UNICO I+D Project
6G-SORUS-RAN

SORUS-RAN-A3.2-E2 (E17)

Final performance evaluation

Abstract

This deliverable explores the algorithmic challenges and performance limits of dynamic reconfiguration in virtualized Radio Access Networks (vRAN) for B5G systems. By integrating realistic, mobility-driven traffic profiles into the analysis, it evaluates how spatial and temporal reconfiguration strategies can optimize throughput and energy efficiency under varying user demands. The study also considers the interplay between vRAN, UAV-assisted access, and Reconfigurable Intelligent Surfaces (RIS), highlighting the benefits of coordinated, mobility-aware orchestration. The results provide key insights into the potential of predictive and adaptive network management, laying the groundwork for intelligent, flexible, and energy-efficient B5G infrastructures.

Financiado por la Unión Europea NextGenerationEU

GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES

Plan de Recuperación, Transformación y Resiliencia

UNICO I+D

Document properties

Document number	SORUS-RAN-A3.2-E2 (E17)
Document title	Final Performance Evaluation
Document responsible	Marco Gramaglia
Document editor	Juan Manuel Montes
Editorial team	Pablo Serrano, José Gallego
Target dissemination level	Public
Status of the document	Final
Version	1.0
Delivery date	30-06-2025
Actual delivery date	30-06-2025

Production properties

Reviewers	María Molina

Disclaimer

This document has been produced in the context of the 6G-SORUS Project. The research leading to these results has received funding from the Spanish Ministry of Economic Affairs and Digital Transformation and the European Union-NextGenerationEU through the UNICO 5G I+D programme.

Contents

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

List of Figures

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

List of Acronyms

5GC: 5G Core

AF: Access Function

AI: Artificial Intelligence

API: Application Programming Interface

ARIMA: Autoregressive Integrated Moving Average

B5G: Beyond 5G

CAPEX: Capital Expenditure

CU: Centralized Unit (in vRAN architecture)

DL: Downlink

DT: Digital Twin

E2E: End-to-End

eMBB: Enhanced Mobile Broadband

ETSI: European Telecommunications Standards Institute

gNB: Next-Generation NodeB (5G Base Station)

IoT: Internet of Things

KPI: Key Performance Indicator

LLC: Logical Link Control

LLC: Low Latency Communications

LSTM: Long Short-Term Memory

MEC: Multi-access Edge Computing

ML: Machine Learning

MLP: Multi-Layer Perceptron

mMTC: Massive Machine-Type Communications

NF: Network Function

NFV: Network Function Virtualization

NR: New Radio (5G air interface)

OPEX: Operational Expenditure

OSS: Operations Support System

QoE: Quality of Experience

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

QoS: Quality of Service

RAN: Radio Access Network

RIS: Reconfigurable Intelligent Surface

RU: Radio Unit (in vRAN architecture)

SDN: Software Defined Networking

SINR: Signal-to-Interference-plus-Noise Ratio

SVM: Support Vector Machine

UE: User Equipment

UAV: Unmanned Aerial Vehicle

UL: Uplink

URLLC: Ultra-Reliable Low Latency Communications

vRAN: Virtualized Radio Access Network

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

Resumen Ejecutivo

Este Entregable analiza los retos algorítmicos y los límites de rendimiento de la virtualización en las redes de acceso radio (vRAN) en el contexto de las redes B5G, centrándose específicamente en las ventajas y limitaciones de la reconfiguración dinámica. Las vRAN están transformando el panorama de las telecomunicaciones al desacoplar las funciones de red del hardware dedicado, lo que permite una flexibilidad sin precedentes para soportar una amplia gama de servicios y aplicaciones, desde la automatización industrial con requisitos de latencia ultrabaja hasta las ciudades inteligentes impulsadas por IoT. Sin embargo, esta flexibilidad también introduce nuevos desafíos, como la gestión de patrones de tráfico altamente variables y de usuarios distribuidos espacialmente.

Basándonos en nuestro trabajo previo sobre modelización de la movilidad y generación de datos sintéticos, integramos perfiles de demanda realistas basados en la movilidad en el estudio de las estrategias de reconfiguración de vRAN. Al capturar los patrones espacio-temporales de conectividad de los usuarios, podemos evaluar cómo optimizar la reconfiguración de la red, tanto a nivel espacial (ajustes por celda) como temporal (asignación dinámica de recursos), con el objetivo de mejorar el rendimiento y la eficiencia energética. Esta aproximación permite realizar evaluaciones más precisas de los límites operativos de la gestión dinámica de vRAN y resalta el potencial de mecanismos de orquestación predictivos sensibles a la movilidad.

En el contexto más amplio de la orquestación B5G, este trabajo se alinea también con los paradigmas emergentes que combinan vRAN con puntos de acceso asistidos por UAV y superficies inteligentes reconfigurables (RIS). Los UAV permiten extender la cobertura bajo demanda en áreas críticas, mientras que las RIS pueden modificar de manera dinámica el entorno de propagación inalámbrica, complementando las capacidades de virtualización de vRAN. La integración de estrategias de reconfiguración informadas por movilidad con despliegues asistidos por UAV y RIS posibilita un funcionamiento de la red más adaptable y energéticamente eficiente, especialmente en escenarios caracterizados por cambios rápidos en la distribución de usuarios o en la demanda de servicios.

Al cuantificar los límites de rendimiento de la reconfiguración dinámica de vRAN y demostrar su interacción con modelos de movilidad realistas, este informe sienta las bases para marcos de coordinación que integren vRAN, UAV y RIS. Este enfoque representa un paso fundamental hacia infraestructuras B5G inteligentes, sostenibles y altamente adaptables, capaces de satisfacer los exigentes requisitos de los servicios de próxima generación.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

Executive Summary

This deliverable investigates the algorithmic challenges and performance limits of virtualization in Radio Access Networks (vRAN) within the context of B5G networks, focusing specifically on the benefits and constraints of dynamic reconfiguration. As vRAN continues to reshape the telecommunications landscape by decoupling network functions from dedicated hardware, it provides unprecedented flexibility for supporting a diverse range of services and applications—from ultra-low-latency industrial automation to IoT-driven smart cities. However, this flexibility also introduces complexity in managing highly variable traffic patterns and spatially distributed users.

Building upon our prior work on mobility modeling and synthetic data generation, we integrate realistic, mobility-driven demand profiles into the study of vRAN reconfiguration strategies. By capturing spatiotemporal patterns of user connectivity, we can evaluate how network reconfiguration—both spatial (cell-level adjustments) and temporal (time-dependent resource allocation)—can be optimized to improve throughput and energy efficiency. These insights enable a more accurate assessment of the operational boundaries of dynamic vRAN management and highlight the potential for predictive, mobility-aware orchestration mechanisms.

In the broader context of B5G orchestration, this work also aligns with emerging paradigms that combine vRAN with UAV-assisted access points and Reconfigurable Intelligent Surfaces (RIS). UAVs can provide on-demand coverage extensions, while RIS can dynamically reshape the wireless propagation environment, both of which complement vRAN's virtualization capabilities. The integration of mobility-aware reconfiguration strategies with UAV- and RIS-assisted deployments can unlock highly adaptive, energy-efficient network operation, particularly in scenarios with rapidly changing user distributions or service demands.

By quantifying the performance limits of vRAN reconfiguration and demonstrating its interplay with mobility-driven insights, this deliverable provides a foundation for coordinated orchestration frameworks that unify vRAN, UAV-based access, and RIS-enabled coverage enhancements. This approach represents a critical step toward achieving intelligent, sustainable, and highly adaptive B5G network infrastructures capable of meeting the stringent demands of next-generation services.

1. Introduction

This deliverable aims to study the algorithmic challenges and state-of-the-art solutions for the virtualization of Radio Access Networks (vRAN) in B5G networks. More specifically, we perform a literature review on relevant scientific work addressing different aspects of vRAN implementations in various applications. This analysis delves into key areas such as the optimization of network resources, the enhancement of network flexibility, innovations in network architecture, advanced and functional network slicing, improvements in interoperability, and the development of energy-efficient protocols.

This study is motivated by the significant transformation the telecommunication industry has witnessed, driven by the rapid evolution of wireless services and the rising demand for flexible and scalable communication infrastructures. At the forefront of this transformation is the emerging trend of virtualizing wireless networks. By decoupling network services from the underlying physical infrastructure, vRAN introduces an unprecedented level of flexibility and programmability, enabling more dynamic and efficient use of network resources.

As virtualization technology continues to evolve, its principles are increasingly applied to wireless network infrastructures, leading to the concept of vRAN. This paradigm shift allows service providers to instantiate virtual networks tailored to specific services or customer requirements, all on a shared physical infrastructure. Such capabilities are crucial in the 5G and B5G era, where a diverse set of applications—including high-definition video streaming, autonomous vehicles, smart cities, and the Internet of Things (IoT)—demand highly agile, adaptable, and performance-aware network architectures.

Building on the insights gained from our previous work on mobility modeling and synthetic data generation, this deliverable leverages a deeper understanding of how users move and interact with the wireless network to better assess the limits of reconfiguration gains in vRAN. Accurate modeling of device mobility enables more realistic evaluation scenarios, where spatial and temporal variations in user demand are explicitly captured. This allows us to analyze how dynamic reconfiguration—such as adjusting cell configurations or reallocating resources in real time—can be informed by user mobility patterns, improving both throughput and energy efficiency. By incorporating mobility-aware demand forecasting, we can better quantify the performance boundaries of vRAN reconfiguration strategies and identify scenarios where mobility-driven adjustments lead to the most significant benefits.

In this way, the lessons learned from mobility modeling not only enhance the realism of our evaluation but also provide a foundation for exploring how predictive, mobility-aware algorithms could be integrated into vRAN orchestration frameworks. Ultimately, this synergy allows us to connect user behavior, network virtualization, and dynamic reconfiguration, advancing the design of highly adaptive and resource-efficient B5G networks.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

2. Related Work

Mobile networks are integral to providing constant connectivity, not only for individual users but also for a wide array of services, devices, and sensors, making them essential in both daily life and business operations. These networks rely on a highly complex support infrastructure that includes thousands of access cells, each of which requires careful configuration and ongoing maintenance by network operators. This complexity results in significant operational costs and a heavy management burden. The behavior of each individual cell is influenced by neighboring cells and ever-changing network conditions, including the number of connected users, traffic types, and mobility patterns. These factors contribute to an increasingly dynamic and challenging network environment as the demand for new services grows and new technologies are integrated [SMM2020].

In the context of RAN and vRAN, this complexity is further magnified. As networks evolve to meet the demands of modern connectivity, operators are tasked with optimizing infrastructure usage while managing operational challenges [AH2015]. The complexity of managing these networks is expected to increase, especially as operators must deliver enhanced services with fewer financial gains. The need to optimize resource allocation in this context is more critical than ever [YDA2020].

Network slicing has emerged as a promising solution, offering greater agility by enabling networks to dynamically adapt to fluctuations in traffic loads and service demands. This evolution, supported by studies demonstrating the benefits of network adaptability through real-world data, allows for better utilization of network resources. However, this flexibility also introduces added complexity in network management. As part of this broader effort to optimize network performance, our research examines how more frequent network configuration updates could improve efficiency, particularly in resource allocation and management [MGF2018].

In this study, we focus on the optimization of network configurations, particularly analyzing the impact of more frequent reconfigurations on downlink traffic performance—a key metric for network operators. By leveraging real-world data, we explore various configuration patterns and scales, evaluating the trade-offs between increased adaptability and the associated management overhead. Our goal is to find a balance that maximizes performance while minimizing complexity, aligning with broader industry trends toward smarter, more efficient network resource allocation in RAN and vRAN environments.

Several studies have explored optimizing the configuration of mobile access networks, using a variety of objectives and methodologies. For instance, [MAG2020] analyzes real network traces to identify anomalies and suggests mitigation strategies through the reconfiguration of cell clusters. Similarly, another study explores the use of neural networks to predict network behavior based on its configuration. While this work considers transmitted traffic as an environmental variable, our study focuses on downlink traffic as the primary Key Performance Indicator (KPI), as defined by the network operator.

Another relevant approach [VSP2023] investigates the use of Digital Twins to test configurations in RANs and train reinforcement learning models. However, the solution is only tested on a small

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

network, consisting of five nodes, limiting its applicability to larger and more complex RAN environments. Additionally, another [PJW2024] study presents a machine learning framework for optimizing configuration parameters of cells in mobile access networks, targeting Signal to Interference and Noise Ratio (SINR) as the KPI. However, this work only optimizes two configuration parameters.

In contrast, [AFI2018] propose a machine learning-based solution aims at maximizing cell coverage while minimizing interference, though it does not account for traffic or user activity in the optimization process. Meanwhile, another study employs a support vector machine (SVM) approach to maximize user throughput in RANs, although the reference scenario for their results is not specified. Similarly, another study investigates load balancing optimization in 5G networks through the configuration of cell parameters, though it is tested on a synthetic network.

Lastly, another study addresses [HN2015] the parameter optimization problem from the end-user perspective, using Quality of Experience (QoE) as the objective function. This solution is tested via simulations on a 19-cell reference network.

Overall, the existing literature provides various methodologies for optimizing cell configuration parameters, targeting different objectives such as coverage, SINR, load balancing, and user throughput. However, to the best of our knowledge, this work represents the first case of optimizing RAN configurations directly involving a real network operator, using real-world data and operational insights. Our findings demonstrate potential increases in network traffic capacity of over 30%, offering valuable insights into which configuration changes provide the greatest improvements in network performance.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

3. Network Overview

In this section, we examine the access network architecture of one of Brazil's largest mobile network operators, providing valuable insights into the resource allocation challenges inherent in modern mobile networks, particularly in the context of LTE and 5G readiness. This network, composed of 1,553 cells across multiple municipalities (Itaguaí, São Gonçalo, and Petrópolis), presents a complex scenario with varying bandwidth configurations to meet fluctuating service demands across urban and suburban areas. Such variability highlights the critical need for optimized resource allocation strategies, especially as operators prepare for the transition to vRAN and enhanced network slicing in 5G.

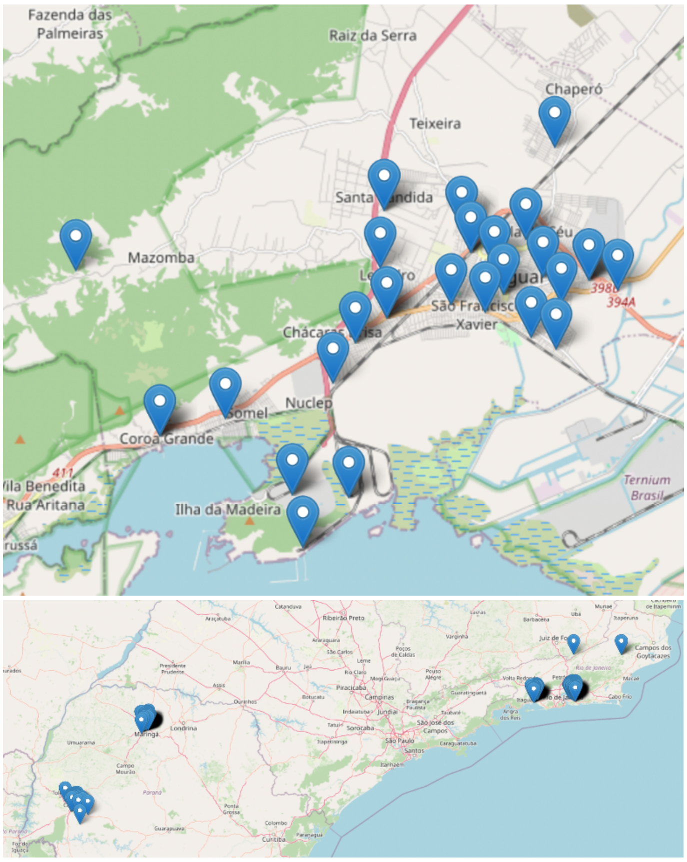

FIGURE 1: LOCATION OF THE CELLS IN THE CONSIDERED NETWORK (TOP), AND IN THE ITAGUAI MUNICIPALITY IN DETAIL (BOTTOM).

The network operates over four distinct frequency bands: 700 MHz (Low Band), 1.8 GHz and 2.1 GHz (Mid Band), and 2.6 GHz (High Band). The heterogeneous nature of these deployments, with each node using a unique combination of frequencies, demonstrates the need for advanced optimization techniques. This distribution of frequencies across cells, where over 42% utilize three bands and the remainder employ a varying mix of one to four frequencies, underscores the complexity of network resource management. Ensuring efficient use of these frequency resources

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

requires dynamic and intelligent allocation systems, particularly as mobile networks evolve towards virtualization through vRAN.

The key performance indicator (KPI) used to assess the efficiency of network configurations is the Downlink Traffic (DL) Volume, a metric prioritized by the network operator due to its direct impact on user experience and network performance. This focus aligns with broader industry trends that emphasize traffic-based resource allocation and performance optimization in RAN and vRAN deployments. Addressing missing or faulty KPI measurements—an issue affecting around 20% of the data—remains a challenge, emphasizing the importance of real-time, fault-tolerant resource management algorithms in ensuring robust service delivery.

3.1. Network Configuration Optimization

Within this heterogeneous network, cell configurations are selected from a predefined set of parameters based on performance data and vendor recommendations. These configurations are crucial in determining how the network manages its resources, particularly in scenarios involving inter-frequency handovers and traffic redistribution across cells. The balance between traffic load and signal quality is critical for ensuring optimal network performance and minimizing energy consumption, two key concerns in the evolution of RAN into vRAN architectures.

The parameters guiding these configurations—such as *Cellreselpriority, QRxLevMin*, and *SNonIntraSearch*—play a significant role in determining how user equipment (UE) interacts with the network. These parameters control when UEs should switch between frequencies or cells based on signal strength and network load, ensuring seamless transitions and balanced resource allocation. This process of adjusting configurations in response to dynamic network conditions mirrors the emerging strategies in vRAN, where intelligent resource orchestration must be flexible enough to optimize both user experience and operational efficiency.

The network's reliance on predefined configurations, while practical, also limits its flexibility and scalability. Increasing the granularity of these configurations, particularly in problem areas, introduces significant management overhead. However, modern vRAN systems are moving towards more sophisticated, automated configuration management through machine learning and predictive analytics, reducing the need for manual adjustments and trial-and-error approaches.

An example of this configuration management process can be seen in Figure 2, where inter-frequency handovers are regulated within a node using parameters like *SNonIntraSearch*. In such scenarios, advanced algorithms could further optimize the thresholds and signal measurements to ensure a more dynamic and energy-efficient handover process. As vRAN deployments become more common, the ability to fine-tune these parameters in real-time without excessive overhead will be critical for achieving the energy efficiency and scalability required by future 5G networks.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

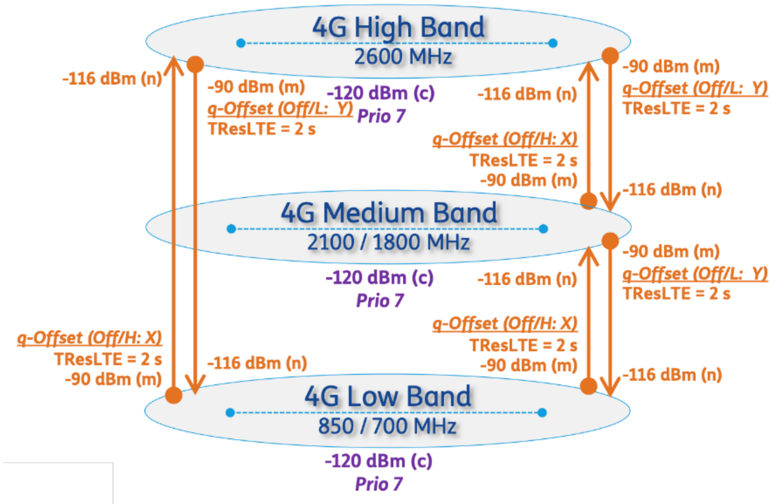

FIGURE 2: EXAMPLE OF NETWORK CONFIGURATION.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

4. Baseline Network Performance

This section provides an overview of the baseline performance of the network using the operator's default configuration, referred to as **configuration 0**. Configuration 0 is widely considered the most effective in terms of maximizing downlink (DL) traffic, the primary Key Performance Indicator (KPI) used by the operator to assess network throughput. By analyzing the baseline configuration, we can better understand how it performs under typical conditions and use this as a benchmark for evaluating alternative network configurations and their impact on performance.

4.1. Downlink Traffic Analysis

To begin the analysis, we examine the hourly DL traffic over the course of a typical week, as shown in Figure 3. The results illustrate the dynamic nature of network traffic, with an average DL traffic volume of approximately 3.4 billion KB per hour. During peak periods, traffic spikes to as much as 5.7 billion KB, while it drops to a minimum of 0.6 billion KB during off-peak hours. This variability highlights the importance of optimizing resource allocation in real-time, as traffic loads fluctuate significantly throughout the day.

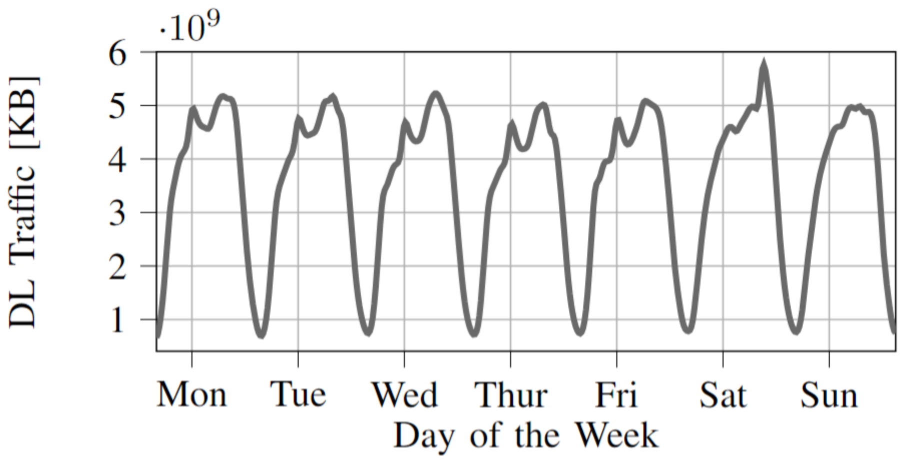

FIGURE 3: HOURLY TRAFFIC ACROSS THE ENTIRE NETWORK OVER A WEEK, USING THE BASELINE CONFIGURATION.

The data in Figure 3 reflects typical day-night traffic cycles, with noticeable peaks during midday and evening hours. The relatively stable traffic patterns observed during weekends suggest that network usage remains high, even as the number of connected users may vary. This consistency underscores the necessity of implementing adaptive resource allocation strategies within the RAN and vRAN, allowing the network to handle diverse traffic loads while maintaining high performance and energy efficiency.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

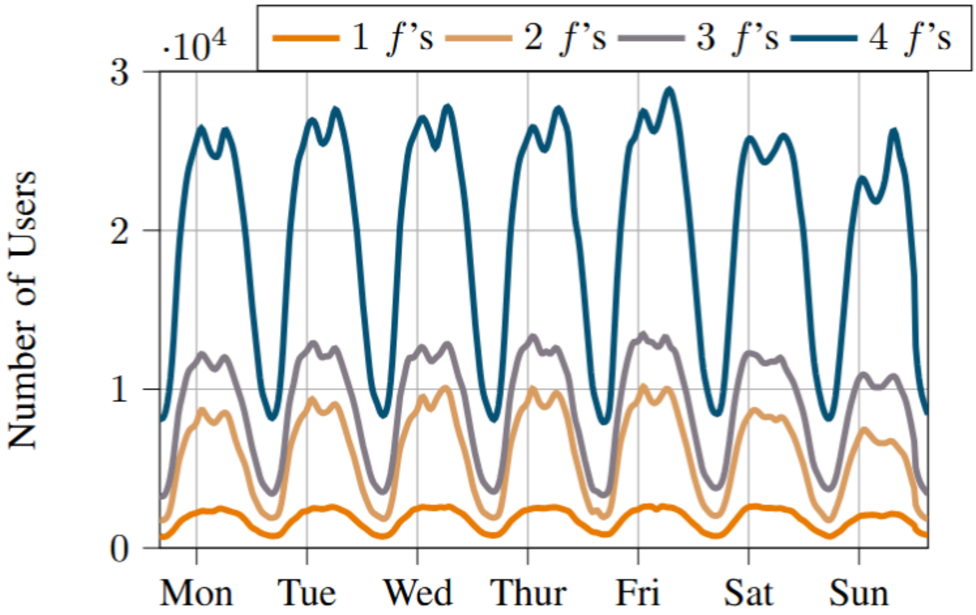

FIGURE 4: USERS CONNECTED PER HOUR AND NUMBER OF FREQUENCIES.

Figure 4 illustrates the number of users connected to antennas with varying numbers of frequency bands (from 1 to 4 frequencies) over the course of a week. A clear periodic pattern emerges, with noticeable peaks during weekdays (Monday to Friday) and lower numbers on weekends (Saturday and Sunday). This suggests that user connectivity is higher on weekdays, likely due to work or school activities, while it decreases over the weekend. Each weekday shows a consistent pattern with two main peaks, which likely correspond to high-usage times, such as morning and evening commuting hours or regular work hours.

In terms of frequency usage, antennas supporting four frequencies (dark blue) attract the majority of users, indicating a preference or necessity for these higher-capacity antennas. Connections to four-frequency antennas remain the highest throughout the week, suggesting that these antennas are essential for handling most of the network's demand. In contrast, antennas with only one frequency (orange) see the fewest connections, suggesting they serve lower-traffic areas or are less optimal for managing high user volumes. Antennas with two and three frequencies (gray and light blue) handle moderate levels of traffic, with three-frequency antennas generally supporting more users than those with two frequencies.

The clear separation in user numbers across antennas with different frequencies suggests a tiered network structure. High-frequency antennas are likely positioned in areas with greater traffic

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

demand, while antennas with fewer frequencies may be placed in lower-demand areas, efficiently matching network resources to user distribution.

4.2. Peak Hour and Frequency Distribution

The double peaks observed in the traffic patterns (Figure 3) suggest that different geographical areas experience peak usage at distinct times. Residential areas likely see increased traffic in the evening, while commercial zones experience their peak during midday. Figure 5 illustrates this distribution of peak hours across network nodes, showing that most peak traffic occurs between 12:00 PM and midnight. This insight is critical for optimizing resource allocation in both RAN and vRAN environments, as it allows operators to anticipate high-traffic periods and adjust resources accordingly to ensure service quality during these peak times.

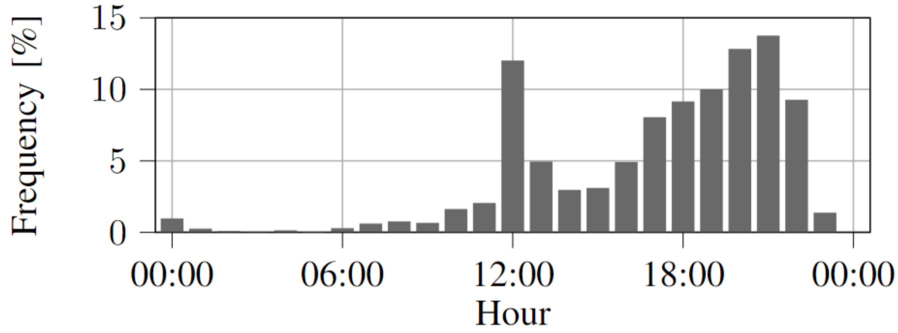

FIGURE 5: DISTRIBUTION OF THE PEAK HOUR FOR THE DIFFERENT NODES IN THE NETWORK.

The distribution of peak hours further highlights the importance of dynamic scheduling algorithms in vRAN systems, which can allocate resources like bandwidth and computing power in near-real-time, based on fluctuating demand. By managing intelligently managing resources during peak hours, operators can prevent network congestion, ensuring smooth user experience and optimized energy use.

4.3. User Traffic Patterns

Next, we assess the number of users connected to the network over the same period, as shown in Figure 6.a. While the number of users decreases on weekends, the overall traffic remains relatively stable, suggesting that individual users consume more data during non-working days. This is an important consideration for RAN and vRAN optimization strategies, as it highlights the need to balance user density with traffic volume when allocating resources.

In Figure 6.b, the variations in user numbers across different days of the week are presented. These fluctuations indicate that certain days, like Tuesdays and Wednesdays, are more stable, offering

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

reliable benchmarks for comparing the performance of different network configurations. Conversely, higher variability on Fridays suggests that real-time adjustments may be more necessary on such days to maintain network performance and resource efficiency.

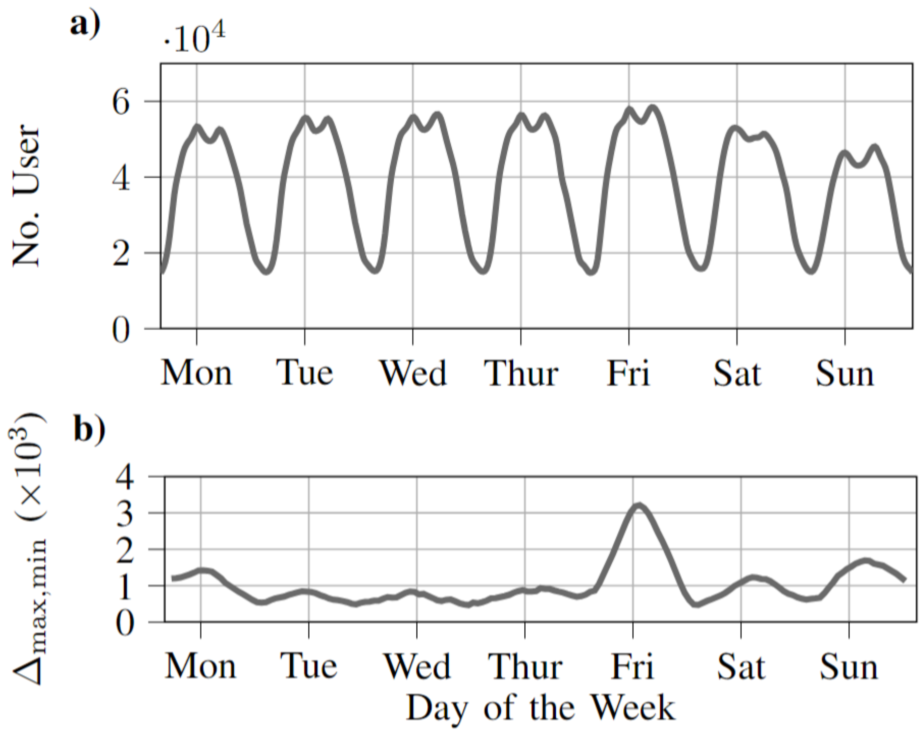

FIGURE 6: (A) AVERAGE NUMBER OF CONNECTED USERS PER HOUR DURING THE BASELINE WEEK. (B) HOURLY DIFFERENCE BETWEEN THE MAXIMUM AND MINIMUM NUMBER OF CONNECTED USERS, TO EVALUATE THE VARIABILITY ACROSS DIFFERENT WEEKS.

4.4. Optimal Network Configuration

To further assess the effectiveness of configuration 0, we compare it with alternative configurations over two comparable days (Table 1). As expected, configuration 0 achieves the highest total DL traffic, but the performance improvement over other configurations is modest, with less than a 4% difference in throughput. Nevertheless, configuration 0 remains the optimal choice for static, network-wide configuration due to its slightly better performance in managing overall traffic load.

This analysis demonstrates the importance of fine-tuning network configurations based on traffic patterns and user behavior. As networks continue to evolve towards vRAN architectures, the ability to implement dynamic, real-time adjustments will become increasingly crucial. While configuration 0 serves as a solid baseline, future developments in network resource allocation should focus on more adaptive and scalable solutions that can respond to the unique demands of modern mobile networks.

Configuration	DL Data [KB]	Difference [%]
0	$17.03 \cdot 10^{10}$	0.0
1	$16.71 \cdot 10^{10}$	-1.89
2	$16.97 \cdot 10^{10}$	-0.36
3	$16.49 \cdot 10^{10}$	-3.25

TABLE 1: BEST STATIC CONFIGURATION

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

5. Configuration Analysis

This section evaluates the impact of using different network configurations on resource allocation and overall network performance, focusing on how configuration changes influence traffic patterns and service quality in the network. In collaboration with the network operator, four distinct configurations were applied sequentially over consecutive weeks from August to October 2023. This period was chosen to avoid major holidays or network changes, allowing for a controlled environment to analyze the effects of each configuration on downlink (DL) traffic, a primary performance metric in Radio Access Networks (RANs) and virtualized RAN (vRAN) systems.

5.1. Temporal Analysis

The first step in analyzing the effectiveness of each configuration is to examine how traffic fluctuates throughout the day and how each configuration responds to these fluctuations. Figure 7.a illustrates the total traffic carried by the network for each configuration across the hours of the day. The data highlights two key traffic peaks—one around midday and another in the evening—demonstrating the dynamic nature of network demand. These peaks emphasize the need for intelligent resource allocation strategies in both RAN and vRAN to handle high traffic volumes efficiently.

Interestingly, the intersection of traffic curves for different configurations suggests that no single configuration consistently outperforms the others throughout the day. As seen in Figure 7.b, configuration 2 proves to be optimal during most hours, while configurations 1 and 3 show better performance during specific times, such as off-peak periods or times when traffic is decreasing. This observation aligns with the current trends in vRAN systems, where dynamic and adaptive configurations are crucial for optimizing network performance in response to real-time traffic conditions.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

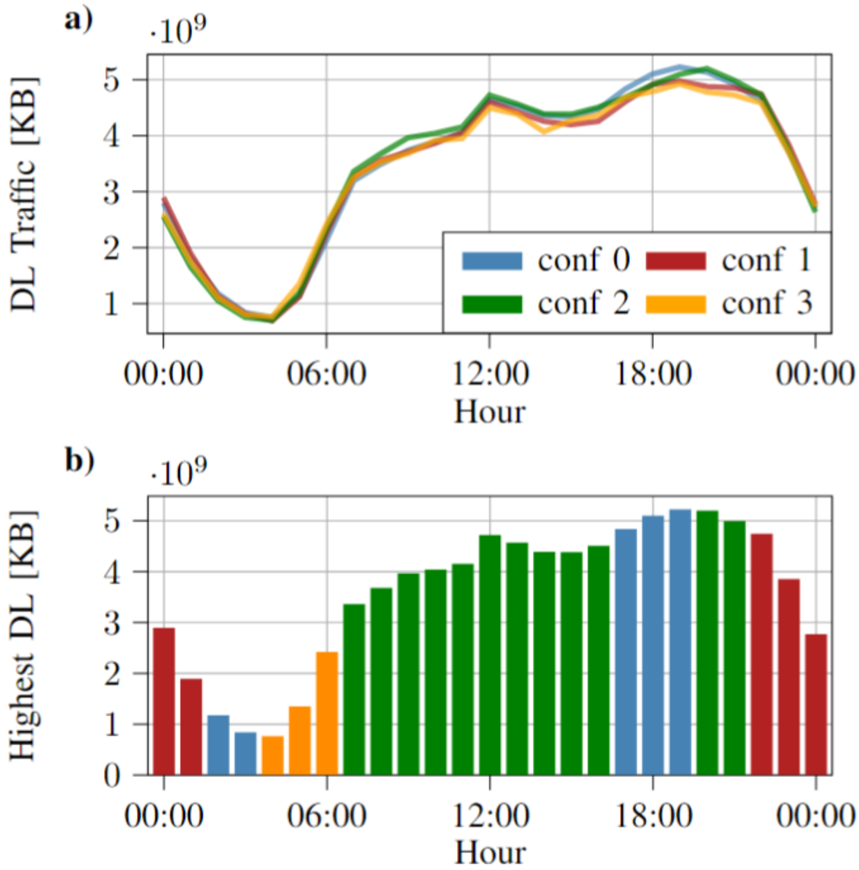

FIGURE 7: A) TOTAL TRAFFIC CARRIED BY THE NETWORK, UNDER THE DIFFERENT CONSIDERED CONFIGURATIONS, AND (B) BEST PERFORMING CONFIGURATION BY HOUR.

Since network traffic fluctuates throughout the day, we need to account how much each hourly-optimized configuration performs relative to a static configuration. This can be seen in Figure 8, which shows the relative traffic gain when using the optimal configuration for each hour compared to configuration 0 (i.e., the best configuration when considering only a single configuration over the whole day). In some hours, the gain exceeds 10%, while for most hours, if there is any gain, it remains below 5%. On average, switching to the optimal configuration each hour results in a 3% traffic increase, which is not significant enough to justify the added complexity of dynamic configuration changes, as this gain is like the difference between the best and worst static configurations for the network.

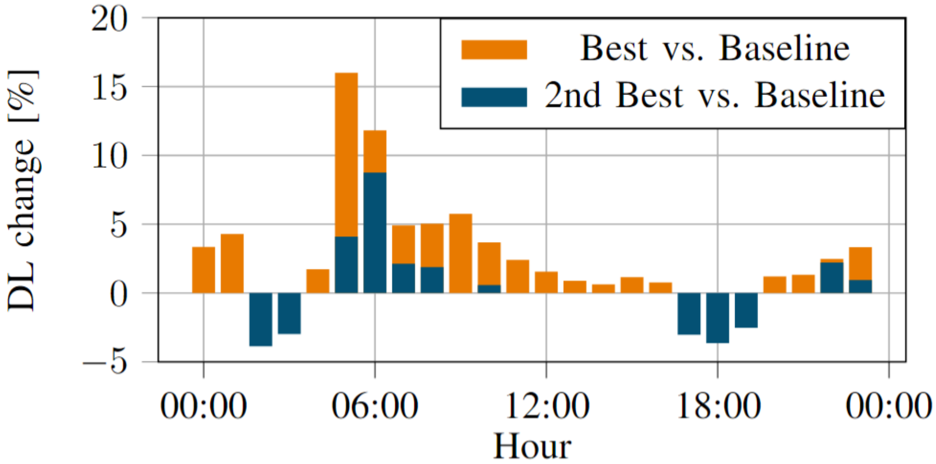

FIGURE 8: TRAFFIC GAIN BY HOUR, WHEN USING THE BEST (AND SECOND BEST) CONFIGURATION FOR EACH HOUR, WITH RESPECT TO USING ALWAYS THE CONFIGURATION 0.

5.2. Spatial Analysis

We now analyze the spatial scale of reconfiguration. It might seem reasonable to assume that the network conditions are similar considering close areas, and, as such, that the same configuration would work for nearby cells.

To investigate this, we randomly selected a set of cells close to each other, though connected different nodes. We then analyzed the optimal configuration for each cell, allowing reconfigurations each 6 hours. We consistently found that the best configuration differed from cell to cell, as shown in Figure 9, where each color represents a different optimal configuration for a cell over 6-hour periods throughout the reference day. Moreover, the distribution of the best configurations for randomly selected cells is mostly uniform; in other words, there are roughly the same number of cells where each configuration is the optimal one, meaning that no single configuration consistently outperforms the others across nearby cells. This suggests that allowing different configurations for different geographical areas allows the network to reach a higher gain in terms of carried traffic. When selecting the best configuration for each one of the 3 municipalities, a traffic increase of 1% is observed, while when selecting the best configuration for each neighborhood, the gain rises to 8.8%, 10.6% for node, reaching 17.5% when selecting the best configuration for each cell. This shows that there is considerable potential when allowing different configurations for different areas of the network, and that the smaller the area, the higher the gain.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

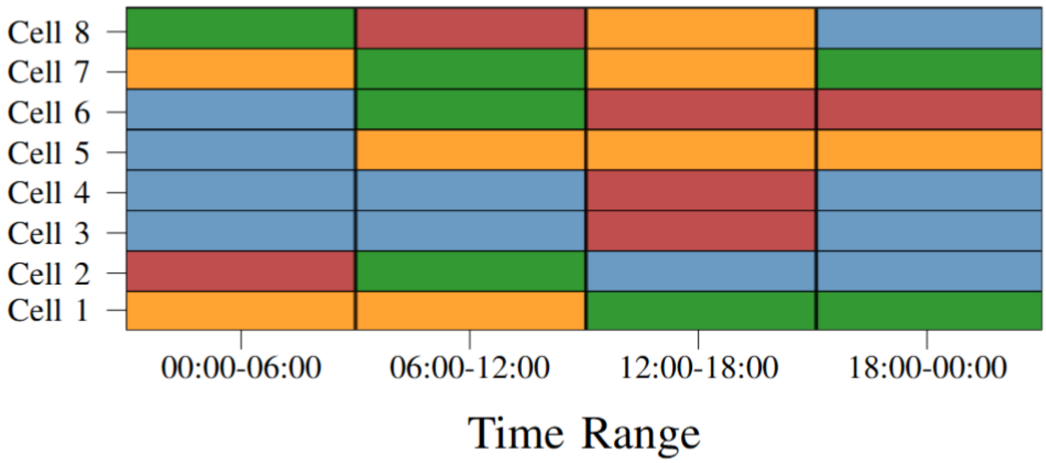

FIGURE 9: BEST CONFIGURATION FOR DIFFERENT NEIGHBOUR CELLS, ON THE REFERENCE DAY (BLUE: CONF0, RED: CONF1, GREEN: CONF2, AND YELLOW: CONF3)

5.3. Frequency Analysis

This figure illustrates the Downlink (DL) Traffic Gain percentage associated with antennas operating with different numbers of frequency bands, ranging from 1 to 4 frequencies. Each bar in the figure represents the average traffic gain achieved by antennas with a specific number of frequencies, while the error bars indicate the variability or range within which these gains are observed.

The traffic gain tends to be highest for antennas with only 1 frequency, reaching approximately 20% or more on average, with significant variability, as indicated by the large error bars. Antennas with 2 frequencies show a slightly lower average traffic gain than those with only 1 frequency, but they still achieve a notable increase in DL traffic, with some degree of variability. As the number of frequencies increases to 3 and 4, the traffic gain percentage continues to decrease. Antennas with 4 frequencies show the lowest traffic gain, around 10% or less, with a relatively smaller range of variability.

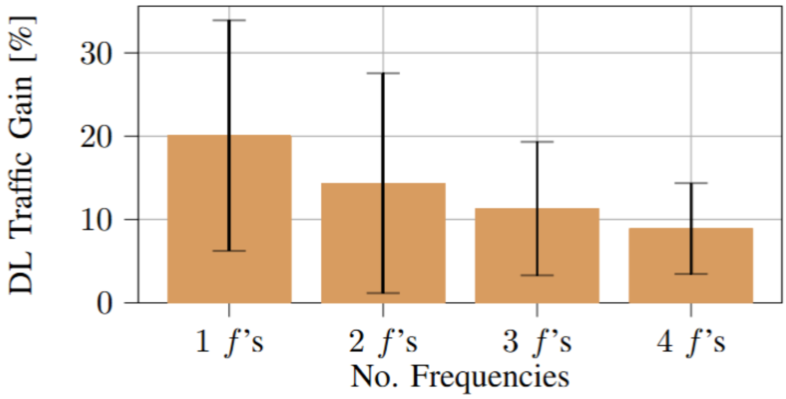

FIGURE 10: TRAFFIC GAINS AS A FUNCTION OF THE NUMBER OF FREQUENCIES.

This trend suggests that antennas with fewer frequency bands, such as those with 1 or 2 frequencies, experience higher relative gains in DL traffic. This may indicate that lower-frequency antennas are more sensitive to flexible configuration changes, possibly because they are in areas with lower baseline traffic demand. Adjustments in configuration for these antennas result in more substantial gains relative to their normal traffic levels. In contrast, antennas with more frequencies, such as those with 3 or 4, may already be handling high traffic volumes, so their relative traffic gains from configuration adjustments are smaller. These higher-frequency antennas may be positioned in high-demand areas where the impact of configuration changes is less pronounced compared to the total traffic handled.

The error bars reveal wide variability in traffic gains for antennas with 1 and 2 frequencies, suggesting that these gains are less consistent and could be influenced by specific contextual factors such as location, time, or network demand. For antennas with 3 or 4 frequencies, the gains are more stable, as shown by the smaller error bars, indicating that traffic increases are more predictable but relatively lower.

5.4. Spatio-temporal Analysis

Now, we evaluate the potential gain in traffic achieved by dynamically adjusting the configuration for both different spatial and temporal scales, being aware that such reconfigurations increase system complexity and management costs. To this extent, we admit, for each dimension of the analysis, different values of the corresponding granularity, and then analyze all the possible intersections, corresponding to different trade-offs. For the time scale, we consider configuration changes that may happen every 1, 2, 4, 8, 12, or 24 hours. On the other hand, for the spatial scale, we consider configurations that may change for each cell, node, neighborhood, municipality, or that must be the same over the whole region. For each point considered in the resulting matrix, we consistently select the optimal configuration (i.e., the one resulting in the highest traffic) for the chosen time and spatial resolution over the study period. The results are presented in Figure 10 which shows the average percentage of traffic gain when selecting the optimal configuration at each spatio-temporal granularity. Each curve represents a different spatial resolution, with the x-axis indicating

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

the time granularity and the y-axis showing the percentage of traffic gain for the corresponding setting.

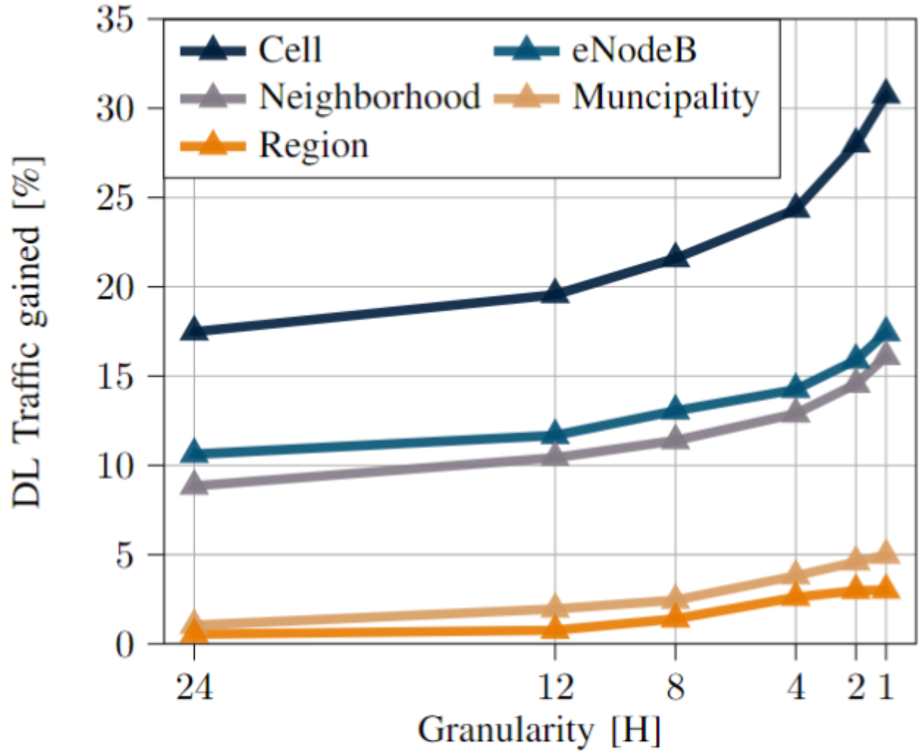

FIGURE 11: AVERAGE PERCENTAGE OF TRAFFIC INCREASE, ALLOWING CONFIGURATION CHANGES WITH DIFFERENT TIME GRANULARITIES (THE X-VALUE INDICATES THE MINIMUM INTERVAL OVER WHICH A CONFIGURATION CANNOT CHANGE), AND AT DIFFERENT GEOGRAPHICAL SCALES (DIFFERENT CURVES).

The simplest scenario is when a single configuration is applied across the entire network, and it is not allowed to change over the 24h period. This scenario corresponds to the left point of the bottom curve (orange), in Figure 10. Here the gain is marginal (0.5%), as it is only achieved by changing configuration once per day, with respect to having a fixed configuration always. Remaining on the same curve, but moving right, we consider scenarios in which a configuration is selected for the whole region, but it is allowed to change more frequently, up to each hour (i.e., right end of the lower curve - orange - also corresponding to the scenario analyzed in Figure 7 and 8). The average gain for this scenario is 3.3%. The curve immediately above (light brown) corresponds to the scenario in which different configurations can be selected for different municipalities. Here the improvement is very marginal, with respect to the regional scenario (i.e., from 1%, allowing a configuration to change every 24h, to 4.9% allowing a configuration change per hour), as each area sharing a configuration (i.e., municipality) still includes high heterogeneity. On the other hand, we can see a significant improvement when configurations are selected at a neighborhood and eNodeB level (grey and light blue curves). When selecting the optimal configuration for the entire day in these scenarios, average performance increases by 8.8% and 10.4%, respectively. These gains rise to approximately 16.1% and 17.4% when configurations are adjusted hourly.

Finally, allowing configuration changes at cell level results in substantial throughput improvements (dark blue curve). In a less complex scenario, where the configuration is fixed for each cell over 24 hours (left end of the curve), we observe a 17.5% traffic gain. In the most dynamic scenario, with hourly adjustments at the cell level (right end of the curve), traffic volume increases by 30.7%.

As expected, both increases in flexibility (i.e., in time and in space) result in significant performance increases for the network. At the same time, allowing for different configurations in different neighborhoods, and for different cells, increases the network performance much more than changing the configuration more often. Finally, the gain brought by the time granularity (i.e., the increase in the same curve going from left - 24h - to right - 1h) increases a lot when enabling higher spatial granularity, passing from about 3%, when allowing region-level granularity, to almost 15%, when allowing cell-level granularity. The same is also true the other way around: the gain brought by space granularity (i.e., different curves for the same x-value) increases then enabling higher temporal granularity, passing from about 17%, when allowing 24h time granularity, to about 30%, when allowing 1h time granularity. This means that not only does each dimension of configuration flexibility increase the network performance, but also that each one helps the other to increase even more the network performance (i.e., the two contributions do not sum up linearly, but more than linearly). This study highlights the importance of flexible cell configurations in mobile networks, particularly in virtualized Radio Access Networks (vRAN), where dynamic resource management is essential. By allowing temporal and spatial adjustments to configurations, networks can handle significantly more traffic—up to 30% more with per-cell flexibility—while optimizing resource allocation. Spatial flexibility, especially at finer geographic scales, proves to be the most effective for improving network performance.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

6. Reconfiguration Problem

To tackle the problem of finding an equilibrium between reconfiguration costs at various spatiotemporal scales and the resulting traffic gains, we can define a function that balances these two aspects. This function will weigh the cost of reconfiguring at different spatial scales (e.g., neighborhoods, nodes, regions) and temporal scales (e.g., every hour, every minute) against the traffic gains achieved.

To define the problem mathematically, let's consider the following variables and objective:

1. $\boldsymbol{G\_s}$: The traffic gain (in percentage) achieved by reconfiguring at spatiotemporal scale s.
2. $\boldsymbol{C\_s}$: The cost of reconfiguration at spatiotemporal scale s.
3. $\boldsymbol{w_s}$: The weight or priority assigned to spatiotemporal scale s based on its importance.
4. $\boldsymbol{\alpha}$: The importance given to the gain with respect to the cost.
5. $\boldsymbol{x_s}$: A binary decision variable, where:
    - $x_s = 1$ if reconfiguration at scale s is selected,
    - $x_s = 0$ otherwise.

We want to **maximize the net benefit**, which is the weighted sum of traffic gains minus the reconfiguration costs across all scales. This can be expressed as:

$$Maximize \sum_s x_s(\alpha w_s \cdot G_s - (1-\alpha)C_s)$$

Constraints:

1. Non-negativity: All gains and costs are non-negative:

$$G_s \geq 0, C_s \geq 0, \forall s$$

2. Budget constraint: If there is a total budget B for reconfiguration, we may impose the following constraint:

$$\sum_s x_s \cdot C_s \leq B$$

The optimization problem can then be formally defined as:

$$\begin{aligned} &Maximize \ \textstyle\sum_s x_s (w_s \cdot G_s - C_s) \\ &subject\ to \ \textstyle\sum_s x_s \cdot C_s \leq B \\ &\qquad G_s \geq 0, C_s \geq 0, \forall s \\ &\qquad x_s \in \{0,1\}, \forall s \end{aligned}$$

6.1. Modelling the Cost as a function of Scale and Frequency

While it is straightforward to define traffic gains, as analyzed in the previous section, defining reconfiguration costs requires a more detailed breakdown. Costs can be categorized into direct operational costs, such as power consumption, equipment wear, and personnel or automation expenses, which increase with frequent reconfigurations. Additionally, network performance and quality impacts, such as latency spikes or service interruptions during transitions, contribute to costs by affecting customer experience and SLAs. There is also a revenue opportunity cost, where reconfiguration may temporarily reduce traffic capacity, leading to potential revenue loss, especially during peak times. Finally, a data-driven approach, using historical data on past reconfigurations (like energy consumption spikes or maintenance needs), can help estimate typical costs across different scales, providing a more realistic picture of reconfiguration expenses.

We define the cost C_s at scale s as:

$$C_s = \alpha E_s + \beta T_s + \gamma R_s$$

- E_s : **Energy cost** at scale s (e.g., based on power usage).
- T_s: **Transition disruption cost**, representing performance degradation or quality impact.
- R_s: **Revenue opportunity cost**, estimating lost traffic or potential revenue.

Here's a hypothetical breakdown for a reconfiguration cost at a neighbourhood-hour scale:

- **Energy Cost**: If reconfiguration at this scale increases energy consumption by 5 kWh and electricity costs 0.10$ per kWh, then:

$$E_s = 5\ kWh \times 0.10\ USD/kWh = 0.5\ USD$$

- **Transition Disruption Cost**: If service degradation leads to a 1% increase in latency complaints or SLA penalties, and each complaint/penalty costs 2$:

$$T_s = 0.01 \times \text{total users affected} \times 2\ \$$$

- **Revenue Opportunity Cost**: If reconfiguration interrupts 2 minutes of peak-hour traffic and leads to an estimated 0.1 TB of lost data transfer, with each TB generating 50$ in revenue:

$$R_s = 0.1\, TB \times 50\, USD/TB = 5\, USD$$

Then, combining these:

$$C_s = \alpha \cdot 0.5 + \beta \cdot T_s + \gamma \cdot 5$$

Where α, β, and γ needs to be adjusted to reflect the relative importance of each cost component.

Interval (Hours)	Cell	eNodeB	Neighborhood	Municipality	Region
1	966909.15	99677.01	57562.39	1857.67	625.91
2	483798.56	49872.13	28811.00	928.80	312.94
4	242133.72	24957.93	14413.36	464.46	156.51
8	121099.49	12475.50	7204.34	232.21	78.25
12	81219.28	8364.67	4845.65	154.91	52.22

TABLE 2: AVERAGE NUMBER OF USERS AFFECTED BY THE RECONFIGURATIONS AT EACH SCALE.

We conducted a hyperparameter optimization to determine the optimal configuration of network reconfiguration parameters for various budget and preference settings. Our goal was to identify the best trade-offs between maximizing traffic gain and minimizing user disruptions across different combinations of budget constraints and alpha values (which control the balance between prioritizing traffic gain and minimizing disruptions). We varied the budget, which represents the total allowable number of users impacted by reconfigurations, across multiple values (e.g., 50,000, 100,000, 150,000 users) and explored a range of alpha values (0 to 1), where alpha = 0 focused solely on minimizing user disruptions, and alpha = 1 prioritized traffic gain. For each combination of budget and alpha, we normalized the traffic gains and user impacts to make them comparable and set up an optimization problem using PuLP to maximize a custom net benefit function that incorporated both traffic gains and disruptions based on the selected alpha. We solved this optimization for each (budget, alpha) pair to find the optimal configuration of time intervals (H) and spatial scales (e.g., Cell, eNodeB) for network reconfigurations. After obtaining the optimal configurations, we visualized the results as a matrix plot, where each cell in the matrix represents an optimal configuration for a specific (budget, alpha) pair. The color of each cell corresponds to the time interval (ranging from red for 1-hour intervals to blue for 24-hour intervals), while the label within each cell displays the configuration's time interval, spatial scale, and resulting traffic gain.

This matrix plot in Figure 12 provides insights into the optimal network reconfiguration configurations for various combinations of budget and importance weight (alpha). Each cell represents the optimal configuration in terms of time interval (H) and spatial scale (e.g., Region, Municipality, Neighborhood, Cell) for a given budget and alpha. The color of each cell reflects the time interval (ranging from red for 1-hour intervals to blue for 24-hour intervals), and the label inside each cell shows the configuration in [H, Scale] format.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

One of the primary observations is the influence of alpha (importance weight) on configuration selection. At low alpha values (e.g., 0.01, 0.1), where minimizing disruptions (user impact) is prioritized over maximizing traffic gain, the optimal configurations tend toward longer intervals (e.g., 24 hours) and larger spatial scales (e.g., Region or Municipality). These configurations affect fewer users because they require fewer reconfigurations over a broader area, aligning with the objective to minimize disruptions. However, as alpha increases, shifting the focus more towards traffic gain, the configurations transition to shorter intervals (e.g., 2-4 hours) and smaller spatial scales (e.g., Neighborhood, Cell). This shift allows for more granular and frequent adjustments, enhancing the network's adaptability and potential for traffic gain, albeit with a higher impact on user experience.

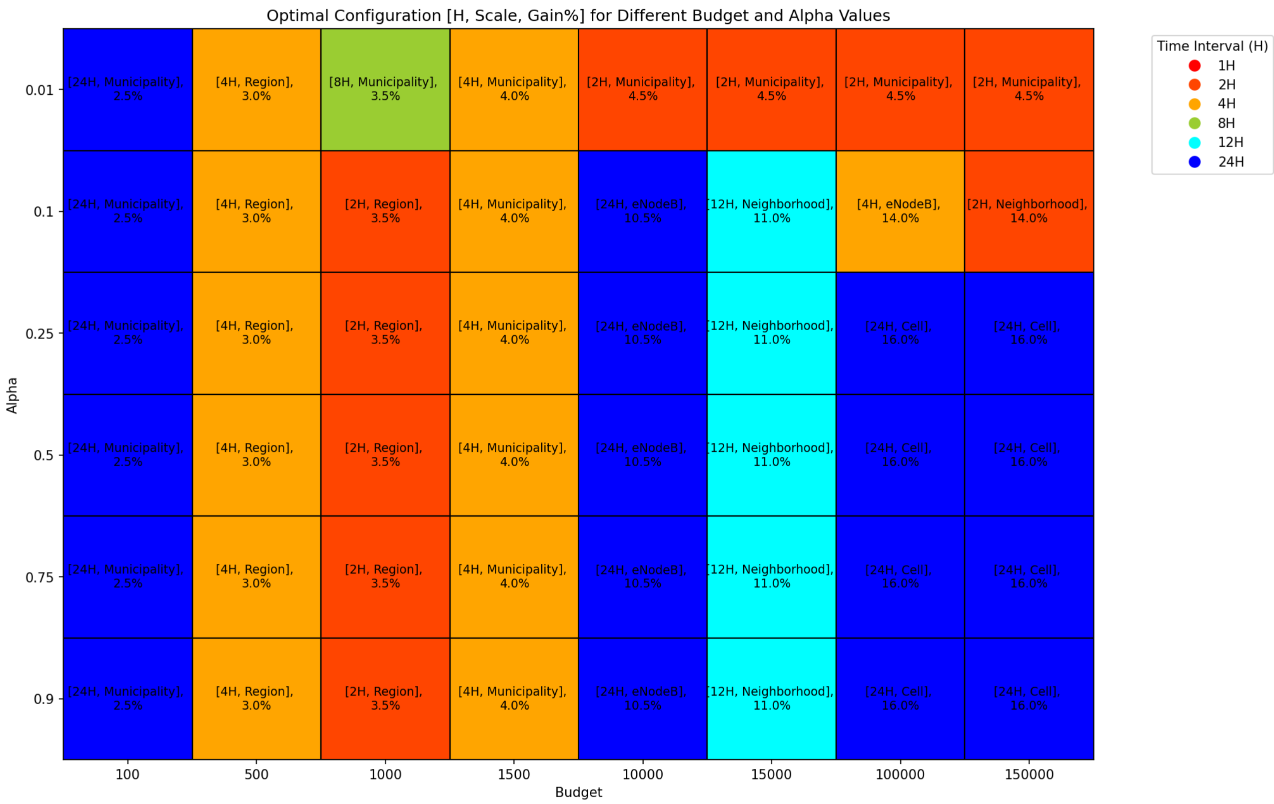

FIGURE 12: OPTIMAL CONFIGURATION AS A FUNCTION OF THE BUDGET AND THE IMPORTANCE TO GAIN AND BUDGET.

The budget also has a significant impact on configuration choices. With very low budgets (e.g., 100), the configurations remain at 24-hour intervals at larger spatial scales (Municipality or Region), regardless of alpha. This constraint indicates that with limited resources, the network can only afford minimal reconfigurations to avoid impacting too many users. As the budget increases (e.g., to 500 or more), there is a clear shift toward shorter time intervals (e.g., 4 hours, 2 hours) at larger alphas, suggesting that, given sufficient budget, the network can afford more frequent reconfigurations to optimize traffic gains while managing user impact.

For intermediate budget levels (e.g., 1000 to 10000), we observe mixed configurations that balance both moderate time intervals (e.g., 4-12 hours) and intermediate spatial scales (e.g., Municipality,

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

Neighbourhood). This suggests that these configurations are optimal for balancing traffic gains and disruptions within reasonable budget constraints. At higher alpha values and higher budgets, configurations increasingly favour frequent reconfigurations (e.g., 2-hour intervals) at finer spatial scales (e.g., Neighbourhood or Cell). This configuration is ideal for maximizing network responsiveness and traffic handling capacity, suitable for scenarios where traffic gain is the main priority.

A preference for larger scales emerges when alpha is low, with optimal configurations favoring larger spatial scales like Region or Municipality, especially when the budget is limited. This suggests that when user disruption is prioritized, reconfiguring larger areas less frequently is more efficient in reducing the impact on individual users.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

7. Summary and Conclusions

The configuration of cells in mobile access networks is a critical aspect for network operators, not only impacting network performance but also directly influencing operational and management costs. As mobile networks evolve towards more complex architectures such as virtualized Radio Access Networks (vRAN), the ability to dynamically and efficiently configure cells takes on even greater significance. In vRAN environments, where virtualized instances of network functions must be managed in real-time, the flexibility of network configuration becomes a crucial tool for optimizing resource allocation, ensuring performance isolation, and maintaining service quality.

In this study, we first investigated how selecting different cell configurations affects a network's primary Key Performance Indicator (KPI), specifically the volume of downlink (DL) traffic carried by the network. Our results show that while static configurations provide a solid foundation, significant traffic gains can be achieved by introducing flexibility in the configuration process. We analyzed two types of flexibility: temporal (adjusting configurations at different times of the day) and spatial (applying different configurations across regions, municipalities, neighborhoods, or cells). The findings reveal that both temporal and spatial flexibility lead to higher carried traffic, but spatial flexibility is far more effective, especially at finer geographic scales.

Allowing different configurations for distinct neighborhoods led to traffic gains of over 15%, while enabling per-cell configuration adjustments resulted in increases of over 30%. These results underscore the importance of granular control over network configurations, particularly in vRAN systems where the precise allocation of resources can significantly enhance overall network performance and efficiency.

In the context of vRAN, where virtualized network functions are dynamically managed and resources can be reallocated on-demand, these findings are particularly relevant. The ability to adapt configurations at the cell level not only improves traffic handling but also aligns with the core objectives of vRAN: improving scalability, resource efficiency, and performance isolation. The potential to increase carried traffic by over 30% through per-cell configuration flexibility is a testament to how vRAN's architectural capabilities can be leveraged for superior network optimization. Moreover, this flexibility allows operators to respond to fluctuating traffic demands in real-time, ensuring optimal use of computational and network resources while reducing the risk of bottlenecks or inefficiencies.

Looking ahead, future work should focus on deepening our understanding of the correlation between network conditions and the optimal configuration for those conditions. This could involve developing predictive models that use real-time data to anticipate traffic patterns and adjust configurations dynamically, further enhancing the efficiency of resource allocation in vRAN systems. Additionally, it will be important to evaluate how different configuration strategies affect other KPIs, such as fairness in traffic distribution across frequency bands and the rate of user drops. By expanding

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

the scope of the analysis to include multiple performance metrics, we can better assess the holistic impact of configuration strategies on network performance and user experience.

In conclusion, this deliverable highlights the critical role that flexible network configuration plays in optimizing performance, particularly in the evolving landscape of vRAN. As operators continue to transition to virtualized architecture, the ability to adapt configurations both temporally and spatially will be key to maintaining high service quality, improving traffic management, and reducing operational costs. These results provide a strong foundation for future work in developing more adaptive, real-time configuration systems that can fully harness the potential of vRAN to meet the growing demands of modern mobile networks.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA
MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL
SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D

References

[SMM2020] J. Shodamola, U. Masood, M. Manalastas, A. Imran. ``A machine learning based framework for KPI maximization in emerging networks using mobility parameters.'' In 2020 IEEE International Black Sea Conference on Communications and Networking, IEEE, Tbilisi, Georgia, 2020.

[AH2015] S. Almowuena and M. Hefeeda. ``Dynamic configuration of single frequency networks in mobile streaming systems.'' In Proceedings of the 6th ACM Multimedia Systems Conference (MMSys '15), ACM, New York, NY, USA, 153–164, 2015.

[YDA2020] A. Yazar, S. Dogan-Tusha, H. Arslan. ``6G vision: An ultra-flexible perspective.'' ITU Journal on Future and Evolving Technologies, 1(1), 121-140, 2020.

[MGF2018] C. Marquez, M. Gramaglia, M. Fiore, A. Banchs, X. Costa-Perez, ``How Should I Slice My Network? A Multi-Service Empirical Evaluation of Resource Sharing Efficiency,'', ACM MobiCom, India, 2018.

[MAG2020] J. Moysen, F. Ahmed, M. García-Lozano, J. Niemelä. ``Big data-driven automated anomaly detection and performance forecasting in mobile networks.'' In 2020 IEEE Globecom Workshops (pp. 1-5), 2020

[VSP2023] I. Vilà, O. Sallent, J. Pérez-Romero. ``On the design of a network digital twin for the radio access network in 5g and beyond.'' Sensors, 23(3), 1197, 2023.

[PJW2024] M. Panek, I. Jablonski, M. Wozniak. ``Modeling configuration-performance relation in a mobile network: a data-driven approach.'' arXiv preprint arXiv:2407.06702, 2024.

[AFI2018] A. Asghar, H. Farooq, A. Imran. ``Concurrent optimization of coverage, capacity, and load balance in HetNets through soft and hard cell association parameters.'' IEEE Transactions on Vehicular Technology, 67(9), 8781-8795, 2018.

[HN2015] R. D. Hegazy, O. A. Nasr. ``A user behavior based handover optimization algorithm for LTE networks.'' In 2015 IEEE Wireless Communications and Networking Conference (WCNC) (pp. 1255-1260), IEEE, Istanbul, Turkey, 2015.

Financiado por la Unión Europea NextGenerationEU
GOBIERNO DE ESPAÑA MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE TELECOMUNICACIONES E INFRAESTRUCTURAS DIGITALES
Plan de Recuperación, Transformación y Resiliencia
UNICO I+D