UNICO I+D Project
Grant No.

6G-SORUS-RIS

SORUS-RIS-A2.3-E2 (E13)

Abstract

This Deliverable explores the integration of Reconfigurable Intelligent Surfaces (RIS) into campus-wide virtualized Radio Access Networks (vRAN), emphasizing the need for scalable control, latency-aware signalling, and resource-efficient architectures to meet the demands of emerging 6G services. We analyze the scalability of mechanically reconfigurable RIS, identifying computational and signalling strategies for real-time management of large-scale deployments. Additionally, we formulate an energy-aware AP ON/OFF optimization problem that leverages demand forecasting and pre-configured RIS steering to achieve significant energy savings without compromising coverage. By combining predictive intelligence in vRAN with RIS control, this work provides a blueprint for intelligent, context-aware resource management. The findings underscore the transformative potential of coordinated RIS and vRAN operation, paving the way for sustainable, flexible, and service-oriented next-generation wireless 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

Document properties

Document number	SORUS-RIS-A2.3-E2 (E13)
Document title	Algorithms for RIS operation: results
Document responsible	Pablo Serrano
Document editor	Eva Rajo
Editorial team	Eva Rajo, Pablo Serrano, Albert Banchs
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	Juan Manuel Montes, Marco Gramaglia

Disclaimer

This document has been produced in the context of the 6G-SORUS-RIS 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

List of Tables

List of Acronyms

Beyond 5G : B5G

Virtualized Radio Access Networks: vRAN

Reconfigurable Intelligent Surfaces: RIS

Unnamed Aerial Vehicles: UAVs

Central Units: CUs

Distributed Units: DUs

Radio Units: RUs

Radio Network Controller: RNC

New Radio: NR

Cloud RAN: CRAN

Machine Learning: ML

Artificial Intelligence: AI

Real-Time RAN: rtran

RAN Intelligent Controller: RIC

Electromagnetic: EM

Base Station: BS

Access Point: AP

Non Line of Sight: NLoS

Multiple-input Multiple-output: MIMO

Reinforcement Learning: RL

Key Performance Indicators: KPI

non-terrestrial network: NTN

Integrated Sensing and Communications: ISAC

signal-to-noise ratio: SNR

Quality of Service: QoS

Hierarchical Reinforcement Learning: HRL

Open RAN: ORAN

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 las complejidades de la gestión de RIS reconfigurables mecánicamente en redes virtualizadas, como las redes de acceso radioeléctrico virtualizadas (vRAN). Analizamos los retos de escalabilidad que surgen al implementar módulos RIS en entornos heterogéneos, destacando los principales cuellos de botella, como la asignación de recursos, la escalabilidad del plano de control y los protocolos de señalización resistentes a la latencia. Partiendo de esta base, exploramos estrategias computacionales y de arquitectura para permitir el control RIS en tiempo real dentro de vRAN. A medida que las redes 6G evolucionan para admitir diversas clases de servicios, incluidos eMBB, mMTC y URLLC, examinamos cómo los marcos predictivos pueden ampliar las vRAN habilitadas para RIS más allá de los objetivos basados en el rendimiento para cumplir con los estrictos requisitos de latencia, fiabilidad y aislamiento del servicio. Las siguientes secciones presentan estos análisis en detalle y proponen modelos y mecanismos para unificar la gestión escalable de RIS con el control inteligente de vRAN.

Por último, formulamos el problema de la eficiencia energética en un escenario con puntos de acceso y RIS, y lo planteamos como un modelo de optimización que tiene como objetivo minimizar el consumo energético de la red al tiempo que se garantiza una cobertura continua. En esta formulación, los AP permanecen apagados a menos que la demanda prevista de los usuarios requiera explícitamente su activación. Para cubrir las lagunas de cobertura durante estos periodos de apagado, se utilizan mosaicos RIS preconfigurados para redirigir las señales de los AP vecinos. El modelo captura restricciones clave como la precisión de la predicción de la demanda de los usuarios, los retrasos en la reconfiguración del RIS y las ganancias de dirección, equilibrando así el ahorro de energía con la necesidad de mantener una conectividad sin interrupciones.

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 first delves into the complexities of managing mechanically reconfigurable RIS in virtualized networks like virtualized Radio Access Networks (vRAN). We analyze the scalability challenges that arise when deploying RIS tiles across heterogeneous environments, highlighting key bottlenecks such as resource allocation, control plane scalability, and latency-resilient signalling protocols. Building on this, we explore computational and architectural strategies to enable real-time RIS control within vRAN. As 6G networks evolve to support diverse service classes, including eMBB, mMTC, and URLLC, we then examine how predictive, risk-aware, and slice-aware frameworks can extend RIS-enabled vRAN beyond throughput-driven objectives to meet stringent latency, reliability, and service isolation requirements. The following sections present these analyses in detail and propose models and mechanisms to unify scalable RIS management with intelligent vRAN control.

Finally, we formulate the problem of energy efficiency in and scenario with access points and RIS, we formulate it as a mixed-integer optimization model that aims to minimize network energy consumption while ensuring continuous coverage. In this formulation, APs remain switched off unless forecasted user demand explicitly requires their activation. To bridge coverage gaps during these off periods, pre-configured RIS tiles are used to redirect signals from neighboring APs. The model captures key constraints such as user demand prediction accuracy, RIS reconfiguration delays, and steering gains, thereby balancing energy savings with the need to maintain seamless connectivity.

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

1. Introduction

The evolution toward 6G and beyond is marked not only by the pursuit of higher data rates, reduced latency, and enhanced reliability, but also by the pressing imperative of energy efficiency. As networks densify and more devices come online, the energy footprint of the Radio Access Network (RAN) grows rapidly, prompting the exploration of fundamentally new paradigms in how wireless systems interact with the environment. Among these paradigms, Reconfigurable Intelligent Surfaces (RIS) have emerged as a groundbreaking solution. By enabling programmable control over electromagnetic wave propagation, RIS can shape the wireless channel itself, providing opportunities to redirect, focus, or suppress signals in a passive or near-passive manner. Their potential to significantly reduce the reliance on active relaying or power-hungry beamforming makes them ideal candidates for the realization of energy-efficient intelligent environments [1].

This deliverable is devoted to exploring the algorithmic challenges and strategies for maximizing energy efficiency in RIS-enabled networks, particularly in scenarios where RIS are not only passive reflectors but intelligent, context-aware agents capable of making informed decisions. Two complementary perspectives guide our investigation. First, we examine mechanical RIS (M-RIS), an emerging class of reconfigurable surfaces that achieve electromagnetic control through physical deformation or mechanical actuation rather than active electronic switching. Examples include kirigami-based surfaces, MEMS-tuned reflectarrays, or soft-material metasurfaces whose geometry can be changed with minimal power consumption. These systems promise ultra-low steady-state power and can potentially operate in a battery-less fashion. However, they also pose unique algorithmic challenges, as their reconfiguration latency (ranging from tens of milliseconds to seconds) and actuation fatigue introduce constraints that are fundamentally different from those of electronically controlled RIS. The need for proactive, long-horizon planning and wear-aware optimisation compels a rethinking of traditional RIS algorithms. In this context, we investigate how motion-aware scheduling, long-term channel prediction, and lightweight heuristics can be adapted or reinvented to work within the unique capabilities and limitations of mechanical RIS [2].

The second perspective we take is more integrative and forward-looking: we explore the joint optimisation of RIS with other network technologies to build energy-aware, cross-domain decision frameworks. In particular, we focus on two promising vectors of integration: (i) virtualised RAN (vRAN) architectures and (ii) Unmanned Aerial Vehicles (UAVs). Virtualised RAN separates network functions into cloud- or edge-hosted containers, enabling global views of network state, traffic forecasts, and user trajectories. This centralised intelligence makes it possible to treat RIS as a dynamically coordinated component in the RAN ecosystem. By leveraging vRAN control and telemetry such as, buffer states, handover maps, or gNB-resident learning models, RIS decisions can be informed by future context, such as expected user movement or traffic bursts. We explore how this foresight enables proactive surface programming and energy-aware trade-offs between computation and propagation control. Moreover, the bidirectional nature of the integration allows RIS telemetry to inform vRAN scheduling, closing the loop in a symbiotic system [3].

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

Similarly, the synergy between RIS and UAVs opens new opportunities for energy-adaptive aerial coverage. UAV-mounted RIS can act as mobile reflectors or passive relays, dynamically repositioned based on spatiotemporal demand or energy constraints. Conversely, ground-deployed RIS can assist UAV communications by creating virtual line-of-sight paths or offloading some of the channel manipulation tasks, thereby reducing UAV transmission power. The joint optimisation of RIS pose, UAV path planning, and beamforming strategy leads to a complex but rewarding design space where decisions across the air and ground domains must be harmonised. From an algorithmic standpoint, this integration demands methods that are robust to mobility, efficient under constrained computation, and capable of real-time adaptation to fast-changing topology and channel conditions [4].

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. Mechanical RIS for Energy Efficiency

Reconfigurable Intelligent Surfaces (RIS) have emerged as a central pillar in the roadmap toward energy-efficient 6G networks. By shaping the wireless channel through programmable meta-atoms, RIS offer a passive or near-passive alternative to traditional relays, enabling directional control of radio waves without the need for high-power RF chains. While most RIS research and prototypes have focused on electronic tuning using CMOS switches, PIN diodes, or varactors, there is growing recognition of the transformative potential of M-RIS [2]. Unlike their electronically tuned counterparts, M-RIS leverage physical deformation or translation to alter electromagnetic response, using actuation mechanisms such as cam-driven reflectarrays, kirigami-based folding metasurfaces, and soft-materials like liquid crystal elastomers (LCEs). These solutions consume power only during reconfiguration and exhibit non-volatile mechanical states, opening the door to ultra-low-power and even battery-less implementations that can remain static for hours or days without drawing any energy [5].

cent research has shown that mechanical RIS can achieve advanced beamforming capabilities while maintaining negligible steady-state power consumption. Unlike electronically controlled RIS, mechanically actuated surfaces, such as MEMS-based reflectarrays, kirigami-inspired metasurfaces, and soft-material structures, leverage physical reconfiguration to manipulate electromagnetic wavefronts without the need for continuous biasing [6]. These systems inherently adopt an event-driven actuation model, where energy is only expended during configuration changes, enabling passive state retention and cutting steady-state power requirements by several orders of magnitude compared to PIN-diode-based RIS. Such approaches have demonstrated efficient control over reflection phase and beam direction, with experimental works reporting low-latency actuation in the millisecond-to-second range and energy savings exceeding 30–50 dB relative to conventional electronic tuning [7].

From an algorithmic standpoint, M-RIS introduce new challenges that demand a departure from standard RIS control models. Traditional beamforming approaches rely on millisecond-scale control loops that fine-tune electronically controlled metasurfaces using high-rate CSI updates. In contrast, mechanical RIS operate with reconfiguration timescales from tens of milliseconds up to several seconds, depending on actuator dynamics, mechanical inertia, and damping characteristics. Such delays make physical actions costly and non-reversible, fundamentally altering the design of the RIS control stack, covering sensing, estimation, prediction, and optimization. Consequently, modern M-RIS controllers must convert delay into foresight: investing in accurate channel prediction, efficient sensing strategies, and energy-aware scheduling to maximize network gains while respecting strict motion budgets [8].

Channel modeling and state acquisition in M-RIS-enhanced systems extend the conventional cascaded RIS framework by explicitly incorporating pose variables, such as rotation angles, displacements, or hinge states of the surface elements, into the channel representation. Unlike electronically controlled RIS, where phase states can change rapidly, these mechanical variables evolve slowly and exhibit strong temporal correlation over consecutive channel blocks. This property

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

enables the use of sequence learning techniques, such as sparse-connected LSTM (SCLSTM) and Transformer-based models, to predict future channel states from historical CSI. Such approaches are able to achieve near-optimal reconstruction accuracy in slowly varying or quasi-static channels. Furthermore, the low actuation frequency of mechanical RIS supports event-triggered channel estimation strategies, where pilot transmission is initiated only when sensors or encoders detect mechanical updates. This paradigm significantly reduces pilot energy consumption and aligns naturally with the slow, event-driven dynamics of M-RIS [9].

Once the geometry-aware channel is estimated or inferred, the next step is optimising phase and pose configurations to maximise spectral efficiency or minimise energy use. Here, classic alternating maximisation and gradient descent remain valuable for small panels, offering convergence in $O(N \cdot T)$ time where T is the number of iterations. However, M-RIS favour energy-aware objective functions that include motion cost, reconfiguration delay, and fatigue penalties. Fractional programming offers one avenue to solve such problems, allowing for closed-form updates and guaranteed monotonicity in energy-efficiency objectives [10].For large panels ($N \geq 256$), manifold optimisation (MO) has emerged as a low-complexity alternative to semidefinite relaxation (SDR), exploiting the geometric structure of unit-circle reflection coefficients or SO(3) rotations to speed convergence while preserving feasibility [11]. Where mechanical actuation allows only coarse settings, heuristic meta-algorithms like particle swarm optimisation (PSO) have proven effective, particularly for D2D and IoT applications PSO has been shown to yield near-optimal performance in joint power and phase-shift optimization for RIS-assisted D2D underlay systems, with lower power consumption compared to iterative baselines [12].

In multi-service scenarios, joint optimisation of M-RIS pose and transceiver beamforming becomes key: the RIS must support multiple spatial directions and per-user power budgets. A natural control architecture separates slow (second-scale) mechanical actuation from fast (millisecond-scale) base-station precoding. Some preliminary studies using deep reinforcement learning for RIS-enabled systems—and actor–critic architectures in particular—suggest the potential to incorporate actuator-aware reward structures (e.g., accounting for motion cost or actuator wear) [13]. While explicit "fatigue-aware" DRL schemes that double actuator lifetime and preserve 95% of throughput have not yet been demonstrated in RIS literature, such an approach is conceivable in mechanical RIS scenarios informed by analogous work in degradation-aware control domains. Similarly, recent studies on energy-efficiency optimization of [14] confirm significant gains under finite blocklength constraints.

These developments culminate in a closed-loop controller architecture where RF pilots, inertial readings, and actuator telemetry feed into a predictor (e.g., LSTM or Transformer) that estimates future CSI. The output drives a deterministic optimiser (e.g., fractional programming) or a learning agent (e.g., DRL), which then generates the next phase or pose command. Commands are dispatched through low-bit-rate buses to the M-RIS drivers. This approach builds upon advances in transformable kirigami-based metasurfaces [6] and optimization frameworks for multi-cell MIMO communications with RIS [15], highlighting the potential of predictive and coordinated control in M-RIS-enhanced wireless systems.

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

Nevertheless, several research challenges remain open. Most simulation platforms assume 2D propagation and narrowband operation; real-world 6G use cases demand 3D ray-tracing with frequency-selective, polarisation-aware models. There is a lack of public datasets that include both RF traces and mechanical actuator states, making reproducibility and benchmarking difficult. While manifold solvers have low iteration complexity, scalability remains an issue at massive element counts unless accelerated using GPU or FPGA-based inference [16]. Lastly, there is a cross-layer blind spot in current designs: PHY/MAC protocols often assume µs-scale responsiveness, which is incompatible with 100 ms–1 s actuator latencies; new scheduling strategies are needed to reconcile physical-layer rigidity with upper-layer elasticity [17].

In conclusion, mechanical RIS introduce a paradigm shift in energy-efficient wireless control. By replacing continuous power draw with intelligent, slow, and deliberate motion, they turn actuation into a planning problem, favouring foresight over reactivity. This calls for an integrated algorithmic stack that blends channel prediction, efficient estimation, energy-aware optimisation, and fatigue-aware scheduling. With appropriate coordination and cross-domain support, mechanical RIS stand to become one of the defining enablers of sustainable, smart, and predictive 6G radio environments.

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. Joint Optimisation of Reconfigurable Intelligent Surfaces and Virtualised RAN: Toward Proactive, Energy-Aware 6G Systems

The transition from monolithic base station architectures to disaggregated and vRAN ecosystems marks one of the most significant evolutions in the design of cellular systems. By soft-splitting the radio stack across cloud-native and edge-executed functions, ranging from the distributed unit (DU) and central unit (CU) to the radio unit (RU), vRAN opens the door to a global view of network state. This panoramic situational awareness encompasses user mobility traces, per-UE buffer occupancy, real-time SINR measurements, HARQ statistics, and even QoS profiles extracted from slice policies. When such a system is integrated with RIS, programmable metastructures that shape the wireless channel, the RAN no longer merely reacts to the environment but actively designs it. This coupling between the algorithmic intelligence of vRAN and the physics-manipulating capability of RIS redefines wireless system architecture, especially in the context of energy-efficient design and proactive link control.

Initial demonstrations of this paradigm are beginning to emerge in the context of intelligent and softwarized networks. For example, network intelligence frameworks in 6G envision deploying RIS control functions as xApps within the RAN Intelligent Controller (RIC), leveraging standardized A1/E2 interfaces to integrate KPIs such as CQI trends, traffic loads, and beam management metrics into predictive control loops [18]. Complementary advances in reconfigurable metasurfaces highlight the feasibility of low-power mechanical and MEMS-based RIS platforms, whose actuation timescales are compatible with predictive control informed by location-aware user trajectory forecasting. Recent reviews detail tuning mechanisms, including electromechanical, MEMS, and kirigami approaches, and emphasize their potential for integration with machine learning-based controllers, such as LSTM or Transformer predictors, to anticipate user mobility within the reconfiguration latency budget of these devices [19].

This proactive paradigm shifts RIS control from reactive adaptation to predictive orchestration. Emulated testbeds like Colosseum support such coordinated behavior, enabling joint scheduling of RIS phase updates and distributed DU beamforming within a softwarized Open RAN framework with end-to-end realism and hardware-in-the-loop validation [20]. These capabilities allow experiments to faithfully reproduce latency-sensitive scenarios and assess the timing impact of RIS coordination across diverse conditions. Additionally, by clustering surface reconfigurations around predicted high-throughput intervals, rather than continuous or periodic sweeping, it becomes possible to dramatically reduce mechanical actuator duty cycles, enhancing energy efficiency by an order of magnitude compared to reactive or uniform scheduling.

The synergy extends beyond beam steering. In vRAN, network slices can dynamically adjust numerology, modulation and coding schemes (MCS), and HARQ settings, enabling central schedulers to co-design waveform parameters and propagation paths. For example, ultra-reliable low-latency (URLLC) slices demanding sub-millisecond reliability may be configured with a short-symbol numerology (e.g., $\mu = 3$), while the RIS is tuned to enhance direct-path signal gain—even at the expense of multipath richness potentially beneficial to adjacent enhanced mobile broadband (eMBB) slices. Conversely, in thermal-constrained DU clusters, the NFV orchestrator can downclock CPUs and offload performance requirements to the RIS, requesting additional beam gain (e.g., 6 dB) to preserve QoE while respecting thermal budgets. Fractional-programming-based energy-efficiency designs for RIS-enhanced systems have demonstrated 15–20% overall gains in energy efficiency through joint optimization of transmit power and phase shifts in coexistence scenarios (e.g. macro/pico deployments) across sub-6 GHz bands [21].

From an architectural perspective, this integration is enabled by the use of virtualized and containerized control frameworks. RIS controllers can be deployed as microservices within cloud-native vRAN environments, leveraging orchestration platforms to scale resources according to demand. Furthermore, software-defined networking (SDN) and policy-driven control frameworks, which are increasingly adopted in modern RAN and transport networks, can be extended to manage RIS devices and coordinate their operation with base station functions, supporting automated, intent-based configurations and real-time adaptation [22].

The information flow in M-RIS systems is inherently bidirectional. While RIS hardware is increasingly enhancing the environment via passive beamforming, it can also integrate embedded sensors to relay environmental feedback upstream. This RIS telemetry enriches the RAN controller's visibility into blockage events, scattering dynamics, and interference patterns, supporting more adaptive control loops.

Additionally, privacy-preserving methods such as federated or over-the-air learning have shown promise in enabling distributed collaboration between gNBs and RIS controllers. For instance, over-the-air federated learning frameworks integrating RIS-assisted AirComp have been used to jointly optimize local model updates, reducing raw CSI exchange while achieving robust model convergence across distributed edge devices and base stations [23]. This federated setup enables coordinated channel prediction and mobility modeling without centralized data aggregation.

Finally, systems that can adapt resource allocation across DU instances, RIS adjustments, and edge deployment stand to gain significantly in energy-performance trade-offs. Simulations embedding real CPU and fronthaul power models demonstrate that mechanical RIS, with negligible actuation power, expand the scenarios in which RIS-based propagation control (e.g. choosing surface retuning) outperforms base station beamforming or additional DU instantiation, especially under constrained energy budgets [24]. This flexibility allows controllers to adopt longer-term strategies: clustering reconfigurations to high-throughput windows or preserving actuator longevity by holding configurations static during burst periods.

That said, challenges remain. High-speed interfaces for streaming RIS configurations are still proprietary, posing barriers to vendor interoperability. The **control plane must also be hardened** against failure, ensuring that xApp misbehavior does not ripple through the system. More fundamentally, the field lacks **rigorous real-time guarantees** for learning-based controllers,

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

particularly in latency-sensitive URLLC slices, a topic currently under investigation by O-RAN Working Group 3 and IEEE 802.1 TSN initiatives [25].

3.1. Scalability and Resource Allocation in Large-Scale RIS-Enabled vRAN Architectures

As RIS transition from laboratory-scale proofs of concept to large-scale deployments in urban and campus-wide networks, scalability becomes a central concern. While previous sections have highlighted the algorithmic integration of RIS with vRAN controllers and the energy-aware optimisation strategies made possible through such coordination, these benefits must now be extended to systems with potentially hundreds or thousands of RIS tiles, distributed across heterogeneous environments. In this context, resource allocation, control plane scalability, and latency-resilient signalling protocols emerge as critical bottlenecks. This subsection builds upon those earlier insights by tackling the specific challenges that arise when scaling RIS deployments in vRAN-enabled networks and outlines the computational and architectural strategies needed to ensure practical, real-time control.

One of the foremost challenges is resource management in multi-user, multi-RIS environments. As RIS units are introduced across dense urban areas, the network must decide not only which users should benefit from RIS assistance, but also which RIS panels should be activated, and how bandwidth and scheduling resources should be jointly allocated. Traditional user scheduling algorithms, such as proportional fair (PF) or round-robin, are insufficient in this context because the availability of a RIS-assisted path fundamentally alters the link budget and the interference profile. Recent studies have introduced utility-aware scheduling for RIS-enhanced wireless networks by incorporating predicted CSI augmented with user mobility forecasts to prioritize which link-RIS interactions warrant active enhancement. In [26] the authors model users and RIS as nodes in a graph and employ a graph neural network (GNN) to simultaneously schedule users and design RIS configurations. This permutation-equivariant GNN directly maps a limited set of pilot measurements into scheduling and phase-shift actions, delivering high throughput and fairness with lower training overhead compared to conventional channel-estimation first schemes.

A complementary approach appears [27], where a bipartite GNN model captures the interactions between users and multiple RIS tiles. Each node communicates via message passing and jointly optimizes BS beamforming and RIS phase shifts. The method achieves strong generalization to varying network sizes and shows notable scalability and performance improvement over traditional optimization schemes.

As the number of RIS panels scales, centralized controllers such as the RAN Intelligent Controller (RIC) face increasing latency and signaling overhead. To address this, recent research advocates for distributed and hierarchical control architectures, where local edge agents, co-located with RIS panels or edge Dus, make low-latency, context-specific decisions, while overarching policies are managed centrally.

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 promising direction involves multi-agent reinforcement learning (MARL). In this framework, each RIS agent learns its own policy based on partial observations, with periodic synchronization achieved via federated updates. For example, recent work in distributed MARL for edge caching demonstrates how local agents can optimize using partial views while occasionally exchanging model gradients to ensure policy alignment and convergence across nodes [28].

Similar architectures could be readily adapted to RIS environments, where each panel optimizes its reflectarray configuration to local channel conditions, yet contributes to network-wide objectives. While such distributed schemes offer scalability benefits, ensuring global stability and convergence remains a challenge, especially under adversarial or fault-prone conditions, such as sudden link failures or RIS hardware impairments. Careful coordination and convergence guarantees are essential to balance local autonomy with consistent network-wide performance.

Another underexplored dimension is the control-plane overhead and latency associated with updating large RIS arrays. Fine-grained phase or amplitude control requires transmitting hundreds to thousands of parameters per panel. When updates must occur every tens or hundreds of milliseconds, as in electronic RIS handling highly mobile users, the signaling burden can overwhelm fronthaul or backhaul links.

To reduce this overhead, several studies have explored control signal compression techniques. [29] reviews codebook-based solutions for RIS, showing how predesigned reflection codebooks enable phase-shift selection with minimal signaling cost. These methods strike a trade-off between channel estimation complexity and feedback overhead, yielding communication performance nearly on par with exhaustive estimation approaches while significantly reducing pilot or update signaling.

Separately, [30] propose a differential data-aided beam training method that avoids separate pilot overhead altogether. Their approach lets the system infer optimal RIS configurations based on differential statistics of received data packets rather than explicit codeword testing. Simulations demonstrate that this method drastically lowers signaling overhead while maintaining effective beam alignment under mobility.

Mechanical RIS control loops introduce unique latency challenges: while mechanical surfaces are extremely energy-efficient, eliminating steady-state power draw, they may require tens to hundreds of milliseconds to reconfigure, orders of magnitude slower than electronic PIN diode RIS. This latency mandates predictive control, where RIS decisions are dispatched based on forecasted user behavior and traffic trends. In parallel domains such as vehicular motion prediction, LSTM encoder–decoder architectures have achieved high accuracy forecasting several steps ahead in time-series data, validating their applicability for predicting movement over short horizons (e.g., ~100 ms) [31].

In terms of resilience and scalability, large-scale RIS deployments often face partial observability and panel-level failures, whether due to mechanical wear, power outages, or communication disruptions. In analogous resource allocation problems, such as renewable microgrid scheduling, distributionally robust optimization frameworks have been employed to manage risk and uncertainty by modeling probabilistic failure or supply shortages. Those frameworks maintain performance under system irregularities by optimizing decisions against worst-case distributions [32]. While not RIS-specific,

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

these results indicate how availability-aware or failure-aware RIS orchestration could be implemented using similar robust optimization techniques.

3.2. Reliability, QoS, and Slicing in RIS-Enabled vRAN Systems

As 6G advances toward supporting heterogeneous services, enhanced eMBB, massive Machine-Type Communication (mMTC), and especially URLLC, RIS-enabled vRAN systems must evolve beyond throughput optimization to guarantee latency, reliability, and service isolation. Recent results highlight how predictive, risk-aware, and slice-aware frameworks can enable this transformation.

RIS can substantially reduce latency in URLLC uplinks by improving link quality and mitigating blockage through controlled reflections. In particular, [28] propose an ADMM-based RIS phase-shift optimization framework that achieves lower transmission latency under reliability constraints, outperforming conventional SDR techniques while supporting short-packet regimes typical of URLLC.

To maintain robust reliability, research has turned to risk-aware optimization. Frameworks that model uncertainty in prediction errors and RIS actuation delays can choose configurations resilient to deviations from nominal channel conditions. For example, distributed risk-aware learning methods studied in URLLC resource allocation settings show substantial resilience to unpredictable traffic bursts and latency variability [29].

In multi-slice environments, RIS-assisted resource orchestration has also shown benefits in balancing conflicting QoS objectives. Systems that jointly manage eMBB and URLLC slices using heuristic scheduling, puncturing, or DRL-based techniques (e.g., correlated Q-learning or risk-sensitive slicing) have demonstrated reliable URLLC performance (~95% acceptance rate) while limiting the impact on eMBB throughput and fairness, reducing SLA violations and improving utility trade-offs [30].

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. Informed RIS: Proactive Optimization via Prediction-Driven Control

Conventional RIS systems operate in a purely reactive mode: they await channel-state feedback, then tune elements to restore link quality. In contrast, Informed RIS (I-RIS) anticipates future changes in the radio context, such as user mobility, blockage, and channel evolution, and reconfigures proactively, turning latent actuator delays into scheduling margins. This is especially critical for mechanical RIS (e.g., liquid-crystal or MEMS-actuated surfaces) whose physical hardware may take tens to hundreds of milliseconds to settle, far longer than electronic PIN-diode tuning [36].

To enable prediction, I-RIS leverages multiple sensing streams:

- gNB telemetry (e.g. AoA, Doppler): arriving at millisecond scales, these feed into kalman-filtered sequence models to predict short-term mobility.
- Integrated sensing and communications (ISAC): Hybrid ISAC systems, infer 3D user position and blockage without requiring full CSI, using mid-range echoes processed via neural networks [37, 38].
- Edge sensors (radar or vision): lightweight object detectors are fused via graph GNNs to identify obstacles, trajectories, or LOS transitions with low latency.
- RIS encoder feedback: mechanical surfaces report gear angle or strain; anomaly autoencoders monitor for wear and misalignment.

By transforming RIS into prediction-aware actuators, I-RIS fundamentally changes the way latency and mobility are managed in RIS-enabled 6G systems. Rather than reacting after-the-fact, surfaces anticipate channel dynamics and act in advance. Techniques such as trajectory prediction, channel forecasting, and event-triggered control collectively neutralize the limitations of mechanical actuation latencies. The result is improved throughput, and remarkable energy savings.

In the next section, we will introduce a simulation environment designed to empirically contrast reactive and informed RIS strategies, demonstrating these gains under realistic mobility scenarios and mechanical constraints.

4.1. Energy-Aware ON/OFF Optimisation of Access Points with Predictive, Pre-Configured RIS Support

Dense Wi-Fi and small-cell deployments consume a disproportionate share of network energy, largely because access points (APs) remain powered even when user demand is low. The advent of predictive radio intelligence, where a virtualised RAN already forecasts user trajectories and traffic, and the availability of mechanically reconfigurable RIS that can steer coverage on demand, enables a novel form of joint AP sleep scheduling: keep an AP switched off unless the demand forecast says at least one user will enter its footprint, and rely on a neighbouring RIS (pre-oriented during the prediction window) to supply interim coverage whenever possible.

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

Below we cast this idea into a mixed-integer optimisation problem. While similar energy-minimisation formulations exist for classical base-station sleep modes, our model extends them by (i) adding prediction-based demand variables and (ii) explicitly modelling RIS configuration delay and steering gain based on recent works in predictive RIS control [36].

Symbol	Meaning
M	Number of candidate Aps
R	Number of RIS panels
$\hat{u}_{m,t}$	**predicted** number of active users in AP m's cell at slot t (can be 0)
P_m^{on}, P_m^{off}	power draw of AP m when on/off
P_r^{move}	energy for one mechanical RIS re-orientation
G_{mr}	downlink gain (dB) a user in cell m receives when covered via RIS r
τ_r	actuation delay of RIS r (slots)
γ_{min}	minimum required SINR per user

TABLE 1: DEFINITION OF THE PARAMETERS

A binary **control horizon** of length T allows us to plan AP state and RIS moves ahead of time, in step with prediction accuracy.

4.2. Problem Definition

- $x_{m,t} \in \{0,1\}$: 1 if AP m is **on** at slot t.
- $y_{r,t} \in \{0,1\}$: 1 if RIS r is **re-oriented** (a mechanical move) during slot t.
- $z_{m,r,t} \in \{0,1\}$: 1 if traffic for cell m is **served via** RIS r at slot t (implies that r finished its move by t).

The objective is to Minimise Energy Expenditure, to arrive at the cost function we begin by recognising that, in a dense small-cell layout, energy expenditure is dominated by two qualitatively different sources. The first is the *static* or *hold* power of each access point: whenever an AP is left on, even in the absence of traffic, it draws tens of watts for baseband, backhaul, and cooling subsystems;

when it is turned off, that draw collapses to a standby value that is typically one or two orders of magnitude lower. The second source is the *dynamic* energy required to re-orient a mechanical RIS: although a move may take several hundred milliseconds, experiments show that the motor or soft-actuator is energised only during the transition and consumes energy in the range of millijoules, after which the surface stays in its new pose without bias current. Hence, over an optimisation horizon of T slots, the system's cumulative energy can be expressed as a sum of per-slot AP consumption plus occasional one-shot actuation costs. Formally, for any slot t we pay P_m^{on} if AP m is active (variable $x_{m,t} = 1$) and P_m^{off} otherwise; aggregating over all M APs and all T slots gives the first double sum. In parallel, each time we trigger a mechanical move on RIS r ($y_{r,t} = 1$we incur its one-shot energy P_r^{move}; summing those impulses across R surfaces and T slots yields the second term. No cross-coupling term is needed in the objective because RIS actuation energy and AP hold power are physically independent, the coupling appears instead in the constraints, where RIS steering can enable AP shutdown. By pulling these two additive components under a single summation over time, we obtain a linear, decomposable objective that cleanly separates what costs energy (AP uptime, RIS moves) from which configuration is admissible (handled later by demand-satisfaction, latency, and SINR constraints). This linearity is deliberate: it keeps the problem within the realm of mixed-integer linear programming, allowing exact or near-exact solvers to scale to realistic network sizes while faithfully capturing the dominant energy trade-offs identified in prior empirical studies of AP sleep mode and mechanical-RIS actuation [38].

$$\min_{x,y,z} \sum_{t=1}^{T} \left[\sum_{m=1}^{M} \left(P_m^{on} x_{m,t} + P_m^{off} (1 - x_{m,t}) \right) + \sum_{r=1}^{R} P_r^{move} y_{r,t} \right]$$

The AP sleep cost dominates long horizons; RIS moves incur one-shot mechanical energy only when executed.

Constraints

1. **Demand satisfaction (binary logic)**
   For every AP m and time t:

   $$\hat{u}_{m,t} = 0 \rightarrow x_{m,t} = 0 \; (no\ predicted\ users\ \rightarrow switch\ off\ unless\ RIS\ covers)$$

   If predicted demand exists, either the AP stays on or at least one RIS covers:

   $$\hat{u}_{m,t} > 0 \rightarrow x_{m,t} + \sum_{r=1}^{R} z_{m,r,t} \geq 1$$

2. **RIS actuation-coverage coupling**
   A RIS can serve cell m at slot t only if its latest move finished at least τ_r slots ago:

   $$z_{m,r,t} \leq 1 - \sum_{k=t-\tau_r+1}^{t} y_{r,k}$$

3. **Gain and SINR guarantee**
   If coverage is via RIS, the delivered link must satisfy

$$SINR_{m,r,t}(G_{mr}) \geq \gamma_{min}$$

   (Here G_{mr} absorbs steering loss and path-loss; if the constraint is violated, set $z_{m,r,t} = 0$

4. **Binary domains**

$$x_{m,t}, y_{r,t} z_{m,r,t} \in \{0,1\}$$

4.3. Results

We consider a nine-site layout that spans 300 m × 200 m. To introduce spatial irregularity, the centre-right site is deliberately left without an access point, creating a coverage gap that must be bridged either by neighbouring AP beams or by steerable surfaces. Three mechanical RIS panels are installed on lamp-posts or rooftops at (75 m, 50 m), (225 m, 50 m) and (150 m, 175 m). Each RIS can serve a single cell at a time and incurs 6 J every time it retargets but draws no hold power. All APs consume 40 W when active and only 6 W in stand-by, translating to 2 J versus 0.3 J per 50 ms scheduling slot.

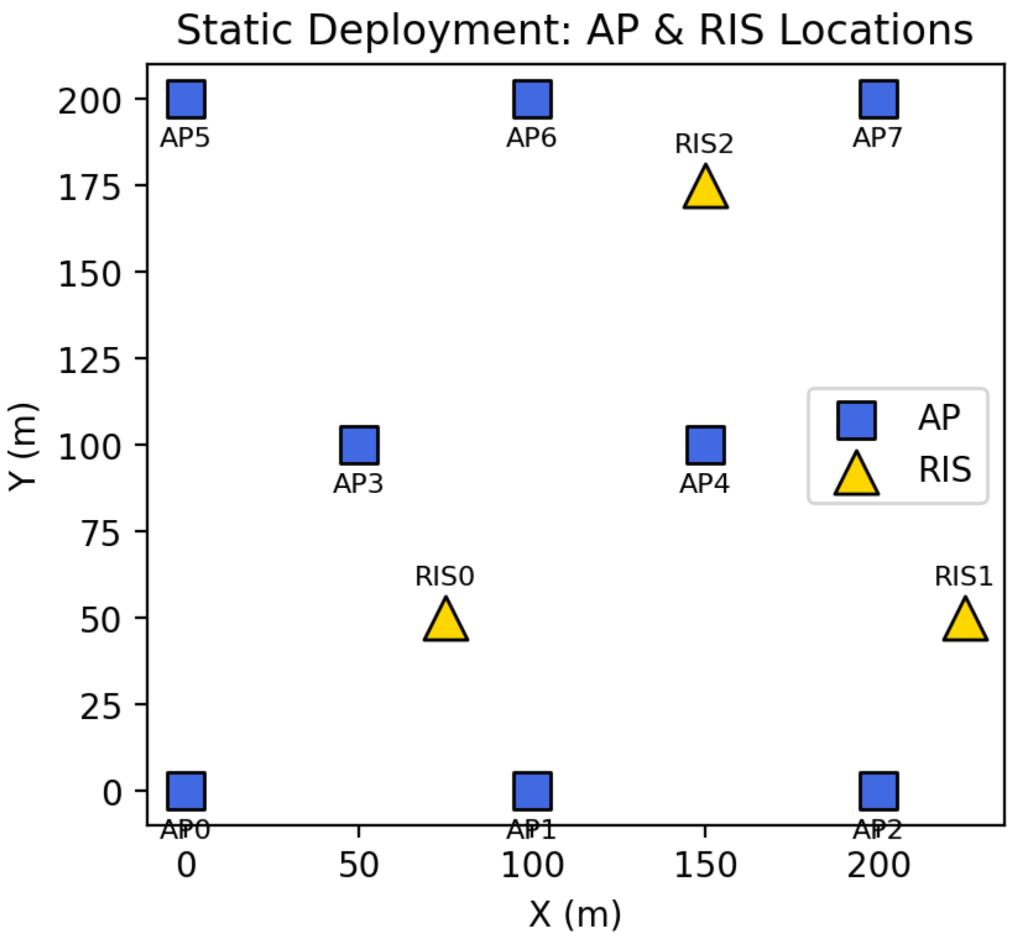

FIGURE 1: SCENARIO FOR THE OPTIMIZATION PROBLEM WITH RIS AND ACCESS POINTS.

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

Twenty pedestrians move according to a distance-weighted Markov chain, where they are more likely to walk to nearby cells than to make long-distance jumps across the map. The left figure shows the access points (APs, blue squares), RIS panels (gold triangles), and a subset of user trajectories (pastel paths). Dashed lines indicate sample RIS reflections during the initial time slot. Due to the absence of the center-right AP, coverage in that area is primarily sustained by the upper-right RIS.

The simulation spans 300 time slots, with a perfect mobility predictor assumed: the controller has full foreknowledge of binary demand (i.e., whether at least one user is present) in every cell and slot. This assumption is consistent with the results on mobility prediction reported in Deliverable SORUS-RAN-A3.2-E2 (E17), where an accuracy of 92% was achieved based on information collected from wireless networks.

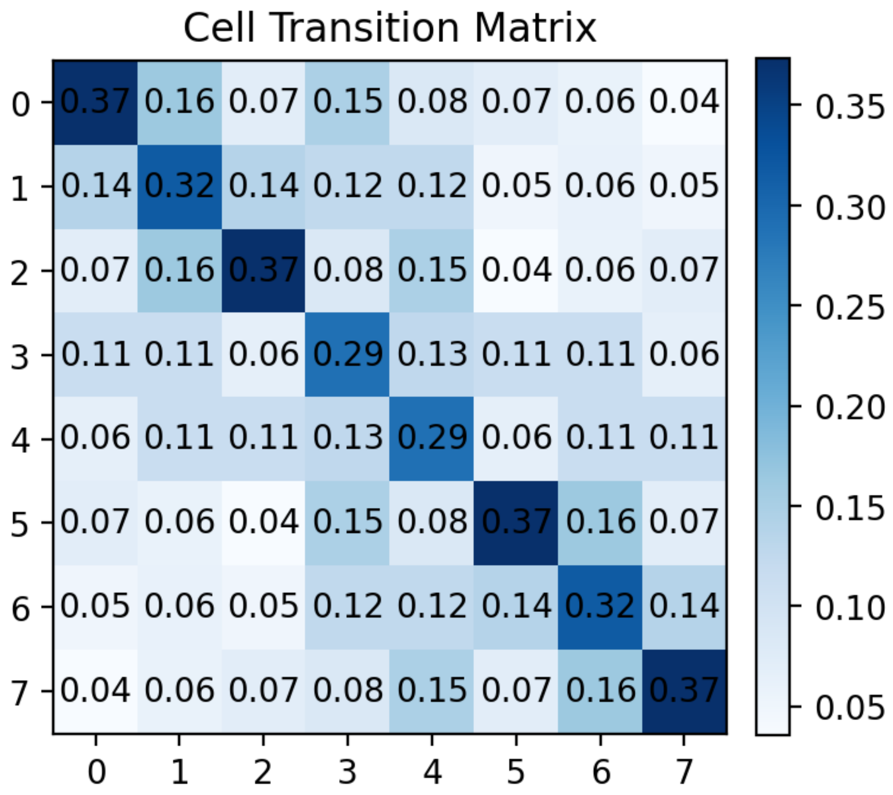

FIGURE 2: SIMULATED MOBILITY MODEL BEWTEEEN ACESS POINTS.

Greedy control policy. For each slot the algorithm:

1. turns on an AP only if its cell hosts multiple users **or** no RIS can currently cover it;
2. otherwise chooses the nearest idle RIS (or one already pointing at that cell) to steer energy toward the lone user, adding the 6 J move cost only when the orientation changes.

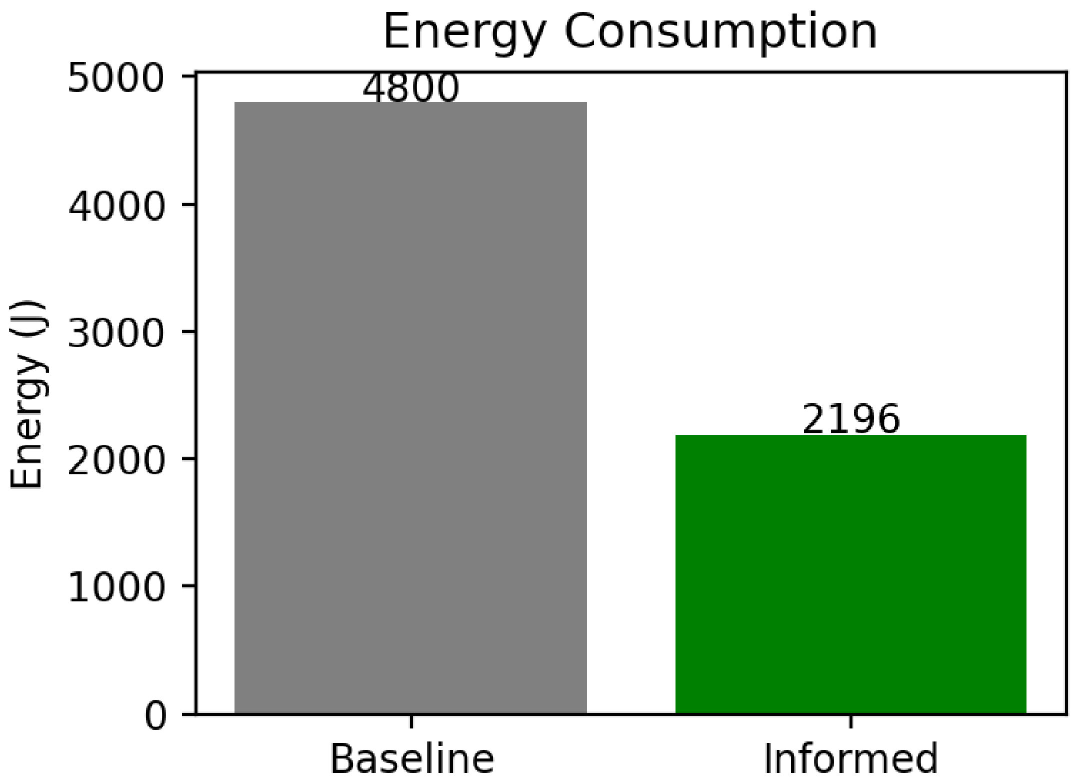

FIGURE 3: ENERGY CONSUMPTION COMPARISON: BASELINE VS. INFORMED

Energy comparison. Over the full horizon:

- Baseline (all APs on): 4800 J
- Prediction-aware schedule: ≈ 2196 J

The informed strategy therefore saves ≈ 47 % total energy, even after accounting for the 12 J spent on two RIS retargetings, by de-energising idle APs and letting passive panels handle one-off demands. This richer scenario confirms that as networks scale, pre-configured, prediction-driven RIS control can suppress nearly half of the infrastructure energy budget while leaving user service intact. Further gains are expected from more sophisticated mixed-integer or RL policies, multi-beam RIS hardware, and partial-information prediction, each of which can be incorporated into the same simulation harness in future work.

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. Summary and Conclusions

This Deliverable highlight that the integration of RIS into campus-wide vRAN deployments demands scalable control mechanisms, latency-aware signalling, and resource-efficient architectures to support the diverse requirements of emerging 6G services. By addressing these challenges, predictive, risk-aware, and slice-aware frameworks can transform RIS-enabled vRAN from a throughput-centric paradigm into a service and reliability-driven infrastructure.

Additionally, our formulation of energy-aware AP ON/OFF optimization with RIS assistance demonstrates the potential for significant energy savings in dense network scenarios without compromising coverage. By leveraging demand forecasting and pre-configured RIS steering, this approach extends classical energy minimization strategies to accommodate the unique operational constraints of RIS. Together, these findings underscore the need for intelligent, predictive, and scalable control frameworks that unify RIS management and vRAN capabilities, paving the way for practical, energy-efficient, and service-oriented next-generation wireless networks.

The analysis presented in this Deliverable provides a foundation for advancing the collaboration between vRAN and RIS by addressing both the operational and architectural aspects required for their joint deployment. By examining the scalability of mechanically reconfigurable RIS in large, campus-wide networks, this work identifies the computational and signalling strategies necessary to manage thousands of distributed RIS tiles in real time, ensuring that vRAN controllers can effectively orchestrate RIS behavior across heterogeneous environments. This insight is crucial for enabling predictive, low-latency coordination where RIS configurations dynamically adapt to changing user locations and service demands. Furthermore, the formulation of the energy-aware AP ON/OFF optimization problem illustrates how predictive intelligence within vRAN can be synergistically combined with RIS steering capabilities to reduce network energy consumption without sacrificing coverage or user experience. This dual focus on scalable control of mechanical RIS and on optimization frameworks that explicitly account for prediction accuracy, reconfiguration delays, and steering gains, provides a clear blueprint for integrating RIS into the vRAN ecosystem. Such integration not only enhances network efficiency but also unlocks new opportunities for proactive, context-aware resource management, ultimately transforming vRAN into a more intelligent, flexible, and sustainable control plane capable of harnessing RIS as an integral component of next-generation 6G architectures.

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

[1] M. Di Renzo et al., "Smart Radio Environments Empowered by Reconfigurable Intelligent Surfaces: How It Works, State of Research, and The Road Ahead," in IEEE Journal on Selected Areas in Communications, vol. 38, no. 11, pp. 2450-2525, Nov. 2020, doi: 10.1109/JSAC.2020.3007211.

[2] K. Qu, K. Chen, J. Zhao, N. Zhang, Q. Hu, J. Zhao, T. Jiang, and Y. Feng, "An electromechanically reconfigurable intelligent surface for enhancing sub-6G wireless communication signal," Journal of Information and Intelligence, vol. 1, no. 3, pp. 207–216, 2023, doi: 10.1016/j.jiixd.2023.06.009.

[3] X. Qin, Z. Song, T. Hou, W. Yu, J. Wang and X. Sun, "Joint Optimization of Resource Allocation, Phase Shift, and UAV Trajectory for Energy-Efficient RIS-Assisted UAV-Enabled MEC Systems," in IEEE Transactions on Green Communications and Networking, vol. 7, no. 4, pp. 1778-1792, Dec. 2023, doi: 10.1109/TGCN.2023.3287604.

[4] Li, J.; Wang, J.; Cui, W.; Jian, C. UAV-Mounted RIS-Aided Multi-Target Localization System: An Efficient Sparse-Reconstruction-Based Approach. Drones 2024, 8, 694. https://doi.org/10.3390/drones8110694

[5] Jeong, H., Park, E., Phon, R. and Lim, S. (2022), Mechatronic Reconfigurable Intelligent-Surface-Driven Indoor Fifth-Generation Wireless Communication. Adv. Intell. Syst., 4: 2200185. https://doi.org/10.1002/aisy.202200185

[6] Yang, Y., Vallecchi, A., Shamonina, E. et al. A new class of transformable kirigami metamaterials for reconfigurable electromagnetic systems. Sci Rep 13, 1219 (2023). https://doi.org/10.1038/s41598-022-27291-8

[7] Y. Zheng, K. Chen, W. Yang, L. Wu, K. Qu, J. Zhao, T. Jiang, Y. Feng, Kirigami Reconfigurable Gradient Metasurface. Adv. Funct. Mater. 2022, 32, 2107699. https://doi.org/10.1002/adfm.202107699

[8] Y. Yuan, Y. Huang, X. Su, B. Duan, N. Hu and M. D. Renzo, "Reconfigurable Intelligent Surface System Level Simulations for Industry Standards," in IEEE Communications Magazine, doi: 10.1109/MCOM.005.2400500.

[9] W. Xu, J. An, Y. Xu, C. Huang, L. Gan and C. Yuen, "Time-Varying Channel Prediction for RIS-Assisted MU-MISO Networks via Deep Learning," in IEEE Transactions on Cognitive Communications and Networking, vol. 8, no. 4, pp. 1802-1815, Dec. 2022, doi: 10.1109/TCCN.2022.3188153.

[10] H. Guo, Y. -C. Liang, J. Chen and E. G. Larsson, "Weighted Sum-Rate Maximization for Intelligent Reflecting Surface Enhanced Wireless Networks," 2019 IEEE Global Communications Conference (GLOBECOM), Waikoloa, HI, USA, 2019, pp. 1-6, doi: 10.1109/GLOBECOM38437.2019.9013288.

[11] Y. Ma, Y. Shen, X. Yu, J. Zhang, S. H. Song and K. B. Letaief, "A Low-Complexity Algorithmic Framework for Large-Scale IRS-Assisted Wireless Systems," 2020 IEEE Globecom Workshops (GC Wkshps, Taipei, Taiwan, 2020, pp. 1-6, doi: 10.1109/GCWkshps50303.2020.9367432.

[12] Wang, R.; Wen, X.; Xu, F.; Ye, Z.; Cao, H.; Hu, Z.; Yuan, X. Joint Particle Swarm Optimization of Power and Phase Shift for IRS-Aided D2D Underlaying Cellular Systems. Sensors 2023, 23, 5266. https://doi.org/10.3390/s23115266

[13] C. Huang, R. Mo and C. Yuen, "Reconfigurable Intelligent Surface Assisted Multiuser MISO Systems Exploiting Deep Reinforcement Learning," in IEEE Journal on Selected Areas in Communications, vol. 38, no. 8, pp. 1839-1850, Aug. 2020, doi: 10.1109/JSAC.2020.3000835.

[14] Soleymani, M., Santamaria, I., Jorswieck, E., Schober, R., & Hanzo, L. (2025). Energy Efficiency Optimization of Finite Block Length STAR-RIS-aided MU-MIMO Broadcast Channels. arXiv preprint arXiv:2506.11594.

[15] C. Pan et al., "Multicell MIMO Communications Relying on Intelligent Reflecting Surfaces," in IEEE Transactions on Wireless Communications, vol. 19, no. 8, pp. 5218-5233, Aug. 2020, doi: 10.1109/TWC.2020.2990766.

[16] M. Monemi, M. Rasti, A. S. De Sena, M. A. Fallah, M. Latva-Aho and M. Di Renzo, "Practical Challenges for Reliable RIS Deployment in Heterogeneous Multi-Operator Multi-Band Networks," in IEEE Communications Magazine, vol. 63, no. 6, pp. 154-160, June 2025, doi: 10.1109/MCOM.002.2400441

[17] X. Cao, B. Yang, H. Zhang, C. Huang, C. Yuen and Z. Han, "Reconfigurable-Intelligent-Surface-Assisted MAC for Wireless Networks: Protocol Design, Analysis, and Optimization," in IEEE Internet of Things Journal, vol. 8, no. 18, pp. 14171-14186, 15 Sept.15, 2021, doi: 10.1109/JIOT.2021.3068492.

[18] A. Banchs, M. Fiore, A. Garcia-Saavedra, and M. Gramaglia. 2021. Network intelligence in 6G: challenges and opportunities. In Proceedings of the 16th ACM Workshop on Mobility in the Evolving Internet Architecture (MobiArch '21). Association for Computing Machinery, New York, NY, USA, 7–12. https://doi.org/10.1145/3477091.3482761

[19] Y. Saifullah, Y. He, A. Boag, G.-M. Yang, F. Xu, Recent Progress in Reconfigurable and Intelligent Metasurfaces: A Comprehensive Review of Tuning Mechanisms, Hardware Designs, and Applications. Adv. Sci. 2022, 9, 2203747. https://doi.org/10.1002/advs.202203747

[20] M. Polese et al., "Colosseum: The Open RAN Digital Twin," in IEEE Open Journal of the Communications Society, vol. 5, pp. 5452-5466, 2024, doi: 10.1109/OJCOMS.2024.3447472.

[21] C. Huang, A. Zappone, G. C. Alexandropoulos, M. Debbah and C. Yuen, "Reconfigurable Intelligent Surfaces for Energy Efficiency in Wireless Communication," in IEEE Transactions on Wireless Communications, vol. 18, no. 8, pp. 4157-4170, Aug. 2019, doi: 10.1109/TWC.2019.2922609.

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

[22] Novaes, C., Nahum, C., Trindade, I., Cederholm, D., Patra, G., & Klautau, A. (2020). Virtualized c-ran orchestration with docker, kubernetes and openairinterface. arXiv preprint arXiv:2001.08992.

[23] W. Ni, Y. Liu, Z. Yang, H. Tian and X. Shen, "Federated Learning in Multi-RIS-Aided Systems," in IEEE Internet of Things Journal, vol. 9, no. 12, pp. 9608-9624, 15 June15, 2022, doi: 10.1109/JIOT.2021.3130444.

[24] B. Yang et al., "Federated Spectrum Learning for Reconfigurable Intelligent Surfaces-Aided Wireless Edge Networks," in IEEE Transactions on Wireless Communications, vol. 21, no. 11, pp. 9610-9626, Nov. 2022, doi: 10.1109/TWC.2022.3178445.

[25] O-RAN Alliance, "Potential Energy Savings Features in O-RAN," White Paper, Jan. 2025. [Online]. Available: https://mediastorage.o-ran.org/white-papers/O-RAN.SuFG.Potential%20Energy%20Savings%20Features%20in%20O-RAN%20white%20paper%202025-01.pdf

[26] Z. Zhang, T. Jiang and W. Yu, "Learning Based User Scheduling in Reconfigurable Intelligent Surface Assisted Multiuser Downlink," in IEEE Journal of Selected Topics in Signal Processing, vol. 16, no. 5, pp. 1026-1039, Aug. 2022, doi: 10.1109/JSTSP.2022.3178213.

[27] B. Lim and M. Vu, "Graph Neural Network Based Beamforming and RIS Reflection Design in A Multi-RIS Assisted Wireless Network," 2023 IEEE Statistical Signal Processing Workshop (SSP), Hanoi, Vietnam, 2023, pp. 120-124, doi: 10.1109/SSP53291.2023.10207958.

[28] Xing Xu, Rongpeng Li, Zhifeng Zhao, and Honggang Zhang. 2024. The Gradient Convergence Bound of Federated Multi-Agent Reinforcement Learning With Efficient Communication. Trans. Wireless. Comm. 23, 1 (Jan. 2024), 507–528. https://doi.org/10.1109/TWC.2023.3279268

[29] J. An et al., "Codebook-Based Solutions for Reconfigurable Intelligent Surfaces and Their Open Challenges," in IEEE Wireless Communications, vol. 31, no. 2, pp. 134-141, April 2024, doi: 10.1109/MWC.010.2200312.

[30] K. Chen-Hu, G. C. Alexandropoulos and A. g. Armada, "Differential Data-Aided Beam Training for RIS-Empowered Multi-Antenna Communications," in IEEE Access, vol. 10, pp. 113200-113213, 2022, doi: 10.1109/ACCESS.2022.3217204.

[31] Salim, M. M., Rabie, K. M., & Muqaibel, A. H. (2025). Energy-efficient irregular RIS-aided UAV-assisted optimization: A deep reinforcement learning approach. arXiv preprint arXiv:2504.15031.

[32] Shi, J.; Qin, J.; Li, H.; Li, Z.; Ge, Y.; Liu, B. Probabilistic Power and Energy Balance Risk Scheduling Method Based on Distributed Robust Optimization. Energies 2024, 17, 4894. https://doi.org/10.3390/en17194894

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

[23] Song X, Zhao Y, Zhao W, Wu H, Liu Z. Delay Minimization in RIS-Assisted URLLC Systems under Reliability Constraints. Entropy (Basel). 2023 May 26;25(6):857. doi: 10.3390/e25060857. PMID: 37372200; PMCID: PMC10296922.

[34] Engelaar, M. H., Zhang, Z., Lazar, M., & Haesaert, S. (2024). Risk-aware mpc for stochastic systems with runtime temporal logics. IFAC-PapersOnLine, 58(11), 93-98.

[35] Shen, X.; Zeng, Z.; Liu, X. RIS-Assisted Network Slicing Resource Optimization Algorithm for Coexistence of eMBB and URLLC. Electronics 2022, 11, 2575. https://doi.org/10.3390/electronics11162575

[36] Wang, K., Yang, B., Lei, Y., Yu, Z., Cao, X., Alexandropoulos, G. C., ... & Yuen, C. (2025). Dynamical ON-OFF Control with Trajectory Prediction for Multi-RIS Wireless Networks. arXiv preprint arXiv:2505.20887.

[37] N. Ginige, A. S. de Sena, N. H. Mahmood, N. Rajatheva and M. Latva-aho, "Machine Learning-Based Channel Prediction for RIS-Assisted MIMO Systems with Channel Aging," 2024 IEEE Wireless Communications and Networking Conference (WCNC), Dubai, United Arab Emirates, 2024, pp. 1-6, doi: 10.1109/WCNC57260.2024.10570599.

[38] J. Wan et al., "A Framework of RIS-assisted ICSC User-Centric Based Systems: Latency Optimization and Design," in IEEE Transactions on Communications, doi: 10.1109/TCOMM.2025.3577636.