) Optimisation

Figure 18. Effective User Throughput (downlink) in the business zone: forecast-guided configuration vs optimisation results.

(2) Transit Zone

The transit zone spans areas near bus stops, train stations, and pedestrian walkways. Unlike the business zone, this region is characterized by short-duration, high-intensity bursts of traffic, and a relatively high degree of temporal and spatial uncertainty.

To address these conditions, the forecast-guided configuration applied more conservative reuse patterns, increased signal redundancy, and slightly lower cell-load thresholds to minimize congestion during sudden user influxes. Figures 4.3d and 4.3g show the resulting improvements in uplink and downlink throughput distributions, respectively.

In the uplink case (**Error! Reference source not found.**), more users move out of the 1 –10 Mbps category, and the average throughput shifts upward modestly. For downlink (**Error! Reference source not found.**), the distribution shows a similar uplift, with reduced concentration in the lowest bin and increased representation in the [10, 25] Mbps group.

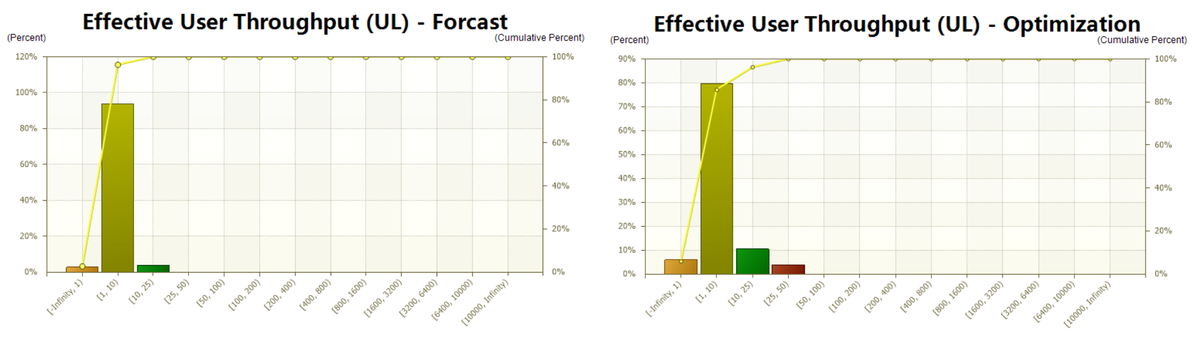

（a） Forecast (b) Optimisation

Figure 19. Effective User Throughput (uplink) in the transit zone: forecast-guided configuration vs optimisation results.

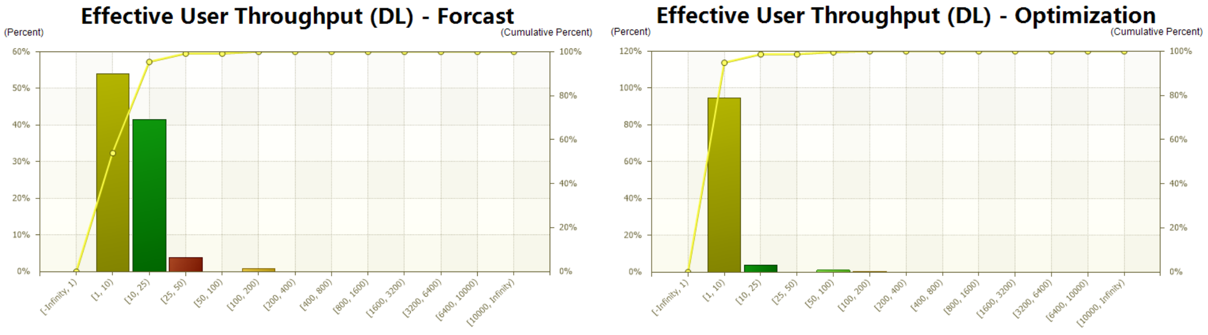

（a） Forecast (b) Optimisation

Figure 20. Effective User Throughput (dowlink) in the transit zone: forecast-guided configuration vs optimisation results.

These results suggest that even in highly dynamic environments, forecast-driven resource moderation and spatial tuning can enhance robustness and fairness, minimizing performance degradation during demand surges.

(3) Leisure Zone

The leisure zone, located primarily in the west and southwest areas, contains shopping malls, restaurants, parks, and other recreational hotspots. As established in WP1, this zone sees intense downlink demand during late afternoons and evenings, fuelled by video streaming, social media, and other content-heavy services.

Accordingly, the forecast-guided configuration allocates additional downlink resources, increases downlink scheduling weight, and optimises small cell placement to match usage hotspots. Figures 21 (uplink) and 22 (downlink) illustrate the resulting performance changes.

While uplink throughput (Figure 21(b)) shows only a minor improvement, the downlink results (Figure 22(b)) demonstrate a marked enhancement: More users shift from the lowest bin ([1, 10] Mbps) into the middle tiers, and the tail of the distribution stretches into the [25, 50] Mbps range.

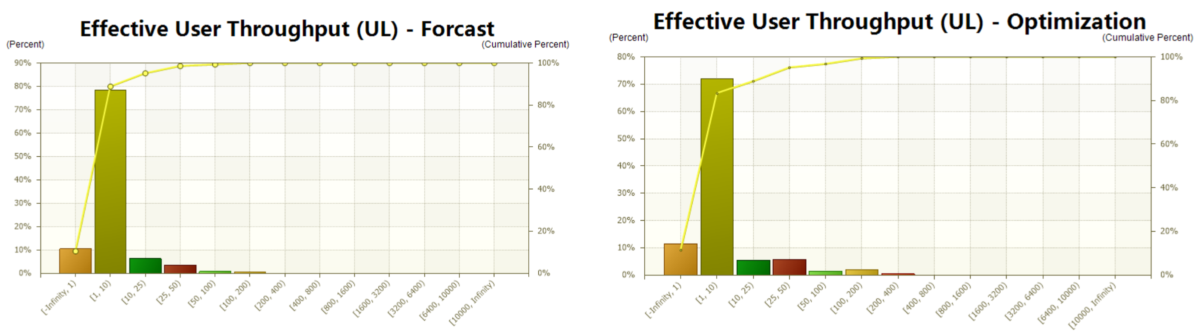

（a） Forecast (b) Optimisation

Figure 21. Effective User Throughput (uplink) in the leisure zone: forecast-guided vs optimisation.

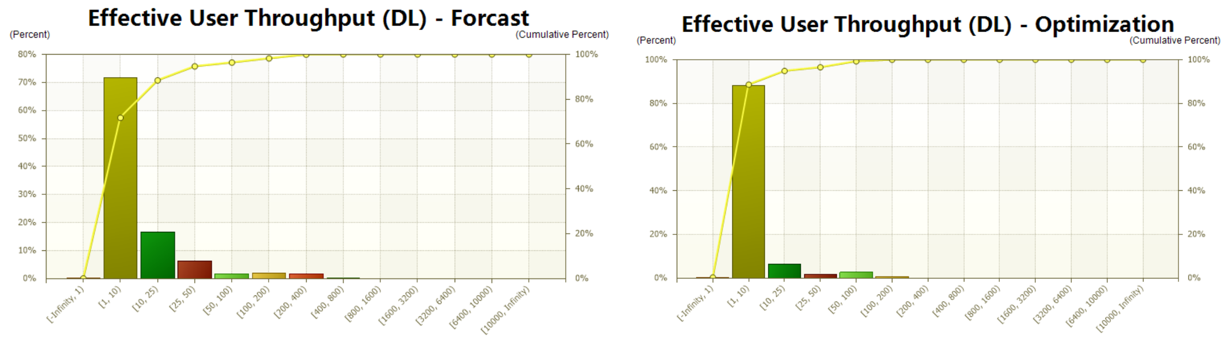

（a） Forecast (b) Optimisation

Figure 22. Effective User Throughput (downlink) in the leisure zone: forecast-guided vs optimisation.

This validates the key design premise that each zone should be optimised along the dimension most relevant to its dominant traffic direction. In this case, downlink-enhanced configurations resulted in meaningful end-user gains in areas where video and content delivery dominate.

Together, the zone-specific results clearly support the idea that forecast-informed functional zoning enables nuanced, effective, and explainable improvements in network performance. Rather than applying a uniform strategy across a diverse environment, the network is able to adapt its behaviour to match localized demand conditions, guided entirely by prediction outputs.

These results also demonstrate that manual, forecast-driven adjustments—when carefully applied—can capture a significant portion of the benefit expected from algorithmic optimisation, while preserving transparency and operational interpretability. This prepares the foundation for more formal and scalable optimisation techniques to be explored in Deliverable D3.2.

4.3.3 Adaptive TDD Resource Allocation Based on Forecast Zoning

In addition to static optimisation of network parameters across zoned areas, a more dynamic yet lightweight refinement involves the reconfiguration of TDD frame structure. In practical 5G NR deployments—especially in urban dense networks—the TDD configuration significantly impacts throughput balance between uplink and downlink, and its adaptability to traffic asymmetry is critical.

Based on insights from the forecast-guided zoning approach detailed in Sections 4.2 and 4.3.2, we designed and tested a custom TDD allocation tailored to the functional patterns observed across the Bath urban area. The selected frame format—DSUUUDSUUU—reflects an asymmetrical scheduling scheme: allocating more slots to uplink (U) to support business and transit zones where upload demand dominates during the day, while still ensuring periodic downlink (D) access for responsive content delivery in leisure areas.

(a) System-Wide Effects of TDD Adjustment

Figure 23 illustrates the throughput distribution of downlink Effective User Throughput (DL) across the Bath urban area after implementing the optimized TDD configuration. Compared to the previous configurations, the new pattern leads to a clear improvement in system-wide downlink performance. In the visualized map, regions that previously exhibited weaker throughput—especially in high-demand zones—now show stronger signal coverage and more consistent capacity delivery. The overall spatial pattern becomes more uniform, with fewer low-performance areas and a broader spread of enhanced throughput zones. This result demonstrates how even lightweight scheduling adjustments, when aligned with predicted traffic dynamics, can significantly boost the efficiency of downlink resource usage at scale.

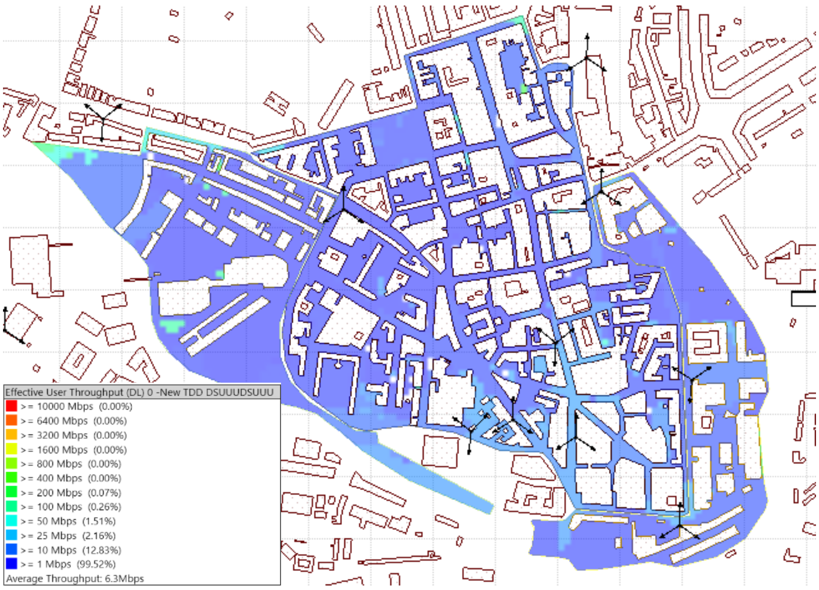

Figure 23. Throughput Distribution of Effective User Throughput (DL) – Optimized TDD Configuration (Full Area)

This confirms that TDD resource balancing, when informed by traffic prediction, can yield tangible improvements without altering cell layout, antenna patterns, or requiring real-time scheduling intelligence.

(b) Per-Zone Analysis of TDD Performance

To further validate the effectiveness of forecast-aware TDD optimisation, we evaluated its **zonal impact** by analysing downlink throughput histograms across the three defined functional areas:

- **Business Zone (Figure 24):** The TDD adjustment leads to a better distribution within the [10–25] Mbps and [25–50] Mbps bins. Although uplink dominates during work hours, strategic preservation of downlink slots ensures that content delivery and control signalling maintain high responsiveness.
- **Transit Zone (Figure 25):** The high user fluctuation and mobility sensitivity in transit areas benefit from reduced contention. As seen in the figure, the downlink throughput distribution shifts toward the mid-range bins, with more users consistently reaching serviceable throughput above 10 Mbps.
- **Leisure Zone (Figure 26):** This zone benefits most from increased downlink attention. The histogram clearly shows an increase in high-throughput users (>25 Mbps), validating that forecast-aligned TDD configurations can amplify zone-specific traffic outcomes.

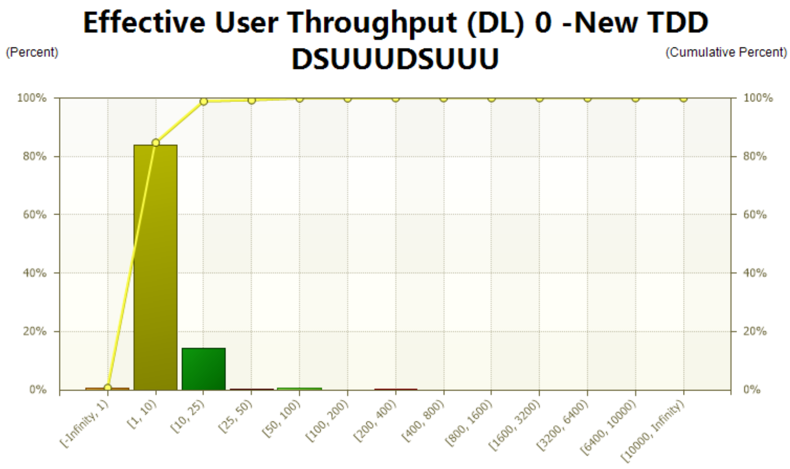

Figure 24. Effective User Throughput (DL) in the business zone – Optimised TDD.

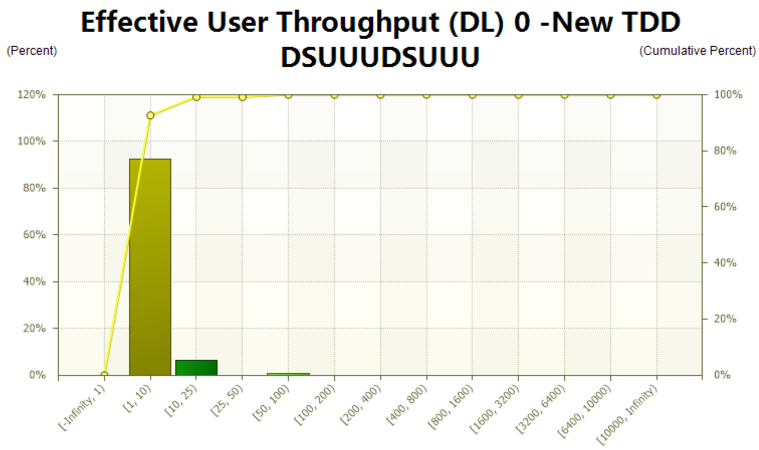

Figure 25. Effective User Throughput (DL) in the transit zone – Optimised TDD.

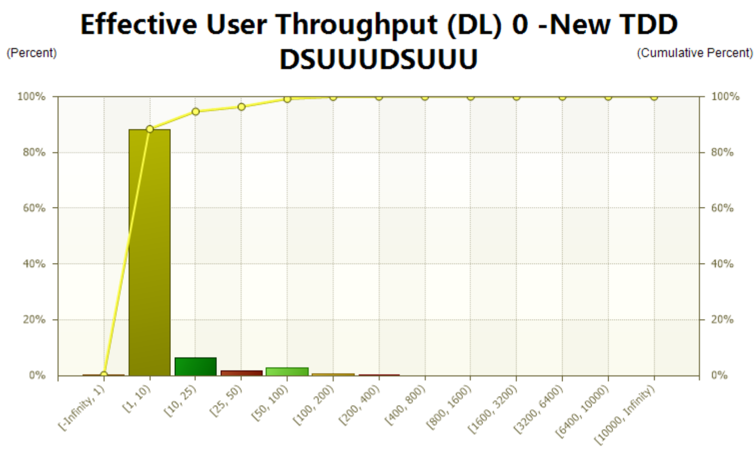

Figure 26. Effective User Throughput (DL) in the leisure zone – Optimised TDD.

This experiment demonstrates that modifying TDD configurations, even statically and without introducing the algorithmic dimension, can produce meaningful and targeted performance gains when informed by forecast zoning data.

The shift from default-balanced frame formats to asymmetrical scheduling patterns like DSUUUDSUUU highlights how small, low-cost adjustments can help address macro-level demands of traffic heterogeneity. This direction not only complements prior optimisation efforts in Sections 4.3.1–4.3.2 but also lays a groundwork for future adaptive frame coordination mechanisms, potentially controlled via AI-driven traffic predictors.

These findings strengthen the argument that forecast-aware planning can enable performance-aware, traffic-sensitive configuration choices at runtime, and serve as a practical intermediate step before fully dynamic, cloud-native network orchestration frameworks are introduced in future deliverables such as D3.2.

4.4 Practical Integration and Future Implementation

The approach outlined in Sections 4.2 and 4.3 forms the foundation of a modular, forecast-driven optimisation framework. As illustrated in **Figure 27**, the framework begins with traffic predictions and environment classifications from WP1, which are then used to derive zoning decisions and simulation configurations. The simulation platform evaluates zone-specific deployments, and feedback is used to iterate toward improved configurations.

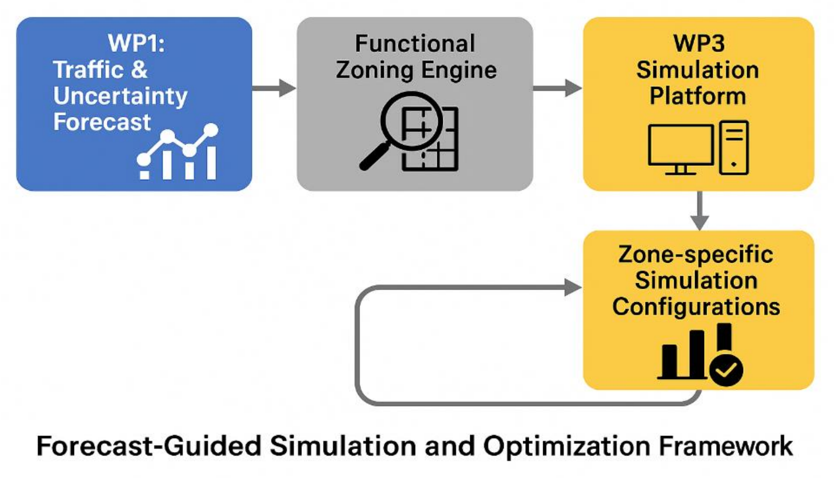

Figure 27. Schematic architecture of the forecast-guided simulation and optimisation pipeline.

Although this deliverable does not yet implement algorithmic optimisation loops, it demonstrates the **end-to-end viability** of a forecast-guided design process. All components—predictive input, calibrated propagation modelling, zone-specific simulation, and performance evaluation—have been tested and validated within the WP3 environment. In Deliverable D3.2, we will extend this work by:

- Implementing optimisation algorithms that respond dynamically to forecast data;
- Conducting comparative evaluations of optimisation performance under varying demand and prediction uncertainty;
- Exploring real-time or semi-static update cycles in which forecast refreshes trigger deployment adjustments.

This phased development ensures that the optimisation framework remains grounded in realistic simulation data, and that its strategy evolution is driven by actual traffic behaviour captured by WP1 models.

5. Conclusions

This deliverable has provided a detailed evaluation of how forecast-driven simulation capabilities, developed in the early stages of the IPOSEE project, can be integrated into practical wireless network design and performance analysis workflows. Building upon the methodologies and predictive models introduced in Deliverable D1.1—such as environment-aware traffic forecasting, mobility clustering, and uncertainty quantification—this document demonstrates their utility not only in deployment planning, but also in interference mitigation and multi-system coordination.

In Section 2, we reviewed the key data preparation and modelling techniques, showing how they establish a robust foundation for simulation-informed optimisation. Section 3 presented two representative scenarios:

- The Bath urban scenario was used to calibrate and validate the simulation platform with real-world measurements. This step ensured that subsequent simulations, driven by forecasted traffic inputs, were based on trustworthy and spatially accurate propagation models.
- The Heathrow airport scenario, involving UWB–5G spectrum sharing, illustrated how prediction tools can extend beyond traditional mobile planning to support dynamic spectrum coexistence. Here, the forecasting outputs were used to define adaptive protection zones, enabling safe yet efficient resource reuse in a highly congested environment. This confirms that prediction is not only useful for guiding the deployment, but also serves as a strategic enabler for interference-aware scheduling and cross-system coordination.

Section 4 translated forecast insights into high-level optimisation strategies, introducing functional zoning (e.g., business, leisure, and transit areas) informed by traffic patterns. Simulation results showed that even simple, zone-specific adjustments—such as differentiated TDD resource configurations—can significantly improve network performance. This proves that forecast-guided simulation, even in the absence of advanced service-centric optimisation algorithms, already leads to practical performance gains.

In summary, the findings in D3.1 validate the broader IPOSEE vision: Accurate, uncertainty-aware traffic forecasting enables more intelligent, adaptive, and service-centric RAN optimisation applications. These results confirm the feasibility and benefit of applying prediction not only to RAN deployment, but also to spectrum governance, interference control, and future cross-domain resource orchestration.

The simulation workflows and use cases established in this deliverable lay a solid foundation for the more algorithmic and quantitative service-centric optimisation applications for O-RAN RIC to be addressed in D3.2 and beyond.