

THIS IS AN AUTHOR-CREATED POSTPRINT VERSION.

Disclaimer: This work has been accepted for publication in *IEEE Transactions on Wireless Communications*.

Copyright: © 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

DOI: 10.1109/TWC.2026.3672338

RIS Control through the Lens of Stochastic Network Calculus: An O-RAN Framework for Delay-Sensitive 6G Applications

Oscar Adamuz-Hinojosa, Lanfranco Zanzi, *Member, IEEE*, Vincenzo Sciancalepore, *Senior Member, IEEE*, Marco Di Renzo, *Fellow, IEEE*, Xavier Costa-Pérez, *Senior Member, IEEE*

Abstract—Reconfigurable Intelligent Surfaces (RIS) enable dynamic electromagnetic control for 6G networks, but existing control schemes lack responsiveness to fast-varying network conditions, limiting their applicability for ultra-reliable low latency communications. This work address uplink delay minimization in multi-RIS scenarios with heterogeneous per-user latency and reliability demands. We propose Delay-Aware RIS Orchestrator (DARIO), an O-RAN-compliant framework that dynamically assigns RIS devices to users within short time windows, adapting to traffic fluctuations to meet per-user delay and reliability targets. DARIO relies on a novel Stochastic Network Calculus (SNC) model to analytically estimate the delay bound for each possible user–RIS assignment under specific traffic and service dynamics. These estimations are used by DARIO to formulate a Nonlinear Integer Program (NIP), for which an online heuristic provides near-optimal performance with low computational overhead. Extensive evaluations with simulations and real traffic traces show consistent delay reductions up to 95.7% under high load or RIS availability.

Index Terms—Stochastic Network Calculus, Smart Surfaces, RIS, Smart Radio Environments, Open Radio, Delay-Sensitive.

I. INTRODUCTION

Reconfigurable Intelligent Surfaces (RIS) are key enablers for Sixth Generation (6G) networks, mitigating signal blockages and multipath fading by electronically reconfiguring the propagation environment [1]. By adjusting the phase and amplitude of reflected signals, RIS devices establish *controlled Line-of-Sight (LoS) links* between User Equipments (UEs) and Base Stations (BSs) [2]. This reconfigurability enables new forms of adaptive radio resource scheduling, particularly suited to delay-sensitive and mission-critical applications.

Funding for open access charge: Universidad de Granada / CBUA. This work is part of the project PID2022-137329OB-C43 funded by MICIU/AEI/10.13039/501100011033 and by FEDER, EU, and also part of the project C-ING-306-UGR23 funded by Consejería de Universidad, Investigación e Innovación and by ERDF Andalusia Program 2021-2027. Furthermore, it has been partially supported by the MultiX project from Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme (Grant 101192521).

Oscar Adamuz-Hinojosa is with the Department of Signal Theory, Telematics and Communications, University of Granada, Granada, Spain (e-mail: oadamuz@ugr.es). Lanfranco Zanzi, Vincenzo Sciancalepore, and Xavier Costa-Pérez are with NEC Laboratories Europe, Heidelberg, Germany. (e-mail: {name.surname}@neclab.eu). Xavier Costa-Pérez is also with i2CAT Foundation and ICREA, Barcelona, Spain (e-mail: xavier.costa@i2cat.net). M. Di Renzo is with Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des Signaux et Systèmes, 3 Rue Joliot-Curie, 91192 Gif-sur-Yvette, France. (marco.di-renzo@universite-paris-saclay.fr), and with King's College London, Department of Engineering - Centre for Telecommunications Research, WC2R 2LS London, United Kingdom (marco.di-renzo@kcl.ac.uk).

Motivation. In addition to advanced radio resource scheduling techniques [3], the application of RIS devices in delay-sensitive scenarios can serve as a complementary approach to enhance performance. In such scenarios, it is crucial to ensure the packet delay budget of each UE is respected, with violation probabilities remaining below predefined thresholds [4]. A possible approach to address this requirement is to dynamically adapt the RIS configuration across time slots to serve different UEs, each experiencing distinct channel conditions, traffic patterns, and delay requirements. This would require adaptive UE–RIS association mechanisms capable of tracking these fast and diverse environmental changes.

To support such adaptive mechanisms, Stochastic Network Calculus (SNC) offers a rigorous mathematical framework for modeling the uncertainties arising from time-varying channels and stochastic traffic arrivals [5]. SNC enables analytical estimation of probabilistic delay bounds [6], which can be used to assess whether a given combination of UE–RIS associations satisfies the required delay constraints.

The envisioned RIS orchestrator framework, supported by SNC-based delay modeling, aligns with the principles of Open Radio Access Network (O-RAN), which promotes disaggregated, programmable, and intelligent control of RAN functions [7]. O-RAN introduces distributed control through the Non-Real Time (RT) and Near-RT RAN Intelligent Controllers (RICs), separating long-term policy optimization from short-term control decisions [8]. Within this architecture, a dedicated RIS orchestrator can leverage real-time network state information accessed via the RICs to perform adaptive UE-RIS association. This enables dynamic RIS reconfiguration in response to changing traffic patterns, mobility, and delay constraints.

Related Works. Prior research on RIS mainly targets system capacity or energy efficiency optimization, often overlooking the stringent delay requirements of ultra-Reliable Low Latency Communication (uRLLC) services. Du et al. [9] analyze a multicast downlink (DL) scenario where a single RIS assists a multi-antenna BS, jointly optimizing transmit covariance and phase shifts to maximize information-theoretic capacity under perfect Channel State Information (CSI). This static, throughput-oriented formulation does not capture stochastic latency or reliability trade-offs. Similarly, Guo et al. [10] maximize weighted sum-rate in multiuser Multiple Input Single Output (MISO) DL systems using fractional programming and block coordinate descent, later extended to imperfect CSI, but remain focused on deterministic spectral-efficiency gains

without accounting for queueing dynamics or time-varying associations. Related works by Du et al. [11] and Liu et al. [12] employ iterative optimization techniques to maximize energy efficiency or sum-rate, respectively, yet similarly rely on static rate or power metrics and do not address probabilistic delay guarantees or dynamic user–RIS reassignment. Xue et al. [13] study user association in dense mmWave multi-connectivity scenarios via power optimization, while Peng et al. [14] propose a two-timescale RIS-aided uRLLC scheme under finite blocklength using statistical CSI to reduce overhead. However, latency in this latter work is modeled deterministically through short-packet rate expressions, without per-user delay guarantees. In a related line of research, Yao et al. [15], [16] investigate hybrid multi-RIS-aided secure Integrated Sensing and Communication (ISAC) systems, jointly optimizing transmit waveforms, receive filters, and discrete reflection coefficients to enhance sensing Signal-to-Interference-plus-Noise Ratio (SINR) and physical-layer security under imperfect CSI. While these works incorporate multi-user and multi-RIS architectures with robust optimization techniques, their formulations remain snapshot-based and performance-driven (e.g., SINR maximization), without modeling stochastic traffic arrivals, queue dynamics, or probabilistic delay guarantees.

A smaller body of work considers delay-aware formulations. Mukherjee et al. [17] and Xia et al. [18] investigate Multi-access Edge Computing (MEC)-enabled networks with a single RIS, aiming to minimize maximum delay by jointly accounting for transmission and computation latency. These problems are solved per snapshot through convex relaxations, but do not model stochastic traffic arrivals or per-user reliability. In the radio interface domain, Almekhlafi et al. [19] and De Souza et al. [20] address coexistence of enhanced Mobile Broadband (eMBB) and uRLLC traffic, focusing on power allocation, precomputed RIS phase sets, and frame-based scheduling. Reliability is assessed via outage probability rather than probabilistic delay bounds, and all assume a static single-RIS setup. Liu et al. [21] study Uplink (UL) multiplexing for joint eMBB/uRLLC services with RIS-assisted scheduling, but do not consider pure uRLLC scenarios under multiple RIS devices, nor delay evolution over time.

More recent studies extend RIS techniques to increasingly dynamic settings with multiple UEs, adaptive configurations, and emerging O-RAN architectures. Soleymani et al. [22] analyze the rate region of uRLLC broadcast channels under different RIS architectures, comparing diagonal and beyond-diagonal designs under reliability constraints. Their follow-up work [23] incorporates Rate Splitting Multiple Access (RSMA) into Multiple Input Multiple Output (MIMO) RIS-aided systems with finite blocklength coding to improve rate–reliability trade-offs. Both formulations evaluate achievable rate and packet error probability under static user–RIS topologies and perfect CSI, without modeling queue dynamics or probabilistic delay bounds.

Beyond differences in optimization objectives, existing RIS frameworks are not directly applicable to probabilistic delay-bound optimization without altering their core assumptions. They are predominantly based on snapshot-driven designs with quasi-static channels, fixed user–RIS associations, and, in some cases, deterministic service rates, which prevents the

modeling of queue dynamics, stochastic traffic arrivals, and delay accumulation over time. As a result, latency is typically treated either as a secondary metric via short-packet or outage-based formulations, or as a deterministic constraint tied to static topologies. Consequently, the resulting configurations (e.g., phase-shift or power settings) are computed offline or per frame and cannot be mapped to per-slot control policies with per-UE probabilistic delay guarantees. Extending these approaches would require introducing explicit traffic and queueing models together with time-coupled decision-making, leading to fundamentally different problem formulations rather than direct adaptations. Moreover, none of the existing works jointly address multiple RIS devices, user mobility, and time-varying traffic while enabling real-time adaptability under standardized O-RAN architectures and an SNC-based framework to provide per-UE probabilistic delay guarantees.

Contributions. This work is the first to design, implement, and evaluate a multi-RIS network that serves mobile UEs with heterogeneous delay budgets and violation probabilities. The main contributions are:

C1) Proposal of Delay-Aware RIS Orchestrator (DARIO), a O-RAN-compliant RIS orchestrator that dynamically configures RIS devices across time slots to provide instantaneous LoS conditions, improving UE channels and reducing delay violation probabilities.

C2) Design a novel SNC-based delay model capturing UL rate variability due to dynamic UE-RIS configurations, allowing analytical estimation of delay violation probabilities under realistic traffic and channel dynamics.

C3) Formulation of the UE–RIS assignment as a Non-linear Integer Programming (NIP) problem, minimizing the ratio between experienced delay and delay bounds, aligned with real-time requirements.

C4) Design of a heuristic algorithm solving the NIP with near-optimal performance and low computational complexity, suitable for real-time use in O-RAN architecture.

C5) Comprehensive evaluation using simulations and real-world traffic traces, showing that DARIO consistently reduces delay across all scenarios, achieving up to 95.71% improvement in high-load or high-RIS-availability deployments.

The remainder of this paper is organized as follows. Section II describes DARIO’s main functional blocks. Section III defines the system model. Section IV introduces the proposed SNC model. Section V details the UE-RIS assignment and the proposed heuristics for solving it. Section VI evaluates DARIO’s performance, considering a scenario where UEs generate traffic following a Poisson distribution. In contrast, Section VII assesses DARIO’s performance based on real traffic traces previously collected from a major mobile operator. Finally, Section VIII summarizes the main conclusions.

II. DARIO FRAMEWORK

This section presents the proposed DARIO framework. First, it describes the framework architecture, its main functional entities, and the logical interfaces that enable interaction with existing O-RAN control components. Then, the orchestration workflow of DARIO is detailed across the Non-RT and Near-RT control loops. Subsequently, the main implementation challenges for practical integration within O-RAN are analyzed. Next, control-loop latency and signaling overhead are

examined to assess the feasibility of the proposed orchestration within the O-RAN temporal hierarchy. Finally, the scalability of DARIO is discussed, outlining how DARIO can be extended to multi-cell environments.

A. Architecture, Functionalities, and Interfaces

DARIO, illustrated in Fig. 1, is a delay and reliability aware orchestration framework designed to dynamically coordinate multiple RISs within the O-RAN ecosystem. Its main function is to determine, for each UE and at every scheduling slot, which single RIS should be configured to provide this UE the most favorable propagation conditions, together with the corresponding radio resource allocation. The selected RIS configuration is dynamically updated across slots, allowing each UE to be served by different RISs over time depending on the predicted channel and traffic states. The objective of DARIO is to minimize the per-packet transmission delay experienced by each UE while ensuring the required reliability level across all scheduling intervals.

From an architectural perspective, DARIO operates as a separate orchestration module colocated with the Service Management and Orchestration (SMO) and fully interoperable with existing O-RAN components. It remains logically distinct from the Non-RT RIC to preserve compatibility with current specifications, while exploiting its analytics as input and providing optimized configuration policies as output to the Near-RT control layer. This design enables seamless integration of DARIO into the O-RAN management hierarchy, relying on the standard interfaces for coordination and introducing dedicated extensions to support RIS-specific control. Internally, DARIO comprises two main entities:

SNC Module: Estimates the statistical delay bound W for each feasible UE–RIS configuration and radio resource allocation using SNC principles. The packet transmission

delay w in the radio interface is modeled as a random variable constrained by $\Pr[w > W] \leq \varepsilon$, where ε is the maximum allowed violation probability.

RIS Scheduler: Iteratively interacts with the SNC module to evaluate candidate UE–RIS mappings and RB allocations. It selects the configuration that minimizes the resulting delay bound W while ensuring $W \leq W^{\text{th}}$ for each UE and satisfying the reliability constraint $\Pr[w > W] \leq \varepsilon$.

To enable full integration with O-RAN control entities, DARIO introduces three logical interfaces that extend those existing from O-RAN specifications:

RO1: connects DARIO with the Non-RT RIC. It carries aggregated analytics produced by Non-RT *rApps*, including predicted traffic load per UE, expected channel quality indicators, and radio resource utilization.

RO2: links DARIO with the Near-RT RIC. It conveys the optimization outcomes related to radio resource allocation and scheduling policies executed by *xApps* in the Near-RT domain.

RO3: interfaces DARIO with local RIS controllers to deliver the specific configuration states computed by the RIS scheduler, which are then translated by the RIS controllers into element-level phase updates for each RIS.

The analytics consumed through RO1 originate from Non-RT *rApps* such as the *Traffic Estimator*, *Channel Capacity Estimator*, and *RB Utilization* modules. These applications collect raw performance data from the network through the O1 interface and produce aggregated indicators that serve as input for DARIO during each optimization cycle. The outputs generated by DARIO are transmitted through RO2 to the Near-RT RIC, which manages radio resource allocation, and through RO3 to local RIS controllers, which apply the selected configuration states to their corresponding surfaces at the element level.

B. Orchestration Workflow within O-RAN Control Loops

The operation of DARIO follows the hierarchical control loops defined by the O-RAN architecture, where optimization and actuation occur at different temporal scales. Two time parameters are defined: the assignment period $i \in \mathcal{I}$ of duration I_{time} , typically a few seconds, and the scheduling period $t \in \mathcal{T}_i$ of duration $T_{\text{time}} = I_{\text{time}}/|\mathcal{T}_i|$, aligned with Near-RT timescales between 10 and 1000 ms. During each assignment period $i \in \mathcal{I}$, DARIO executes its optimization process, while the corresponding UE–RIS configurations are enforced at every scheduling period $t \in \mathcal{T}_i$.

At the **Non-RT control loop**, DARIO performs the end-to-end orchestration procedure: 1) Collects aggregated analytics from the Non-RT RIC through RO1, including predicted traffic demand per UE, expected channel quality, and radio resource utilization generated by dedicated *rApps*. 2) Executes an iterative process between the SNC module and the RIS scheduler to compute, for each scheduling period $t \in \mathcal{T}_i$, the optimal UE–RIS association that minimizes latency under reliability constraints, while defining a single RB allocation fixed throughout the current assignment period $i \in \mathcal{I}$. 3) Distributes the orchestration results to the corresponding control entities: the radio resource allocation policies are transmitted to the Near-RT RIC through RO2, and the RIS configuration plan




Fig. 1. DARIO integrated within the O-RAN architecture.

for each scheduling period is delivered directly to the local RIS controllers through RO3.

At the **Near-RT control loop**, the precomputed decisions are applied as follows: 1) The Near-RT RIC enforces the radio resource allocation and scheduling policies received from DARIO by means of its *xApps*. 2) Each RIS controller applies the corresponding phase-state settings to its RIS devices via internal control links. These configurations are updated dynamically at the beginning of every scheduling period $t \in \mathcal{T}_i$ of duration T_{time} . 3) The resulting RIS phase profiles create controlled LoS paths between the selected UEs and serving BSs, adapting over time to statistically maintain the latency and reliability levels defined by DARIO.

In this hierarchical process, DARIO performs optimization at the Non-RT scale, while Near-RT entities enforce RIS and scheduling configurations at millisecond granularity. Although PIN-diode RIS elements can switch in sub-microseconds [24], the effective reconfiguration latency is dominated by channel estimation, computation, and O-RAN signaling [7], making updates feasible only at Near-RT intervals.

C. Implementation Challenges for O-RAN Integration

DARIO is designed to operate within the O-RAN framework, enabling hierarchical delay-aware RIS orchestration across multiple time scales. However, real-world integration introduces practical challenges that must be addressed for reliable operation. Below, we outline key issues:

Management Plane Latency: Although RIS devices support microsecond-scale switching [24], the end-to-end management pipeline, including Non-RT and Near-RT RICs as well as local RIS Controllers, may introduce delays due to processing, queuing, and transmission. This challenge concerns ensuring that telemetry collection, Non-RT optimization, and configuration dissemination comply with the temporal hierarchy defined by O-RAN. Our study demonstrates that DARIO meets the execution-time requirements of the Non-RT control loop, but the design and dimensioning of the management network responsible for telemetry exchange and configuration delivery lie beyond the scope of this work. These mechanisms must be properly engineered in real implementations to sustain the required control-loop timing.

Interface Compatibility: Although DARIO leverages existing O-RAN interfaces (e.g., O1) and conceptually extends them through the proposed RO1, RO2, and RO3 links to support RIS-specific control, their practical realization would require either vendor-specific implementations or future standardization efforts coordinated by the corresponding O-RAN Working Groups to ensure interoperability and compliance across different equipment vendors.

D. Quantitative Considerations on Control-Loop Latency and Signaling Overhead

The operation of DARIO adheres to the hierarchical control loops defined by the O-RAN architecture, where orchestration is performed at the Non-RT level and enforcement takes place at Near-RT time scales. The feasibility of the proposed framework is governed by the time-critical orchestration loop executed by DARIO at each assignment period of duration

I_{time} , which must satisfy the timing constraint $T_{DARIO} + T_{signal} + T_{act} \leq I_{time}$, where T_{DARIO} denotes the execution time of the DARIO optimization process, T_{signal} accounts for inter-layer signaling latency, and T_{act} represents the Near-RT enforcement latency.

Traffic and channel-related metrics are generated at fine time granularity within the RAN but are exposed to the Non-RT domain through the O1 interface in an aggregated and asynchronous manner. These analytics are processed by Non-RT *rApps* and reused across multiple assignment periods to derive statistical estimations and predictions. As a result, telemetry acquisition is decoupled from the time-critical orchestration loop and does not lie on its critical timing path.

The assignment period I_{time} is determined by the dynamics of the propagation environment and the mobility of the UEs. For pedestrian mobility scenarios, where the average channel conditions evolve slowly over time, I_{time} can be assumed to be on the order of a few seconds, e.g., $I_{time} \approx 2$ s.

The inter-layer signaling latency T_{signal} is associated with the dissemination of aggregated orchestration decisions through the proposed RO2 and RO3 interfaces. These interfaces operate at the same functional level as the A1 interface in O-RAN and are used for policy and configuration dissemination. The signaling latency is therefore expected to lie in the range $T_{signal} \in [10, 100]$ ms. The Near-RT enforcement latency T_{act} corresponds to the effective application of orchestration decisions in the RAN. It includes *xApp* execution and enforcement through the E2 interface. This process involves computation and synchronization at Near-RT time scales. As a result, the actuation latency can be bounded as $T_{act} \in [30, 150]$ ms.

Considering the most restrictive case within the above ranges, i.e., $T_{signal} = 100$ ms and $T_{act} = 150$ ms, and assuming a pedestrian mobility scenario with $I_{time} \approx 2$ s, the execution time of the DARIO optimization must satisfy $T_{DARIO} \leq I_{time} - (T_{signal} + T_{act}) \approx 1.75$ s. This inequality defines the execution-time threshold that DARIO must satisfy, which is shown to be met in Section VI-D.

E. Scalability of DARIO in Multi-Cell Environments

The current design of DARIO centralizes the coordination of multiple RISs within a single BS domain. Extending its operation to multi-cell networks could be achieved by deploying one DARIO instance per cell, each responsible for the RISs and UEs served by its BS. The orchestration logic would remain identical, as each instance would perform the same delay and reliability aware optimization within its own coverage area.

Two configurations would be possible for analytics provisioning. In moderately sized urban networks, a single set of Non-RT *rApps*, i.e., *Traffic Estimator*, *Channel Capacity Estimator*, *RB Utilization*, would generate per-cell telemetry and feed all DARIO instances through indexed data streams. In larger city-wide or regional deployments, these *rApps* would also be instantiated per zone or per cell to reduce telemetry transport delay and better exploit local monitoring sources. In both cases, each DARIO instance would operate autonomously, without requiring coordination among cells.

The physical placement of DARIO instances may be adapted to the scale and latency constraints of the network. For compact deployments, all instances may run as parallel

processes within the same Non-RT RIC. For large-scale urban networks, instances can be distributed geographically, for example at edge nodes closer to the corresponding BSs. At the logical level, the SMO and Non-RT RIC remain single and global, providing unified management and analytics while supporting multiple DARIO instances in parallel.

III. SYSTEM MODEL

We consider the UL operation of a Single-Input Multiple-Output (SIMO) system in a dense urban environment, providing delay-critical connectivity to a set $\mathcal{U}$ of UEs. Each UE $u \in \mathcal{U}$ generates packets with arbitrary size and arrival distribution. The UEs are equipped with a single antenna, while the BS has N_{ant} antennas, operating in the sub-6 GHz band. Within the coverage area, a set $\mathcal{R}$ of RIS devices are deployed, each composed of L passive elements. The RIS configuration during *scheduling period* $t \in \mathcal{T}_i$ is described by the phase-shift matrix $\Phi_{r,t} = \text{diag}(\kappa_{r,1,t}e^{j\phi_{r,1,t}}, \dots, \kappa_{r,L,t}e^{j\phi_{r,L,t}})$, where $\kappa_{r,l,t} \in [0, 1]$ and $\phi_{r,l,t} \in [0, 2\pi)$ are the amplitude and phase-shift of element $l \in [1, L]$ in RIS $r \in \mathcal{R}$. Each RIS r has a local RIS Controller that applies $\Phi_{r,t}$ according to the configurations provided by DARIO at the beginning of each *scheduling period* $t \in \mathcal{T}_i$.

Table I summarizes the main notation used in the system model and subsequent analysis.

A. Radio Resource Scheduling Model

The serving Orthogonal Frequency-Division Multiple Access (OFDMA) cell has $N_{\text{cell}}^{\text{RB}} = \left\lfloor \frac{B}{N_{\text{sc}}\Delta_f}(1 - OH) \right\rfloor$ Resource Blocks (RBs), where B is bandwidth, N_{sc} the subcarriers per RB, Δ_f the subcarrier spacing, and OH the control signaling overhead. RB allocation among active UEs follows a Weighted Round Robin (WRR) policy¹, enabling proportional RB assignment at each *assignment period* $i \in \mathcal{I}$ based on pre-assigned weights $\pi_{u,i}$, normalized such that $\sum_{u \in \mathcal{U}} \pi_{u,i} = 1$. Each UE receives $N_u^{\text{RB}} = N_{\text{cell}}^{\text{RB}} \pi_{u,i}$ RBs. While exploring optimal RB scheduling algorithms is beyond this paper's scope, see e.g., [25], WRR enables differentiation of resource assignments according to UE's delay bound and reliability requirements.

The BS adapts the Modulation and Coding Scheme (MCS) per UE considering the instantaneous channel quality, yielding spectral efficiencies η_m bps/Hz indexed by $m \in [1, N_c]$, where N_c is the total number of available MCSs. Each MCS is valid over an Signal-to-Noise Ratio (SNR) range $[\gamma_m^{\min}, \gamma_m^{\max}]$. For each UE $u \in \mathcal{U}$, $p_{u,m,n}$ denotes the probability of achieving η_m in RB $n \in [1, N_{\text{cell}}^{\text{RB}}]$ during *scheduling period* $t \in \mathcal{T}_i$, reflecting channel quality and impacting transmission delay including buffering. Note that the values of η_m , $\gamma_m^{\min}$, and $\gamma_m^{\max}$ follow the 3GPP TS 38.214 specification [26]. Two scenarios are considered in a *scheduling period* $t \in \mathcal{T}_i$, $\forall i \in \mathcal{I}$: *Scenario 1 (S1)*: the UE is not assigned to any RIS device, due to RIS occupancy by other UE or lack of LoS with any RIS. *Scenario 2 (S2)*: the UE is assigned to a RIS, establishing a dynamically reconfigurable LoS to the BS via the RIS.

¹Note that O-RAN is assumed to execute this WRR policy dynamically in the Near-RT RIC via an *xApp*, while DARIO operates at the Non-RT level and uses this assumption to infer the expected share of radio resources per UE for its delay bound and reliability estimations.

The following subsections detail the computation of $p_{u,m,n}$ for scenarios *S1* and *S2*. Our analysis assumes the adoption of Inter-Cell Interference Cancellation (ICIC) techniques to mitigate UL interference between UEs in different cells [27].

B. Channel Model Scenario 1: Non-RIS Assignment

When DARIO does not assign a RIS to UE u , RIS devices act as passive scatterers. The received flat-fading signal at the BS on RB n is $\mathbf{y}_{u,n}^{\text{S1}} = \sqrt{P_{u,n}^{\text{tx}} P_{u,n}^{\text{pl,S1}}} \mathbf{h}_{u,n}^{\text{S1}} s_{u,n} + \mathbf{n}_o$, where $s_{u,n}$ is the transmitted symbol, $P_{u,n}^{\text{tx}}$ its power, and $P_{u,n}^{\text{pl,S1}}$ the path loss between UE u and the BS. The Rayleigh fading vector $\mathbf{h}_{u,n}^{\text{S1}} \in \mathbb{C}^{N_{\text{ant}} \times 1}$ has i.i.d. entries $\sim \mathcal{CN}(0, 1)$, and $\mathbf{n}_o \sim \mathcal{CN}(0, N_o \mathbf{I}_{N_{\text{ant}}})$ is AWGN.

Assuming the BS applies Maximal Ratio Combining (MRC), the received vector $\mathbf{y}_{u,n}^{\text{S1}}$ is projected onto the conjugate of the channel vector $\mathbf{h}_{u,n}^{\text{S1}}$, assuming perfect CSI² at the BS, in order to maximize the output SNR. This results in $r_{u,n}^{\text{S1}} = \sqrt{P_{u,n}^{\text{tx}} P_{u,n}^{\text{pl,S1}}} \|\mathbf{h}_{u,n}^{\text{S1}}\|^2 s_{u,n} + \mathbf{h}_{u,n}^{\text{S1},H} \mathbf{n}_o$, with instantaneous SNR $\gamma_{u,n}^{\text{S1}} = \bar{\gamma}_{u,n}^{\text{S1}} \|\mathbf{h}_{u,n}^{\text{S1}}\|^2$, where $\bar{\gamma}_{u,n}^{\text{S1}} = P_{u,n}^{\text{tx}} P_{u,n}^{\text{pl,S1}} / N_o$.

Given the i.i.d. $\mathcal{CN}(0, 1)$ structure of $\mathbf{h}_{u,n}^{\text{S1}}$, the squared norm $\|\mathbf{h}_{u,n}^{\text{S1}}\|^2$ follows a chi-squared distribution with $2N_{\text{ant}}$ degrees of freedom. Therefore, $\gamma_{u,n}^{\text{S1}}$ follows a Gamma distribution with Probability Density Function (PDF):

$$f_{\gamma_u}^{\text{S1}}(\gamma) = \frac{1}{\bar{\gamma}_u^{\text{S1}}} \cdot \frac{1}{(N_{\text{ant}} - 1)!} \left(\frac{\gamma}{\bar{\gamma}_u^{\text{S1}}} \right)^{N_{\text{ant}} - 1} \exp\left(-\frac{\gamma}{\bar{\gamma}_u^{\text{S1}}}\right), \quad (1)$$

The probability of achieving a target spectral efficiency η_m , where $m \in [1, N_c]$, is:

$$p_{u,m}^{\text{S1}} = F_{\Gamma}(\gamma_m^{\max}; N_{\text{ant}}, \bar{\gamma}_u^{\text{S1}}) - F_{\Gamma}(\gamma_m^{\min}; N_{\text{ant}}, \bar{\gamma}_u^{\text{S1}}), \quad (2)$$

where $F_{\Gamma}(\cdot; k, \theta)$ is the Gamma Cumulative Distribution Function (CDF) with shape $k = N_{\text{ant}}$ and scale $\theta = \bar{\gamma}_u^{\text{S1}}$.

C. Channel Model Scenario 2: RIS Assignment

DARIO assigns the RIS $r \in \mathcal{R}$ to the UE $u \in \mathcal{U}_r$, where $\mathcal{U}_r$ is the set of UEs with LoS to r . The remaining RIS devices $r' \in \mathcal{R} \setminus \{r\}$ act as passive scatterers and are not explicitly modeled in the channel equations, as their impact is assumed negligible due to random phase alignment. The received signal at the BS on RB n is $\mathbf{y}_{u,r,n}^{\text{S2}} = \sqrt{P_{u,n}^{\text{tx}} P_{u,r,n}^{\text{pl,S2}}} s_{u,n} \sum_{l=1}^L \mathbf{g}_{u,r,l,n} \psi_{r,l,n} \mathbf{h}_{u,r,l,n} + \mathbf{n}_o$, where $P_{u,r,n}^{\text{pl,S2}}$ is the path loss over the cascaded UE-RIS-BS link, and $\psi_{r,l,n} = \kappa_{r,l,n} e^{j\phi_{r,l,n}}$ is the reflection coefficient of the l -th RIS element, with amplitude $\kappa_{r,l,n} \in [0, 1]$ and phase $\phi_{r,l,n} \in [0, 2\pi)$. The term $\mathbf{h}_{u,r,l,n} \in \mathbb{C}$ is the fast-fading coefficient from UE u to element l of RIS r , and $\mathbf{g}_{u,r,l,n} \in \mathbb{C}^{N_{\text{ant}} \times 1}$ is the vector channel from element l to the BS antennas. Both links follow Rician fading due to LOS conditions [30]:

²The accuracy of the CSI depends on the estimation techniques (e.g., [28], [29]) implemented by the Channel Capacity Estimator rApp in the Non-RT RIC. Since CSI is only used to derive long-term service statistics required by the SNC model and is not a decision variable of DARIO, CSI inaccuracies affect the numerical tightness of the estimated delay bounds but do not alter the operation of the DARIO framework nor the relative behavior among compared baseline solutions. Accordingly, our analysis assumes the most favorable case of perfect CSI estimation.

TABLE I
SUMMARY OF NOTATION AND PARAMETERS.

Symbol	Definition	Symbol	Definition
$\mathcal{U}$	Set of UEs	$\mathcal{R}$	Set of RISs
$u \in \mathcal{U}$	UE index	$r \in \mathcal{R}$	RIS index
$i \in \mathcal{I}$	Assignment period index	$t \in \mathcal{T}_i$	Scheduling period index
I_{time}	Assignment period duration [s]	T_{time}	Scheduling period duration = $I_{time}/ \mathcal{T}_i $ [s]
N_{ant}	BS receive antennas	L	Elements per RIS
$\Phi_{r,t}$	RIS- r phase-shift matrix at period t	$\kappa_{r,l,t}$	Amplitude of RIS element l
$\phi_{r,l,t}$	Phase of RIS element l [rad]	B	Phase-quantization bits per element
d	Inter-element spacing at RIS [m]	θ_0	Reflection angle (broadside reference) [rad]
B_{sys}	System bandwidth [Hz]	N_{sc}	Subcarriers per RB
Δ_f	Subcarrier spacing [Hz]	OH	Overhead fraction for control plane in 5G
N_{cell}^{RB}	RBs in cell	N_u^{RB}	RBs assigned to UE u
$\pi_{u,i}$	WRR weight of UE u at period i	t_{slot}	TTI duration [s]
$P_{u,n}^{tx}$	UE u TX power on RB n [W]	N_0	Noise PSD [W/Hz]
$P_{u,n}^{pl,S1}$	Path loss UE-BS in S1	$P_{u,r,n}^{pl,S2}$	Cascaded UE-RIS-BS path loss in S2
$\mathbf{h}_{u,n}^{S1}$	UE-BS fast fading in S1	$\mathbf{g}_{u,r,l,n}$	RIS- l -BS vector channel
$h_{u,r,l,n}$	UE-RIS- l scalar channel	$K_{u,r}, K_r$	Rician factors UE-RIS, RIS-BS
$\bar{\gamma}_{u,n}^{S1}$	Avg. SNR in S1	$\bar{\gamma}_{u,r,n}^{S2}$	Avg. SNR in S2
$\gamma_{u,n}^{S1}$	Instantaneous SNR in S1	$\gamma_{u,r,n}^{S2}$	Instantaneous SNR in S2
η_m	Spectral efficiency of MCS m [bit/s/Hz]	$[\gamma_m^{\min}, \gamma_m^{\max}]$	SNR region of MCS m
$p_{u,m}^{S1}$	$\Pr\{\text{MCS} = m \mid S1\}$	$p_{u,r,m}^{S2}$	$\Pr\{\text{MCS} = m \mid S2 \text{ with RIS } r\}$
ω_u^{S1}	Time fraction in S1 for UE u	$\omega_{u,r}^{S2}$	Time fraction in S2 using RIS r
$x_{u,r,t}$	Binary: RIS r assigned to u at t	G_{loss}	Avg. amplitude reduction
λ	Noncentrality parameter in S2	$Q_{N_{ant}}(\cdot, \cdot)$	Marcum- Q function
$A_u(\tau, t)$	Arrival process [bit]	$S_u(\tau, t)$	Service process [bit]
$B_u(i)$	Arrivals at TTI i [bit]	$C_u(j)$	Served bits at TTI j [bit]
$MA_u(\theta)$	MGF of arrivals	$M_{S_u}(-\theta)$	Negative MGF of service
ρ_A, σ_A	EBB rate and burst params.	ρ_S, σ_S	EBF rate and burst params.
ε	Delay violation target	W_u	Delay bound of UE u [s]
W_u^{th}	Target delay bound [s]	δ, θ	SNC tuning parameters
f_{obj}	Objective function in Eq. (31)		

$h_{u,r,l,n} = \sqrt{\frac{K_{u,r}}{K_{u,r}+1}} \bar{h}_{u,r,l,n} + \sqrt{\frac{1}{K_{u,r}+1}} \tilde{h}_{u,r,l,n}$, with Rician factor $K_{u,r}$, deterministic LOS component $\bar{h}_{u,r,l,n}$, and scattered component $\tilde{h}_{u,r,l,n} \sim \mathcal{CN}(0, 1)$. Similarly, $\mathbf{g}_{u,r,l,n} = \sqrt{\frac{K_r}{K_r+1}} \bar{\mathbf{g}}_{r,l,n} + \sqrt{\frac{1}{K_r+1}} \tilde{\mathbf{g}}_{r,l,n}$, where K_r is the Rician factor, $\bar{\mathbf{g}}_{r,l,n}$ the LOS component, and $\tilde{\mathbf{g}}_{r,l,n} \sim \mathcal{CN}(\mathbf{0}, \mathbf{I}_{N_{ant}})$ the scattered term. We assume the RIS elements introduce uncorrelated scattering in both the UE-RIS and RIS-BS links. Accordingly, the components of $\tilde{h}_{u,r,l,n}$ and $\tilde{\mathbf{g}}_{r,l,n}$ are modeled as i.i.d. across elements. This assumption simplifies the statistical characterization of the cascaded channel.

RIS Phase Quantization and Frequency Response Effects. Each RIS element applies a quantized phase shift $\phi_{r,l,n}$ using a uniform B -bit phase shifter [31], constant across all RBs assigned to UE u . Thus, phase values lie in the set $\phi_{r,l,n} \in \{0, \Delta\phi, \dots, (2^B - 1)\Delta\phi\}$, $\Delta\phi = \frac{2\pi}{2^B}$. The RIS controller computes a continuous ideal phase $\phi_{r,l,n_0}^{\text{ideal}}$ based on the central RB n_0 , which is then quantized as $\hat{\phi}_{r,l,n} = \arg \min_{\phi \in \{0, \Delta\phi, \dots\}} |\phi - \phi_{r,l,n_0}^{\text{ideal}}|$.

Due to channel frequency selectivity, the ideal phase at RB n differs by a frequency-induced deviation $\xi_{r,l,n}$: $\phi_{r,l,n}^{\text{ideal}} = \phi_{r,l,n_0}^{\text{ideal}} + \xi_{r,l,n}$. The total phase error is then

$$\epsilon_{r,l,n} = \hat{\phi}_{r,l,n} - \phi_{r,l,n}^{\text{ideal}} = \underbrace{\hat{\phi}_{r,l,n} - \phi_{r,l,n_0}^{\text{ideal}}}_{\epsilon_{r,l,n}^{\text{quant}}} - \underbrace{\xi_{r,l,n}}_{\epsilon_{r,l,n}^{\text{freq}}}, \quad (3)$$

with quantization $\epsilon_{r,l,n}^{\text{quant}}$ and frequency $\epsilon_{r,l,n}^{\text{freq}}$ components assumed independent, as they arise from separate mechanisms:

the former due to discrete phase control, and the latter arising from the frequency-dependent phase progression across the RIS aperture, which produces a deterministic spatial phase rotation for subcarriers away from the central frequency f_c . These phase errors impair the coherent combining gain of the RIS. The expected complex reflection gain for a single element is $\mathbb{E}[e^{j\epsilon_{r,l,n}}] = \mathbb{E}[e^{j\epsilon_{r,l,n}^{\text{quant}}}] \cdot \mathbb{E}[e^{-j\xi_{r,l,n}}]$.

Assuming $\epsilon_{r,l}^{\text{quant}} \sim \mathcal{U}[-\frac{\Delta\phi}{2}, \frac{\Delta\phi}{2}]$ [32], we get $\mathbb{E}[e^{j\epsilon_{r,l}^{\text{quant}}}] = \text{sinc}(\Delta\phi/2)$. The second term accounts for the deterministic frequency-dependent phase deviation among subcarriers caused by the fixed phase configuration of the RIS at the central frequency f_c . For the m -th subcarrier, this deviation is

$$\xi_{r,l,m} = \frac{2\pi(f_m - f_c)d \sin \theta_0}{c} (l - 1), \quad (4)$$

where d is the inter-element spacing, θ_0 the reflection angle, c the speed of light, and l the element index along the RIS aperture. The coherent gain over all $12N_u^{\text{RB}}$ subcarriers assigned to the UE is then given by

$$\mathbb{E}[e^{-j\xi_{r,l,n}}] \approx \frac{1}{12N_u^{\text{RB}}} \sum_{m=1}^{12N_u^{\text{RB}}} \left| \frac{\sin(\frac{L\varsigma_m}{2})}{L \sin(\frac{\varsigma_m}{2})} \right|, \quad (5)$$

where $\varsigma_m = \frac{2\pi(f_m - f_c)d \sin \theta_0}{c}$.

Combining both effects, the average amplitude reduction per RIS element is

$$G_{\text{loss}} = \text{sinc}\left(\frac{\Delta\phi}{2}\right) \cdot \frac{1}{12N_u^{\text{RB}}} \sum_{m=1}^{12N_u^{\text{RB}}} \left| \frac{\sin\left(\frac{L\zeta_m}{2}\right)}{L\sin\left(\frac{\zeta_m}{2}\right)} \right|. \quad (6)$$

Received Signal and SNR with Phase Errors. The phase errors impair the coherent superposition of the reflected paths. The effective combining vector at the BS is $\mathbf{z}_{u,r,n} = \sum_{l=1}^L \mathbf{g}_{u,r,l,n} h_{u,r,l,n} e^{j\epsilon_{r,l,n}}$, and the received signal vector becomes $\mathbf{y}_{u,r,n}^{S2} = \sqrt{P_{u,n}^{\text{tx}} P_{u,r,n}^{\text{pl},S2}} s_{u,n} \mathbf{z}_{u,r,n} + \mathbf{n}_0$, and the scalar output after MRC combining is $r_{u,r,n} = \mathbf{z}_{u,r,n}^H \mathbf{y}_{u,r,n}^{S2} = \sqrt{P_{u,n}^{\text{tx}} P_{u,r,n}^{\text{pl},S2}} s_{u,n} \|\mathbf{z}_{u,r,n}\|^2 + \mathbf{z}_{u,r,n}^H \mathbf{n}_0$.

The instantaneous SNR is $\gamma_{u,r,n}^{S2} = \bar{\gamma}_{u,r,n}^{S2} \|\mathbf{z}_{u,r,n}\|^2$, with $\bar{\gamma}_{u,r,n}^{S2} = P_{u,n}^{\text{tx}} P_{u,r,n}^{\text{pl},S2} / N_0$ the average SNR.

Probability of Selecting MCS Index. Due to the LOS components in both UE-RIS and RIS-BS links, $\mathbf{z}_{u,r,n}$ comprises independent Rician-distributed terms, and thus $\|\mathbf{z}_{u,r,n}\|^2$ follows a non-central chi-square distribution with $2N_{\text{ant}}$ degrees of freedom. Its non-centrality parameter is $\lambda = L^2 \cdot \frac{K_{u,r}}{K_{u,r}+1} \cdot \frac{K_r}{K_r+1} \cdot G_{\text{loss}}^2$, where $K_{u,r}$ and K_r are the Rician factors. We adopt $K_{u,r} = 3$ dB and $K_r = 6$ dB, representing typical partial and stable LOS conditions for UE-RIS and RIS-BS paths, respectively.

The PDF of $\gamma_{u,r,n}^{S2}$ is

$$f_{\gamma_{u,r,n}^{S2}}(\gamma) = \frac{1}{\bar{\gamma}_{u,r,n}^{S2}} \left(\frac{\gamma}{\lambda \bar{\gamma}_{u,r,n}^{S2}} \right)^{\frac{N_{\text{ant}}-1}{2}} \cdot \exp\left(-\frac{\gamma + \lambda \bar{\gamma}_{u,r,n}^{S2}}{\bar{\gamma}_{u,r,n}^{S2}}\right) I_{N_{\text{ant}}-1}\left(\frac{2\sqrt{\lambda\gamma}}{\bar{\gamma}_{u,r,n}^{S2}}\right), \quad (7)$$

where $I_{N_{\text{ant}}-1}(\cdot)$ is the modified Bessel function of the first kind. The CDF is $F_{\gamma_{u,r,n}^{S2}}(\gamma) = 1 - Q_{N_{\text{ant}}}\left(\sqrt{\lambda}, \sqrt{\frac{\gamma}{\bar{\gamma}_{u,r,n}^{S2}}}\right)$, where $Q_{N_{\text{ant}}}(\cdot, \cdot)$ is the Marcum-Q function. Thus, the probability of selecting MCS index m , defined over $[\gamma_m^{\min}, \gamma_m^{\max}]$, is

$$p_{u,m}^{S2} = Q_{N_{\text{ant}}}\left(\sqrt{\lambda}, \sqrt{\frac{\gamma_m^{\min}}{\bar{\gamma}_{u,r,n}^{S2}}}\right) - Q_{N_{\text{ant}}}\left(\sqrt{\lambda}, \sqrt{\frac{\gamma_m^{\max}}{\bar{\gamma}_{u,r,n}^{S2}}}\right). \quad (8)$$

IV. PER-UE DELAY BOUND MODELING VIA STOCHASTIC NETWORK CALCULUS

This section derives the per-UE delay bound W using the SNC framework for a multi-RIS deployment scenario under dynamic traffic and channel conditions. We briefly revisit the fundamental principles of SNC and the procedure to derive arrival and service envelopes and the corresponding delay bound. The traffic model is characterized from empirical observations, while the service process is formulated from the cell capacity under dynamic RIS operation. In contrast to conventional SNC formulations assuming a stationary and ergodic service process, the proposed analysis extends the framework to configuration-dependent and time-varying conditions induced by RIS reconfiguration, by modeling the service

process as an explicit function of UE-RIS assignment decisions. This leads to a finite-mixture service characterization whose parameters evolve with the scheduling configuration, enabling direct analytical evaluation of how different RIS configurations reshape the resulting delay bound.

A. Introduction to Stochastic Network Calculus (SNC)

SNC provides an analytical framework for modelling and analyzing the performance of mobile networks, characterized by stochastic packet arrivals and service rates. Let $A(\tau, t)$ and $S(\tau, t)$ denote the cumulative arrival and service processes over the interval $(\tau, t]$, respectively, representing the total number of bits that arrive to, and can be served by, the node during this time. SNC uses the concepts of Exponentially Bounded Burstiness (EBB) and Exponentially Bounded Fluctuation (EBF) [5] to define an upper bound $\alpha(\tau, t)$, known as arrival envelope, and a lower bound, $\beta(\tau, t)$, known as service envelope, as specified in the following

$$\mathbb{P}[A(\tau, t) > \alpha(\tau, t)] \leq \varepsilon_A, \quad \mathbb{P}[S(\tau, t) < \beta(\tau, t)] \leq \varepsilon_S, \quad (9)$$

where ε_A and ε_S are the overflow and deficit profiles [5]. We assume affine functions for $\alpha(\tau, t)$ and $\beta(\tau, t)$ which are defined in Eqs. (10) and (11). The parameters $\rho_A > 0$, $\rho_S > 0$ and $b_A \geq 0$, $b_S \geq 0$ are the rate and burst parameters for these functions. Additionally, $[x]_+$ denotes $\max\{0, x\}$. Finally, $\delta > 0$ is a sample path argument considered by EBB and EBF models [5].

$$\alpha(\tau, t) = (\rho_A + \delta)[t - \tau] + b_A, \quad (10)$$

$$\beta(\tau, t) = (\rho_S - \delta)[t - \tau - b_S/\rho_S]_+. \quad (11)$$

Representing $\alpha(\tau, t)$ and $\beta(\tau, t)$ over a $t - \tau$ axis, we can obtain the delay bound W as the horizontal deviation between these envelopes, as shown in Eq. (12). This deviation occurs when the slope of the service envelope is greater than the slope of the arrival envelope, i.e., $\frac{\partial \beta(\tau, t)}{\partial (t - \tau)} > \frac{\partial \alpha(\tau, t)}{\partial (t - \tau)}$, provided that the stability condition $\rho_S - \delta > \rho_A + \delta$ is satisfied, ensuring that the delay bound remains finite.

$$W = \frac{b_A + b_S}{\rho_S - \delta}. \quad (12)$$

For further details on SNC fundamentals, the reader is referred to [33], [34].

B. Methodology for Delay Bound Derivation with SNC

We summarize the generic methodology introduced in [5, Sections II.A and II.B] to derive the arrival and service envelopes, and subsequently compute the corresponding delay bound:

1) Compute the Moment Generating Functions (MGFs) for the arrival and service processes, i.e., $M_A(\theta)$ and $M_S(-\theta)$. The MGF of a process X is $\mathbb{E}[e^{\theta X}]$ with a tunable parameter θ .

2) Define upper bounds for the MGFs as shown in Eqs. (13) and (14). They are characterized by the rate parameters ρ_A and ρ_S , and the burst parameters σ_A and σ_S . Note that ρ_A and ρ_S match those in Eqs. (10) and (11). Fit the MGFs computed in Step 1 to the exponential bounds defined in Eq. (13) and (14) to extract the parameters ρ_A , σ_A , ρ_S , and σ_S .

$$M_A(\theta) \leq e^{\theta(\rho_A[t - \tau] + \sigma_A)}, \quad (13)$$

$$M_S(-\theta) \leq e^{-\theta(\rho_S[t-\tau]-\sigma_S)}. \quad (14)$$

3) The EBB and EBF models defined in Eq. (9), and the MGFs are directly connected by the Chernoff bound [35]. Based on this, we can obtain b_A and b_S as shown in Eqs. (15) and (16). We equally distribute the target violation probability among ε_A and ε_S , i.e., $\varepsilon_A = \varepsilon_S = \varepsilon/2$.

$$b_A = \sigma_A - \frac{1}{\theta} \left[\ln(\varepsilon_A) + \ln(1 - e^{-\theta\delta}) \right], \quad (15)$$

$$b_S = \sigma_S - \frac{1}{\theta} \left[\ln(\varepsilon_S) + \ln(1 - e^{-\theta\delta}) \right]. \quad (16)$$

4) Combining the expressions for b_A and b_S with the delay bound expression in Eq. (12), the delay bound W can be reformulated as:

$$W = \frac{\sigma_A + \sigma_S - \frac{2}{\theta} \left[\ln\left(\frac{\varepsilon}{2}\right) + \ln(1 - \exp[-\theta\delta]) \right]}{\rho_S - \delta}. \quad (17)$$

SNC models typically yield conservative estimations of the delay bound [5]. Hence, it is necessary to optimize the tuneable parameters θ and δ to minimize W and obtain a bound as tight as possible.

C. UE Traffic Model for the Arrival Process

In a slotted system with Transmission Time Interval (TTI) duration t_{slot} , the cumulative arrival process $A_u(\tau, t)$ for UE u is given by

$$A_u(\tau, t) = \sum_{i=\lceil \tau/t_{slot} \rceil + 1}^{\lfloor t/t_{slot} \rfloor} B_u(i), \quad (18)$$

where $B_u(i)$ represents the number of bits that arrive to the transmission buffer in the i -th TTI.

In this work, we consider that each UE may generate traffic following an arbitrary distribution. For this reason, we assume that the Probability Mass Function (PMF) of $B_u(i)$ can be estimated by using samples of the incoming bits per TTI in the last T_{OBS} TTIs. Specifically, we define the sample vector $\vec{x}_{B_u} = \{b_{u,1}, b_{u,2} \dots b_{u,T_{OBS}}\}$, where $b_{u,i}$ denotes the number of bits that arrive to the transmission buffer in the TTI i for the UE u . We assume the elements of $\vec{x}_{B_u}$ are independent and identically distributed random variables, taken from a sliding window of the most recent T_{OBS} TTIs. Note that in our prior work [36], we demonstrated that $T_{OBS} \geq 4000$ is required to accurately estimate the PMF of the generated traffic. Additionally, $b_{u,i} = \sum_{j=1}^{J_{u,i}} l_j$, where $J_{u,i}$ is the number of packets generated by the UE u in the TTI i and l_j the size of the packet j . Note that the computation of $\vec{x}_{B_u}$ is a task performed by the Traffic Estimator *rApp*.

With these assumptions, the empirical MGF of $B_u(i)$ over the observation window is given by

$$M_{B_u}(\theta) = \frac{1}{T_{OBS}} \sum_{i=1}^{T_{OBS}} \exp[\theta b_{u,i}]. \quad (19)$$

Since the samples of $B_u(i)$ are i.i.d., the MGF of the arrival process $A_u(\tau, t)$ is $M_{A_u}(\theta) = E[M_{B_u}(\theta)^{N_{slot}(\tau, t)}]$, being $N_{slot}(\tau, t) = \lfloor t/t_{slot} \rfloor - \lceil \tau/t_{slot} \rceil$ the number of TTIs fully contained in the interval $(\tau, t]$. By exponentiating the logarithm of the expectation's argument for $M_{A_u}(\theta)$, we get $M_{A_u}(\nu) =$

$E[e^{\nu N_{slot}(\tau, t)}]$, with $\nu = \ln[M_{B_u}(\theta)]$. Since $N_{slot}(\tau, t)$ is a constant, $M_{A_u}(\nu) = e^{\nu(t-\tau)/t_{slot}}$. We can replace ν and get

$$M_{A_u}(\theta) = e^{\left[\ln\left(\frac{1}{T_{OBS}} \sum_{i=1}^{T_{OBS}} \exp[\theta b_{u,i}]\right) \frac{(t-\tau)}{t_{slot}} \right]}. \quad (20)$$

Finally, by equaling the left and right sides of Eq. (13), we obtain $\rho_{A_u}(\theta)$ as Eq. 21 shows. Note that $\sigma_{A_u}(\theta) = 0$.

$$\rho_{A_u}(\theta) = \frac{\ln\left[\frac{1}{T_{OBS}} \sum_{i=1}^{T_{OBS}} \exp[\theta b_{u,i}]\right]}{\theta t_{slot}}. \quad (21)$$

D. Cell Capacity Model for the Service Process

The service process $S_u(\tau, t)$ represents the accumulated capacity provided to UE u and is defined by Eq. (22). $C_u(j)$ is the number of bits served by the cell to UE u in TTI j , as specified in Eq. (23). Additionally, $\eta_{u,n}(j)$ denotes the spectral efficiency (bps/Hz) for UE u in TTI j on RB n . We assume $\eta_{u,n}$ is an i.i.d. random variable across RBs and TTIs, which simplifies the derivation. Although in practice, spectral efficiency may exhibit spatial and temporal correlation, this assumption enables tractable MGF-based analysis under SNC.

$$S_u(\tau, t) = \sum_{j=\lceil \tau/t_{slot} \rceil + 1}^{\lfloor t/t_{slot} \rfloor} C_u(j), \quad (22)$$

$$C_u(j) = \sum_{n=1}^{N_u^{RB}} N_{sc} \Delta_f t_{slot} \eta_{u,n}(j). \quad (23)$$

The negative MGF for $\eta_{u,n}(j)$ is defined in Eq. (24). It depends on the probabilities $\omega_{u,r}^{S1}$ and $\omega_{u,r}^{S2}$, $\forall r \in \mathcal{R}_u$, that characterize the fraction of scheduling periods in which UE u operates without RIS assistance or is served through a specific RIS, respectively, together with the corresponding MCS spectral efficiencies η_m and selection probabilities $p_{u,m}^{S1}$ and $p_{u,m}^{S2}$. This formulation yields a finite-mixture service characterization in which the mixture weights $\omega_{u,r}^{S1}$ and $\omega_{u,r}^{S2}$ are set by the UE-RIS configuration, while CSI only affects the MCS selection probabilities $p_{u,r,m}^{S2}$ and $p_{u,m}^{S1}$.

$$M_\eta(-\theta) = \sum_{r \in \mathcal{R}_u} \omega_{u,r}^{S2} \sum_{m=1}^{N_c} \exp(-\theta \eta_m) p_{u,r,m}^{S2} + \omega_u^{S1} \sum_{m=1}^{N_c} \exp(-\theta \eta_m) p_{u,m}^{S1}. \quad (24)$$

The parameter $\omega_{u,r}^{S2} = (\sum_{t \in \mathcal{T}_i} x_{u,r,t}) / |\mathcal{T}_i|$ is the probability UE u being assigned a specific RIS $r \in \mathcal{R}_u$ during the scheduling period $t \in \mathcal{T}_i$, and $x_{u,r,t} \in \{0, 1\}$ is a scheduling decision variable indicating whether RIS r is assigned to UE u at scheduling period t . Finally, $\omega_u^{S1} = 1 - \sum_{r \in \mathcal{R}_u} \omega_{u,r}^{S2}$. Note that Section V details how $x_{u,r,t}$ is set to minimize the packet transmission delay for all UEs. The MGF of $C_u(j)$ is defined in Eq. (25). We represent $M_C(-\theta)$ in terms of $M_\eta(-\theta)$, leveraging on the MGF properties related to linear combination of independent random variables. Next, by substituting $\ln(M_\eta(-N_{sc} \Delta_f t_{slot} \theta)) = \nu$, we obtain the MGF of the process N_u^{RB} as a function of the free parameter ν .

$$\begin{aligned}
M_C(-\theta) &= E \left[[M_\eta(-N_{sc}\Delta_f t_{slot}\theta)]^{N_u^{RB}} \right] \\
&= E \left[e^{\nu N_u^{RB}} \right] = M_{N_u^{RB}}(\nu).
\end{aligned} \tag{25}$$

Since $M_{N_u^{RB}}(\nu) = e^{N_u^{RB}\nu}$, we can define $M_{C_u}(-\theta)$ as follows

$$\begin{aligned}
M_{C_u}(-\theta) &= \left(\sum_{r \in \mathcal{R}_u} \omega_{u,r}^{S2} \sum_{m=1}^{N_c} e^{(-\theta N_{sc}\Delta_f t_{slot}\eta_m)} \right. \\
&\quad \left. p_{u,r,m}^{S2} + \omega_u^{S1} \sum_{m=1}^{N_c} e^{(-\theta N_{sc}\Delta_f t_{slot}\eta_m)} p_{u,m}^{S1} \right)^{N_u^{RB}}.
\end{aligned} \tag{26}$$

Considering $M_{C_u}(-\theta)$, we can compute the negative MGF of the service process $S_u(\tau, t)$ as the following $M_{S_u}(-\theta) = e^{\ln(M_{C_u}(-\theta)) \frac{(t-\tau)}{t_{slot}}}$. Finally, by equaling the left and right sides of Eq. (14), we obtain $\rho_{S_u}(\theta)$ as follows. Note that $\sigma_{S_u}(\theta) = 0$.

$$\rho_{S_u}(\theta) = \frac{-\ln[M_{C_u}(-\theta)]}{\theta t_{slot}}. \tag{27}$$

E. Delay Bound Estimation for Packet Transmission

Using the results from Sections IV-C and IV-D, we can define the delay bound W_u as a function of the tuneable parameters θ and δ , as specified in Eq. (28). To estimate W_u , we need to solve the following optimization problem:

Problem DELAY_BOUND_PACKET_TRANSMISSION:

$$\min_{\theta, \delta} \quad W_u = \frac{2t_{slot} \left[\ln\left(\frac{\varepsilon_u}{2}\right) + \ln(1 - \exp[-\theta\delta]) \right]}{-\theta t_{slot} \rho_{S_u}(\theta) + \delta \theta t_{slot}}. \tag{28}$$

$$\text{s.t.:} \quad \theta, \delta, \rho_{A_u}(\theta), \rho_{S_u}(\theta) > 0, \tag{29}$$

$$\rho_{S_u}(\theta) - \rho_{A_u}(\theta) > 2\delta. \tag{30}$$

The above-mentioned problem involves a non-convex objective function within a non-convex region, leading to multiple local minima and making exhaustive search computationally infeasible. To address the non-convexity, we adapt a heuristic inspired by [36]. However, unlike [36], which focuses on non-RIS systems, our method supports dynamic multi-RIS configurations and mobile UEs. These aspects introduce new constraints and variability in the delay model that are not considered in prior work.

V. UE-RIS ASSIGNMENT PROBLEM FORMULATION

For each *assignment period* $i \in \mathcal{I}$, DARIO determines in advance a one-to-one UE-RIS assignment for every *scheduling period* $t \in \mathcal{T}_i$. The UE-RIS-RB_ASSIGNMENT problem is formulated as follows.

Problem UE-RIS-RB_ASSIGNMENT:

$$\min_{x_{u,r,t}, N_u^{RB}} \quad f_{obj} = \max_{u \in \mathcal{U} \setminus \mathcal{U}_{\mathcal{R}}} \left\{ \frac{W_u[\omega_u]}{W_u^{th}} \right\} + \max_{u \in \mathcal{U}_{\mathcal{R}}} \left\{ \frac{W_u[\omega_u]}{W_u^{th}} \right\}, \tag{31}$$

$$\text{s.t.:} \quad \sum_{t \in \mathcal{T}_i} \sum_{r \notin \mathcal{R}_u} x_{u,r,t} = 0, \quad \forall u \in \mathcal{U}_{\mathcal{R}}, \tag{32}$$

$$\sum_{r \in \mathcal{R}} x_{u,r,t} \leq 1, \quad \forall t \in \mathcal{T}_i, \quad \forall u \in \mathcal{U}_{\mathcal{R}}, \tag{33}$$

$$\sum_{u \in \mathcal{U}_{\mathcal{R}}} x_{u,r,t} \leq 1, \quad \forall t \in \mathcal{T}_i, \quad \forall r \in \mathcal{R}, \tag{34}$$

$$\sum_{u \in \mathcal{U}} N_u^{RB} = N_{cell}^{RB}. \tag{35}$$

In this problem, Eq. (31) minimizes the sum of two worst-case delay ratios: the first term considers the maximum ratio of the experienced delay bound W_u to the target delay bound W_u^{th} over the UEs without LoS to any RIS (i.e., $u \in \mathcal{U} \setminus \mathcal{U}_{\mathcal{R}}$), while the second term considers the maximum ratio over the subset $\mathcal{U}_{\mathcal{R}}$ of UEs with LoS to at least one RIS. Note that W_u is computed with respect to a target violation probability $\varepsilon_u \forall u \in \mathcal{U}$. Focusing on a single UE $u \in \mathcal{U}$, the experienced delay bound $W_u[\omega_u]$ depends on the probability vector $\omega_u = \{\omega_u^{S1}, \omega_{u,r}^{S2}\} \forall r \in \mathcal{R}$, as defined by Eqs. (26), (27) and (28). This dependence is explicit and analytical, enabling the optimization variables to directly influence the SNC service process through the configuration-dependent mixture weights. In turn, these probabilities depend on the binary and integer decision variables $x_{u,r,t}$ and N_u^{RB} , respectively, which are constrained by the following: Constraint (32) enforces that any UE u lacking LoS to RIS r , i.e., $r \notin \mathcal{R}_u$, cannot be assigned that RIS; Constraint (33) enforces that one RIS is assigned to UE u in a *scheduling period* $t \in \mathcal{T}_i$; Constraint (34) enforces that each RIS $r \in \mathcal{R}$ serves a single UE in a *scheduling period* $t \in \mathcal{T}_i$; and Constraint (35) enforces that the total number of RBs allocated across all UEs equals the available RBs in the cell.

The objective function in Eq. (31) depends on $W_u \forall u \in \mathcal{U}$, and involves a non-linear relationship with the weights $\{\omega_u^{S1}, \omega_{u,r}^{S2}\} \forall r \in \mathcal{R}$, and in turn with the decision variables $x_{u,r,t} \forall r \in \mathcal{R}, \forall t \in \mathcal{T}_i$ and the integer RB allocation variables $N_u^{RB} \forall u \in \mathcal{U}$, as shown in Eqs. (26), (27) and (28), making the UE-RIS-RB_ASSIGNMENT problem a NIP problem [37]. Given (a) the nonlinearity of the objective function, (b) the large number of decision variables, and (c) the need to compute the delay bound $W_u \forall u \in \mathcal{U}$ by invoking a heuristic (see Section IV-E) for every possible UE-RIS assignment (exhaustive search) at each *scheduling period* $t \in \mathcal{T}_i$, the problem becomes computationally prohibitive. The total number of UE-RIS assignment combinations is $|\mathcal{C}|^{|\mathcal{T}_i|}$, where $|\mathcal{C}|$ is the number of possible assignments within a single *scheduling period*. Because RB allocation to UEs is performed once and remains fixed across all *scheduling periods* $\mathcal{T}_i$, the total number of possible RB allocation is $\binom{N_{cell}^{RB}-1}{|\mathcal{U}|-1}$. Therefore, the total number of feasible joint configurations is $|\mathcal{C}|^{|\mathcal{T}_i|} \times \binom{N_{cell}^{RB}-1}{|\mathcal{U}|-1}$, highlighting the exponential growth and computational intractability of exact optimization.

To reduce computational complexity, we propose a two-stage heuristic: (1) RB allocation among UEs assuming an

Algorithm 1: RB allocation with equiprobable UE-RIS assignment

```

1 Inputs:  $W_u^{th}, \forall u \in \mathcal{U}$ , initial RB allocation
    $N_u^{RB} = N_{cell}^{RB}/|\mathcal{U}|$ ;
2 Initialization:  $f_{obj}^{(old)} = \infty, f_{obj}^{(new)} = \infty, Cont = False$ ;
3 Compute  $W_u$  and  $f_{obj}^{(old)}$  = Eq. (31);
4 while not  $Cont$  do
5   Select  $u^{max} = \arg \max_{u \in \mathcal{U}} (W_u/W_u^{th})$ ;
6   Select  $u^{min} = \arg \min_{u \in \mathcal{U} \text{ s.t. } N_u^{RB} > 1} (W_u/W_u^{th})$ ;
7   Update RB allocation:  $N_{u^{max}}^{RB} = N_{u^{max}}^{RB} + 1$  and
      $N_{u^{min}}^{RB} = N_{u^{min}}^{RB} - 1$ ;
8   Recompute  $W_u$  and  $f_{obj}^{(new)}$ ;
9   if  $f_{obj}^{(new)} < f_{obj}^{(old)}$  then
10     $f_{obj}^{(old)} = f_{obj}^{(new)}$ ;
11  else
12    Revert the last RB exchange:  $N_{u^{max}}^{RB} = N_{u^{max}}^{RB} - 1$ 
    and  $N_{u^{min}}^{RB} \leftarrow N_{u^{min}}^{RB} + 1$ ;
13     $Cont = True$ ;
14  end
15 end
16 return:  $N_u^{RB}$ ;
```

equiprobable UE-RIS assignment (i.e., no optimized RIS association), as performed in Algorithm 1; and (2) dynamic UE-RIS association based on the fixed RB allocation obtained in the first stage, aiming to improve delay performance. These stages correspond to Algorithm 1 and Algorithm 2, respectively, described below.

Algorithm 1. It iteratively redistributes RBs among UEs to improve the global delay performance. At each step, the UE with the highest normalized delay ratio receives one RB from the UE with the lowest ratio, as long as the donor UE retains at least one RB (lines 5-7). The objective function is updated after each exchange (line 8), and the process continues until no further improvement is observed (lines 9-13).

Algorithm 1 has a worst-case complexity of $\mathcal{O}(N_{cell}^{RB} \cdot |\mathcal{U}|)$. Each iteration scans all UEs to find u^{max} and u^{min} , and recomputes the objective function with cost $\mathcal{O}(|\mathcal{U}|)$. Only RB exchanges reducing the objective are accepted; since the objective decreases discretely and is bounded below, convergence occurs after a finite number of steps. Conservatively, we upper bound iterations by N_{cell}^{RB} for this complexity estimation.

Algorithm 2. It takes as input target delay bounds W_u^{th} and violation probabilities $\varepsilon_u \forall u \in \mathcal{U}_{\mathcal{R}}$ (line 1) and initializes the state variables (line 2). It begins with a random UE-RIS assignment (line 3), and it continues with the calculation of the weight vector ω_u and the delay bound W_u for each UE $u \in \mathcal{U}_{\mathcal{R}}$, determining the initial objective function value (line 4). The algorithm then iteratively aims to reduce the objective function (lines 5-32). In each iteration, the UE u' with the highest ratio $W_{u'}/W_{u'}^{th}$ is selected (line 6), along with the UE u'' with the lowest ratio $W_{u''}/W_{u''}^{th}$ among those experiencing LoS with at least one common RIS device (line 7). The aim is for UE u' to use a common RIS r during a *scheduling period* $t \in \mathcal{T}_i$ previously assigned to the UE u'' . If no such UEs exist (lines 8-9), UE u' is removed from the set of UEs that can improve their ratio W_u/W_u^{th} . If a suitable UE exists, another iterative loop starts (lines 12-28), selecting common RIS devices, preferentially RIS devices closest to u' (line 13). Once a RIS r is chosen, a *scheduling period* $t \in \mathcal{T}_i$ is sought

Algorithm 2: RIS scheduling

```

1 Inputs:  $W_u^{th}, \varepsilon_u \forall u \in \mathcal{U}_{\mathcal{R}}$ ;
2 Initialization:  $\mathcal{U}_{eva} = \emptyset, \mathcal{U}_{don} = \emptyset, f_{obj}^{(old)} = \infty,$ 
    $f_{obj}^{(new)} = \infty, Cont = False$ ;
3 Set  $x_{r,t,u}$  and  $x'_{r,t,u} = x_{r,t,u}$ ;
4 Compute  $\omega_u, W_u \forall u \in \mathcal{U}_{\mathcal{R}}$ , and  $f_{obj}^{(old)}$ ;
5 while not  $Cont$  do
6   Select UE  $u' \in \mathcal{U}_{\mathcal{R}} \setminus \mathcal{U}_{eva}$  with maximum  $W_{u'}/W_{u'}^{th}$ ;
7   Select UE  $u'' \in \mathcal{U}_{\mathcal{R}} \setminus (\mathcal{U}_{eva} \cup \mathcal{U}_{don})$  with minimum
      $W_{u''}/W_{u''}^{th}$ ;
8   if  $u'' = \emptyset$  then
9      $\mathcal{U}_{eva} = \mathcal{U}_{eva} \cup u'$ 
10  else
11    Set  $\mathcal{R}_{dis} = \emptyset$  and  $stop = False$ ;
12    while  $\mathcal{R}_{dis} \neq \mathcal{R}_{u'} \cup \mathcal{R}_{u''}$  or  $stop == False$  do
13      Set  $r = \arg \min \{d(r', u')\}$ 
         $\forall r' \in \mathcal{R}_{u'} \cup \mathcal{R}_{u''} \setminus \mathcal{R}_{dis}$ ;
14      Randomly select  $t \in \mathcal{T}_{i,u'',r}$ ;
15      if  $t \neq \emptyset$  then
16        Set  $x'_{u',r,t} = 1$  and  $x'_{u'',r,t} = 0$ ;
17        Compute  $W_{u'}, W_{u''}$ , and  $f_{obj}^{(new)}$ ;
18        if  $f_{obj}^{(new)} \leq f_{obj}^{(old)}$  then
19          Update  $x_{u,r,t} = x'_{u,r,t}, W_u = W_{u'}$ , and
             $f_{obj}^{(old)} = f_{obj}^{(new)}$ ;
20        else
21          Set  $\mathcal{U}_{don} = \mathcal{U}_{don} \cup u''$ ;
22        end
23        Set  $stop = True$ ;
24      else
25         $\mathcal{R}_{dis} = \mathcal{R}_{dis} \cup r$ ;
26      end
27    end
28  end
29  if  $\mathcal{U}_{eva} == \mathcal{U}_{\mathcal{R}}$  then
30     $Cont = True$ ;
31  end
32 end
33 return:  $x_{u,r,t}, W_u$  and  $f_{obj}^{(old)}$ ;
```

for u' to use it (line 14). If such a period exists, the exchange is made (line 16), weights and delay bounds are updated, and the new objective function value is computed (line 17). If the objective function improves, the new UE-RIS assignment is updated (line 19). If not, UE u'' is discarded, and the next iteration (outer loop) will attempt another UE. If no suitable *scheduling period* $t \in \mathcal{T}_i$ is found, the nested loop exits (line 21), and another RIS $r \in \mathcal{R}_{u'} \cup \mathcal{R}_{u''} \setminus \mathcal{R}_{dis}$ is considered (line 25). This process repeats until no UEs u' can improve the ratio between the experienced delay bound and the target one (line 30). Finally, the algorithm returns the UE-RIS assignment with the lowest objective function (line 33).

The worst-case complexity of Algorithm 2 is $\mathcal{O}(|\mathcal{U}_{\mathcal{R}}|^2 + |\mathcal{U}_{\mathcal{R}}| \cdot |\mathcal{R}|)$. Each outer iteration selects the UE with the highest normalized delay ratio in $\mathcal{O}(|\mathcal{U}_{\mathcal{R}}|)$, then finds a UE with shared LoS RIS also in $\mathcal{O}(|\mathcal{U}_{\mathcal{R}}|)$. The inner loop explores at most $2|\mathcal{R}|$ RIS devices, performing objective evaluations with cost $\mathcal{O}(|\mathcal{R}|)$ per iteration. With at most $|\mathcal{U}_{\mathcal{R}}|$ outer iterations, total complexity is as stated.

The computational complexities derived for the proposed algorithms refer to a single DARIO instance operating within one base-station domain. In large-scale deployments, multiple DARIO instances can be deployed in parallel, as described in Section II.D. Each instance executes the same two-stage

TABLE II
SIMULATOR PARAMETERS CONFIGURATION

Parameters	Value	Parameters	Value
Cell Layout		UE config	
Size	250x250 m ² [38]	Assignment Period I_{time}	2 s
5G NR BS		Number of UEs ($ \mathcal{U} $)	{15, 20, 25, 30, 35}
Number of antennas (N_{ant})	8	UE height	1.8 m
Frequency band	n79 band (central frequency 4.7 GHz)	UE Mobility Model	Manhattan [40]. Speed $v_u \in [1, 2]$ m/s
Bandwidth	100 MHz (135 RBs)	UE Traffic Pattern	
Subcarriers per RB (N_{sc})	12	Traffic Distribution	Poisson dist. Average rate $\lambda_u \in [450, 550]$ packets/s
Subcarrier spacing (Δ_f)	60 KHz	Packet Size	{64, 128, 256, 512, 1024} bytes. Uniform Dist.
BS height	25 m	uRLLC Requirements	
MAC scheduler	Round robin	Target Delay Bound	{5, 10, 15, 20, 25, 50, 100} ms
RIS		Violation Probability	{ 10^{-3} , 10^{-4} , 10^{-5} }
Number of RIS ($ \mathcal{R} $)	{5, 10, 15, 20, 25, 30}	Channel Parameters	
RIS elements (L)	100	UE transmission power ($P_{u,u}^{tx}$)	24 dBm
RIS height, i.e., placement	3 m	Thermal noise power	-174 dBm/Hz
Number of phase bits (B)	3	Loss Propagation model	3GPP UMa model (TR 38.901 V18.0.0) [38]
DARIO Conf.		Spectral efficiency per MCS ($\eta_{u,k}$)	See 3GPP TS 38.214 Table 5.2.2.1-3 [26]
Scheduling Period T_{time}	100 ms	SNR range per MCS ($[\gamma_k^{min}, \gamma_k^{max}]$)	$[2^{\eta_{u,k}/0.6} - 1, 2^{\eta_{u,k+1}/0.6} - 1]$

optimization procedure independently, preserving the same per-instance complexity. The total computing resources required therefore scale approximately linearly with the number of cells, while the algorithmic complexity of each instance remains unchanged.

VI. PERFORMANCE EVALUATION

We evaluate DARIO using a two-stage approach: first, we perform simulations with synthetic traces to rigorously analyze its performance under controlled conditions, then, we validate these results through real-world experiments described in Section VII.

In the simulation phase, we model a dense urban environment spanning 250×250 m², featuring a single macro cell [38] and multiple RIS devices. The BS operates at 4.7 GHz (Band n79) with 100 MHz bandwidth and 60 KHz subcarrier spacing, providing 135 available RBs to serve delay-critical traffic. Each simulated RIS consists of antenna elements arranged in a 10×10 rectangular grid [39].

Simulated scenarios include a varying number of RIS devices and UEs, where UEs move throughout the urban area following the Manhattan mobility model [40]. Each UE $u \in \mathcal{U}$ generates traffic packets following a Poisson distribution with mean λ_u , and packet sizes d_u are drawn from an arbitrary distribution. These traffic characteristics are defined analogously to the model presented in Section IV-C. Additionally, UEs are assigned specific delay requirements W_u^{th} and violation probabilities ε_u .

Simulations were performed using a Python-based simulator on a computing platform equipped with 16 GB RAM and a quad-core Intel Core i7-7700HQ @ 2.80 GHz processor. Table II summarizes the key parameters of the simulation setup.

A. SNC Model Validation under Poisson Traffic

In the first experiment, we validate the proposed SNC-based model for estimating the delay bound W_u of an UE $u \in \mathcal{U}$, generating packets with an average rate of $\lambda_u = 2000$ packets/s, across four Validation Scenarios (VS). In VS1, we vary the number of RBs allocated to the UE, accounting for different distances from the UE to the BS, and considering both cases with and without exclusive RIS allocation, i.e., $\omega_{u,r}^{S2} = 1$. For all cases, we consider a target violation probability $\varepsilon_u = 10^{-3}$. In VS2, we fixed the number of RBs at $N_u^{RB} = 5$ and vary the violation probability for the previous cases. In VS3, we kept fixed $N_u^{RB} = 5$ and $\varepsilon_u = 10^{-3}$, and analyze how W_u varies with different UE positions, considering both the presence and absence of a RIS device. In VS4, we fixed two UE positions and varied the probability of RIS assignment $\omega_{u,r}^{S2}$ within an *assignment period* $i \in \mathcal{I}$.

Fig. 2 shows the SNC model consistently offers a conservative estimation of W_u . Specifically, the SNC model yields a relative overestimation close to 230%, which aligns with the conservative nature of SNC modeling frameworks [5]. One of the main goals of SNC-based models is to account for complex arrival and service processes. In our study, the arrival process is constructed from empirical traffic traces collected at each TTI, capturing realistic burstiness patterns. Furthermore the service process is defined in Eq. (26) and is based on a probabilistic model. This model captures the UL capacity a UE may obtain depending on the dynamic configuration of RIS devices, which are managed by DARIO to establish temporal LoS links between these RIS devices and the UE. Considering arrival and service processes that are closer to real-world scenarios justifies the use of SNC, which, rather than aiming for an exact match between model and simulator results, provides an upper-bound, conservative estimation of W_u [5]. This conservative estimation is ideal for resource planning, as it ensures the fulfillment of UEs' delay-critical




Fig. 2. Estimated vs. real delay bounds across multiple validation scenarios (VS). VS1: varying number of RBs, UE distances, and RIS assignment; VS2: fixed RBs, varying violation probability; VS3: fixed RBs and violation probability, varying UE positions with/without RIS; VS4: fixed UE positions, varying RIS assignment probability.

requirements in real-world scenarios.

The analysis of the different validation scenarios reveals that assigning a RIS to an UE significantly reduces its experienced delay bound W_u . In VS1 and VS2, a UE using a RIS device at 90m from the BS experiences better delay bound than a UE without RIS at 40m from the BS. In VS3, the use of a RIS device maintains the estimated delay bound nearly invariant across varying UE-BS distances, outperforming scenarios without RIS even when the UE is close to the BS. In VS4, the delay bound increases as the probability of using a RIS within an assignment period decreases, with significant effects at greater distances from the BS.

B. Delay Bound Sensitivity to RIS Phase Quantization

We perform a sensitivity analysis to evaluate how the quantization of RIS phase shifts affects the delay bound W_u . The configuration follows the experimental setup described at the beginning of this section, except for the geometric layout and the distance range. The BS is placed at (0, 0) m, a single RIS is located at (0, 500) m, and one UE moves along the x -axis from 0 to 4 km. The UE is assigned 5 RBs per slot, and the delay bound is computed for a violation probability of 10^{-3} . This setup allows examining how the delay bound varies with the RIS phase quantization level under different propagation conditions.

Figure 3 depicts the evolution of the delay bound W_u as a function of the UE-BS distance for different quantization resolutions $B = \{1, 2, 3, 4, 5\}$ and numbers of RIS elements $L = \{8, 16, 32, 64\}$. The results show that W_u increases exponentially with distance. For small RIS configurations such as $L = 8$, coarse quantization (few discrete phase states) leads to higher values of W_u , since the residual phase error weakens the coherent gain of the reflected components. As L increases, the RIS gain compensates for these errors and the delay bound decreases. Moreover, differences between quantization levels become progressively smaller, and from $B = 3$ bits onward (eight discrete phase states) the quantization error becomes negligible. This confirms that $B = 3$ provides a suitable operating point for RIS devices.

In this scenario, the UE, the RIS, and the BS are in mutual LoS, which makes the deterministic LoS component dominant




Fig. 3. Delay bound W_u versus UE-BS distance for different RIS phase quantization levels B and numbers of reflecting elements L ($\epsilon_u = 10^{-3}$, $N_u^{RB} = 5$ RBs per slot).

TABLE III
OPTIMALITY GAP ANALYSIS OF DARIO

$ \mathcal{T}_i / N_{cell}^{RB}$	Method	Obj. Func.	Relative Error (%)	Iterations	Execution Time
1 / 9	DARIO	3.0623	9.1848	9	178.52 ms
	Brute Force	2.8046		1288	51.21 s
2 / 9	DARIO	1.996	2.0154	10	121.64 ms
	Brute Force	1.9568		29624	16 min. 39 s
3 / 9	DARIO	1.8847	1.0547	11	126.69 ms
	Brute Force	1.8650		681352	6 h. 33 min. 25 s
1 / 10	DARIO	2.0701	8.7060	7	101.73 ms
	Brute Force	1.9043		2898	1 min. 39 s
2 / 10	DARIO	1.7546	5.3681	10	112.69 ms
	Brute Force	1.6652		66654	37 min. 35 s
3 / 10	DARIO	1.5954	4.7488	11	126.69 ms
	Brute Force	1.5231		1533042	14 h. 45 min. 52 s

in the overall channel response. As a result, the sensitivity of the delay bound W_u to the quantization level is substantially reduced. For distances below approximately 1 km when $L = 8$ and up to 2 km when $L = 64$, the effect of the phase resolution on W_u is minimal, since the constructive gain provided by LOS propagation compensates for the discretization error. Therefore, in typical urban deployments with a high density of RIS devices placed close to the UEs and with clear visibility to both the BS and the served UEs, coarse phase quantization has a negligible impact on the achievable delay bound.

C. Evaluating DARIO Against the Optimum

In this experiment, we assess the heuristic algorithms used by DARIO, measuring its deviation from the optimal solution. Due to the problem's complexity, we benchmark our approach against a brute-force algorithm that evaluates all the possible candidate UE-RIS and RBs assignments. For the sake of comprehensibility, we consider a small-scale scenario consisting of $|\mathcal{R}| = 2$ RIS devices and $|\mathcal{U}| = 6$ UEs. In this setup, 2 UEs have LoS with only one RIS, another 2 UEs have LoS with the other RIS, and the remaining 2 UEs have LoS with both RIS devices. This leads to a total of $|\mathcal{C}| = 23$ possible UE-UE assignment combinations in a single *scheduling period* $t \in \mathcal{T}_i$. Furthermore, we consider $N_{cell}^{RB} \in \{9, 10\}$ which leads $\binom{N_{cell}^{RB}-1}{|\mathcal{U}|-1}$ RB allocations, and therefore $23^{|\mathcal{T}_i|} \times \binom{N_{cell}^{RB}-1}{|\mathcal{U}|-1}$ combinations over $|\mathcal{T}_i|$ *scheduling periods*. We evaluate both methods considering different numbers of *scheduling periods* $|\mathcal{T}_i|$ and available RBs. The results, summarized in Table III, indicate that in the worst case the optimal value obtained via the brute force approach and the suboptimal solution from the proposed heuristic exhibit a relative error of 9.19%. Such a discrepancy becomes even more reasonable when considering that the brute force method takes nearly 15 hours to execute for $|\mathcal{T}_i| = 3$ and $N_{cell}^{RB} = 10$, while DARIO only requires < 180 ms. Details about DARIO's computational complexity are discussed in Section VI-D.

We also analyze the convergence of the proposed heuristic over $|\mathcal{I}| = 1800$ assignment periods, considering 40 moving UEs, 20 RIS devices, and *scheduling periods* $|\mathcal{T}_i| \in [2, 10]$ for $\forall i \in \mathcal{I}$. The results, depicted in Fig. 4, reveal that as the number of *scheduling periods* $|\mathcal{T}_i|$ increases, more iterations are needed to extract a solution in a single assignment period while improving the objective function value. The bottom plot displays the CDF of the objective function values collected over the $|\mathcal{I}| = 1800$ assignment periods. While increasing the number of *scheduling periods* generally reduces the objective function value, similar performance is achieved for $|\mathcal{T}_i| \geq 6$.




Fig. 4. Top plot: Convergence of the objective function with the number of scheduling periods $|\mathcal{T}_i|$. Bottom plot: CDF of the objective function values over $|\mathcal{I}| = 1800$ assignment periods.




Fig. 5. Boxplot representation of the distribution of the execution time of DARIO for different network scenarios.

Based on these observations, we empirically set $|\mathcal{T}_i| = 10$ scheduling periods for the remaining experiments. These results indicate that further subdivision of the assignment period into smaller scheduling periods provides negligible delay-bound improvement, supporting the design choice of keeping each RIS configuration fixed within a scheduling period as finer reconfiguration would only yield marginal objective-function improvements.

D. Computational Complexity Analysis of DARIO

We evaluate the computational complexity of DARIO and the number of iterations required by DARIO to determine the UE-RIS assignment for an assignment period $i \in \mathcal{I}$. The boxplot in Fig. 5 shows the distribution of execution time and iteration count across heterogeneous deployment scenarios. The execution time slightly increases with the number of RIS devices and UEs. This is due to a higher probability of multiple UEs having LoS with the same group of RIS devices, which requires DARIO to explore more candidate solutions

to minimize the objective function. The number of iterations typically ranges from 17 to 35, with a maximum of 51. DARIO shows an affordable execution time (< 1.55 s), below the duration of the considered assignment period ($I_{time} = 2$ s). This result confirms that DARIO completes its optimization within the Non-RT control interval defined by O-RAN (i.e., $\leq I_{time} = 2$). The remaining ~ 0.45 s per I_{time} can be used to disseminate the orchestration outputs, namely the radio resource allocation policies and RIS configuration plans, through the proposed RO2 and RO3 interfaces before the next assignment period begins.

VII. REAL-WORLD EVALUATION

We evaluate DARIO's *rApps*, implemented in Python, using empirical traffic traces collected from a real-world deployment, thereby extending the analysis beyond synthetic simulations. We begin by describing the collected traffic traces, detailing their characteristics and how they were integrated into our evaluation framework to assess DARIO's performance. Additionally, we outline the methodology used to adapt these traces to align with the network configuration and delay-sensitive requirements considered in this study. Subsequently, we present the results obtained using these empirical traces, comprehensively comparing DARIO against three reference solutions. This additional testing highlights DARIO's ability to maintain low-latency guarantees and stable performance under realistic traffic fluctuations.

A. RIS Device: Configuration and Deployment

For the empirical evaluation, we use five commercially available off-the-shelf RIS devices [41]. Each RIS operates at 5 GHz and is equipped with a 10×10 patch antenna array supporting electronic phase-shift control for dynamic beam steering. The devices are managed through a custom Python-based control system developed for this study, which applies RIS-wide phase configurations selected from a codebook of 589 predefined phase profiles, each corresponding to a distinct angular response.

A centralized controller orchestrates time-synchronized configuration across all RIS units, ensuring coherent operation during the experiments. Each RIS measures 40×40 cm², which introduces practical deployment constraints related to space availability and LoS conditions. To address this, a site-specific assessment was conducted considering coverage, visibility, and environmental reflections, and the RIS devices were strategically positioned to maximize beamforming effectiveness.

After deployment, extensive measurements were collected under multiple RIS configurations, capturing the resulting signal propagation characteristics. These measurements, used in the subsequent evaluation, provide quantitative evidence of the impact of RIS-assisted operation on latency performance and validate DARIO under realistic conditions.

B. Capturing Urban Real-Traces

We collected a dataset from an operational mobile network deployment in an urban area of Madrid, Spain, using advanced traffic tracking tools [42]. From this original dataset,




Fig. 6. Realization of real traffic traces for 3 UEs. Top: PMF generated traffic. Bottom: Generated traffic per TTI.

we derived the preprocessed traffic traces used in the real-world evaluation presented in this paper. The released dataset³ contains the processed per-user uplink traffic time series at TTI granularity employed in our experiments.

To extend the dataset's scope, we conducted measurements under both RIS-inactive and RIS-active conditions in the network, allowing us to isolate and analyze the impact of RIS configurations on network performance. Specifically, the dataset includes scenarios where RIS devices were inactive, serving as a baseline, and scenarios where RIS devices were dynamically configured to enhance signal propagation. This dual set of measurements enables a detailed comparison, quantifying the performance gains introduced by RIS-assisted propagation. To explore large-scale deployments, scenarios with more than five RIS devices were emulated using measurement-driven models derived from the physical setup. This ensures the extended experiments remained grounded in realistic, empirically validated configurations and performance metrics.

For the experiments, we extracted the UL traffic patterns from all UEs and estimated the empirical PMF of their per-TTI packet arrivals. Using these distributions, we randomly selected a subset of UEs and synthetically generated traffic following the observed patterns. Fig. 6 presents the empirical PMFs of three representative UEs, along with corresponding synthetic realizations.

By incorporating both baseline and RIS-active scenarios, this dataset captures not only empirical mobile network dynamics, but also provides critical insights into the tangible benefits of RIS deployment, further supporting the practical applicability of DARIO in urban-based dynamic network environments.

C. SNC Model Validation Against Real Traces

This experiment validates the proposed SNC delay model using the empirical traffic traces collected from the urban deployment previously described. In this evaluation, the collected traces are used as input to a controlled emulation framework that reproduces the observed traffic dynamics while computing the resulting end-to-end delays under the modeled 5G-RIS configuration. The goal is to verify that the model provides reliable and consistent delay-bound estimations under realistic, trace-driven conditions and to demonstrate that its performance remains stable when compared to a synthetic scenario driven by Poisson arrivals.

³The dataset used in this study is publicly available at IEEE DataPort (<https://iee-dataport.org/documents/preprocessed-urban-mobile-traffic-traces>) to ensure reproducibility and facilitate further research.




Fig. 7. Analytical and emulated delay bounds W_u^{mod} and W_u^{emp} as functions of the violation probability ϵ . The secondary axis depicts the ratio $G(\epsilon) = W_u^{\text{mod}}/W_u^{\text{emp}}$.

For both traffic types, the analytical delay bound $W_u^{\text{mod}}(\epsilon)$ and the emulated delay $W_u^{\text{emu}}(\epsilon)$ were evaluated as functions of the violation probability $\epsilon \in [10^{-6}, 10^{-3}]$, considering a single 5G BS with and without RIS assistance. The comparison also includes the ratio $G(\epsilon) = W_u^{\text{mod}}(\epsilon)/W_u^{\text{emu}}(\epsilon)$, which quantifies the conservatism of the analytical bound with respect to the delay obtained through trace-driven emulation.

Figure 7 shows the obtained results. For both real-trace and Poisson traffic, the analytical and emulated delay curves exhibit a similar evolution across all values of ϵ , with $G(\epsilon)$ consistently taking values between 3 and 5. This behaviour indicates that the delay-bound estimation produced by the SNC model remains consistent regardless of whether the input follows a memoryless Poisson process or a more complex real-trace pattern. In particular, the model, whose arrival envelope is derived from the empirical MGF, responds similarly under both traffic types, suggesting that its operation is largely insensitive to the specific statistical structure of the incoming traffic, including burstiness and temporal correlations.

A slight increase of $G(\epsilon)$ at low violation probabilities is observed in the RIS-assisted case, which is explained by the higher variability of the service process resulting from the cascaded UE-RIS-BS channel and the quantization-induced phase errors. This variability affects the tail behavior of the delay distribution, slightly widening the analytical bound.

D. Performance Evaluation of DARIO against Reference Solutions

We evaluate the performance of DARIO using the real traffic traces described in Section VII-B, thereby assessing its practical effectiveness under realistic and time-varying network conditions. Since there are no existing RIS control frameworks that jointly address dynamic UE-RIS assignment, user mobility, stochastic traffic arrivals, and probabilistic delay-bound guarantees, we compare DARIO against three benchmark solutions: *i)* No RIS Assignment: The baseline scenario where the BS operates without any UE-RIS assignment. This configuration isolates the performance gains attributable solely to the integration of RIS devices; *ii)* SNR-Based RIS Assignment: A static, one-to-one UE-RIS assignment strategy, where each UE is paired with the RIS providing the highest average SNR over the UE-RIS-BS link. The static nature of




Fig. 8. Top plot: Evolution of the average traffic load potentially served by three arbitrary RIS devices. Bottom plot: Evolution of the objective function over the first 160 assignment periods. For these measurements, we consider a scenario with $|\mathcal{U}| = 40$ UEs and $|\mathcal{R}| = 15$ RIS devices.

this approach means the assignment remains fixed throughout all *scheduling periods* $t \in \mathcal{T}_i$ within a single *assignment period* $i \in \mathcal{I}$; *iii*) Delay-Aware Static RIS Assignment: A partially delay-aware solution that considers UE delay requirements for one-to-one UE-RIS assignment during the entire *assignment period* i . However, it lacks the flexibility of DARIO in supporting dynamic, many-to-many UE-RIS assignments across *scheduling periods*. All these schemes are evaluated under identical channel and CSI assumptions, so the observed performance differences are attributable to the orchestration and scheduling strategies.

To evaluate DARIO, we measured the objective function value over $|\mathcal{I}| = 1800$ *assignment periods* under varying configurations of UEs and RISs. Fig. 8 illustrates network dynamics during the first 160 assignment periods for a scenario with $|\mathcal{U}| = 40$ UEs and $|\mathcal{R}| = 15$ RISs. The top plot shows average traffic (blue curve) served by three RIS devices, alongside the number of UEs (red bars) with LoS to those RIS devices over time, highlighting the dynamic traffic conditions DARIO is designed to handle. The bottom plot shows the temporal evolution of the objective function, where RIS-enabled solutions (green, blue, and red curves) consistently outperform the baseline (orange curve) across all assignment periods. Notably, DARIO (red curve) yields the lowest objective function values throughout, highlighting its superior adaptability to traffic dynamics.

To provide a holistic comparison, we also evaluated DARIO across different scenarios where UEs and RISs vary in number. Fig. 9 presents the CDFs of the objective function under four extreme scenarios: Scenario (i): $|\mathcal{U}| = 30$, $|\mathcal{R}| = 5$ (few UEs and RIS); Scenario (ii): $|\mathcal{U}| = 50$, $|\mathcal{R}| = 5$ (many UEs and few RIS); Scenario (iii): $|\mathcal{U}| = 30$, $|\mathcal{R}| = 30$ (few UEs and many RIS); and Scenario (iv): $|\mathcal{U}| = 50$, $|\mathcal{R}| = 30$ (many UEs and many RIS).

In all scenarios, DARIO consistently outperforms the other strategies, followed by Delay-Aware Static, SNR-Based, and finally No RIS. This ordering is expected: No RIS does not leverage RIS technology; SNR-Based, although using RIS, ignores delay requirements; and among the delay-aware approaches, DARIO benefits from dynamic many-to-many UE-RIS assignment within each *assignment period*, unlike Delay-




Fig. 9. CDFs of the values of the objective function when it is measured over $|\mathcal{I}| = 1800$ assignment periods for different network scenarios.

Aware Static, which uses fixed mappings.

In Scenario (i), all CDFs are closely aligned due to the low number of UEs and RIS devices, which limits the likelihood of a UE experiencing LoS with a RIS during an *assignment period*. Consequently, RIS-enabled strategies provide only modest gains over No RIS. Even so, DARIO achieves a 90th percentile (P90) gain of 79.08% over No RIS, 68.99% over SNR-Based, and 52.20% over Delay-Aware Static; 50th percentile (P50) gains are 64.90%, 46.30%, and 36.62%, respectively.

In Scenario (ii), the number of UEs increases while RIS count remains unchanged, raising network load and shifting all CDFs to the right. Here, DARIO's P90 gains are smaller than in Scenario (i), 5.61% over No RIS, 47.66% over SNR-Based, and 16.53% over Delay-Aware Static, because more UEs compete for the same RIS devices. Nonetheless, P50 gains (81.73%, 68.24%, and 47.37%, respectively) still show that DARIO provides a clear advantage.

In Scenario (iii), RIS availability increases substantially, shifting the CDFs of RIS-enabled solutions to the left, while No RIS moves further right. With a nearly 1:1 UE-RIS ratio, all RIS-based approaches benefit, but the dynamic mapping of DARIO extracts more value from a better use of RIS devices

than static methods. As a result, DARIO's P90 gains (95.71%, 73.32%, and 55.10% over No RIS, SNR-Based, and Delay-Aware Static, respectively) are larger than in Scenarios (i) and (ii). P50 gains are 87.76%, 55.52%, and 32.55%, respectively.

Finally, Scenario (iv) consider the higher number of UEs and RIS devices, increasing load relative to Scenario (iii) but also creating more opportunities for dynamic UE-RIS mapping. Here, the performance gaps between RIS-based strategies widen, with DARIO achieving substantial P90 gains over No RIS (78.38%), SNR-Based (75.82%), and Delay-Aware Static (46.61%). P50 gains reach 91.34%, 72.40%, and 52.35%, respectively, confirming that per-scheduling-period RIS reassignment is particularly beneficial under heavy load.

In addition, Fig. 9 also reports the CDF of the average per-UE throughput (dashed curves, top x-axis), where curves shifted to the right indicate higher throughput. Across all scenarios, DARIO achieves per-UE throughput comparable to the SNR-based and no-RIS baselines, while consistently outperforming Delay-Aware Static RIS assignment at medium and high percentiles. This confirms that the delay-bound improvements provided by DARIO are not obtained at the expense of a significant throughput loss: dynamic RIS reallocation primarily reshapes the service process to reduce delay violations, while preserving the long-term average throughput experienced by the UEs.

VIII. CONCLUSIONS

In this paper, we presented DARIO, an O-RAN-compliant RIS scheduler that dynamically configures multiple RIS devices across time slots to establish controlled LoS links between UEs and the BS. By adaptively assigning UEs to RIS devices, DARIO reduces packet delays and supports delay-sensitive requirements under dynamic channel and traffic conditions. This is enabled by a novel SNC-based model for estimating delay violation probabilities, which guides near-optimal UE-RIS assignments formulated as a NIP and solved via a computationally efficient heuristic.

We validated DARIO using real traffic traces and measurements from an operational deployment with five commercial RIS devices, and leveraged this data to evaluate larger-scale scenarios. The results demonstrate the scalability and practical benefits of DARIO in realistic environments, consistently outperforming all benchmark strategies across all scenarios. Performance gains depend on traffic load and RIS availability, reaching up to 95.71% over No RIS, 75.82% over SNR-Based, and 55.10% over Delay-Aware Static at the 90th percentile, thereby confirming the advantage of dynamic many-to-many UE-RIS assignment over static approaches.

REFERENCES

- [1] W. Saad, M. Bennis, and M. Chen, "A Vision of 6G Wireless Systems: Applications, Trends, Technologies, and Open Research Problems," *IEEE Netw.*, vol. 34, no. 3, pp. 134–142, 2020.
- [2] M. A. Kishk and M.-S. Alouini, "Exploiting Randomly Located Blockages for Large-Scale Deployment of Intelligent Surfaces," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 4, pp. 1043–1056, 2021.
- [3] A. Anand, G. de Veciana, and S. Shakkottai, "Joint Scheduling of URLLC and eMBB Traffic in 5G Wireless Networks," *IEEE/ACM Trans. Netw.*, vol. 28, no. 2, pp. 477–490, 2020.
- [4] M. Fidler, "Survey of deterministic and stochastic service curve models in the network calculus," *IEEE Commun. Surv. Tutor.*, vol. 12, no. 1, pp. 59–86, 2010.
- [5] M. Fidler and A. Rizk, "A Guide to the Stochastic Network Calculus," *IEEE Commun. Surveys Tuts.*, vol. 17, no. 1, pp. 92–105, 2015.
- [6] O. Adamuz-Hinojosa, V. Sciancalepore, P. Ameigeiras, J. M. Lopez-Soler, and X. Costa-Pérez, "A Stochastic Network Calculus (SNC)-Based Model for Planning B5G uRLLC RAN Slices," *IEEE Trans. Wireless Commun.*, vol. 22, no. 2, pp. 1250–1265, 2023.
- [7] M. Polese, L. Bonati, S. D'Oro, S. Basagni, and T. Melodia, "Understanding o-ran: Architecture, interfaces, algorithms, security, and research challenges," *IEEE Commun. Surv. Tutor.*, vol. 25, no. 2, pp. 1376–1411, 2023.
- [8] O. Adamuz-Hinojosa, L. Zanzi, V. Sciancalepore, and X. Costa-Pérez, "MAREA: A Delay-Aware Multi-time-Scale Radio Resource Orchestrator for 6G O-RAN," *IEEE Trans. Commun.*, pp. 1–1, 2025.
- [9] L. Du, S. Shao, G. Yang, J. Ma, Q. Liang, and Y. Tang, "Capacity Characterization for Reconfigurable Intelligent Surfaces Assisted Multiple-Antenna Multicast," *IEEE Trans. Wireless Commun.*, vol. 20, no. 10, pp. 6940–6953, 2021.
- [10] H. Guo, Y.-C. Liang, J. Chen, and E. G. Larsson, "Weighted Sum-Rate Maximization for Reconfigurable Intelligent Surface Aided Wireless Networks," *IEEE Trans. Wireless Commun.*, vol. 19, no. 5, pp. 3064–3076, 2020.
- [11] L. Du, W. Zhang, J. Ma, and Y. Tang, "Reconfigurable Intelligent Surfaces for Energy Efficiency in Multicast Transmissions," *IEEE Trans. Veh. Technol.*, vol. 70, no. 6, pp. 6266–6271, 2021.
- [12] S. Liu, R. Liu, M. Li, Y. Liu, and Q. Liu, "Joint BS-RIS-User Association and Beamforming Design for RIS-Assisted Cellular Networks," *IEEE Trans. Veh. Technol.*, vol. 72, no. 5, pp. 6113–6128, 2023.
- [13] Q. Xue, H. Xia, J. Mu, Y. Xu, L. Yan, and S. Ma, "User-Centric Association for Dense mmWave Communication Systems With Multi-Connectivity," *IEEE Trans. Green Commun. Netw.*, vol. 8, no. 1, pp. 177–189, 2024.
- [14] Q. Peng, H. Ren, C. Pan, M. Elkhassan, A. García Armada, and P. Popovski, "Two-Timescale Design for Reconfigurable Intelligent Surface-Aided URLLC," *IEEE Trans. Wireless Commun.*, vol. 23, no. 10, pp. 13 664–13 677, 2024.
- [15] Y. Yao *et al.*, "Hybrid RIS-Enhanced ISAC Secure Systems: Joint Optimization in the Presence of an Extended Target," *IEEE Trans. Commun.*, vol. 73, no. 12, pp. 15 688–15 704, 2025.
- [16] Y. Yao, Z. Zhu, P. Miao, X. Cheng, F. Shu, and J. Wang, "Optimizing Hybrid RIS-Aided ISAC Systems in V2X Networks: A Deep Reinforcement Learning Method for Anti-Eavesdropping Techniques," *IEEE Trans. Veh. Technol.*, vol. 74, no. 6, pp. 9224–9239, 2025.
- [17] M. Mukherjee, V. Kumar, S. Kumar, C. X. Mavromoustakis, Q. Zhang, and M. Guo, "RIS-assisted Task Offloading for Wireless Dead Zone to Minimize Delay in Edge Computing," in *IEEE Globecom*, 2022, pp. 2554–2559.
- [18] L. Xia *et al.*, "Delay Minimization for RIS-NOMA Assisted MEC Networks With SWIPT," in *IEEE Globecom*, 2022, pp. 4631–4636.
- [19] M. Almekhlafi, M. A. Arfaoui, M. Elhattab, C. Assi, and A. Ghayeb, "Joint Resource Allocation and Phase Shift Optimization for RIS-Aided eMBB/URLLC Traffic Multiplexing," *IEEE Trans. Commun.*, vol. 70, no. 2, pp. 1304–1319, 2022.
- [20] J. H. I. De Souza, V. Croisfelt, R. Kotaba, T. Abrão, and P. Popovski, "Uplink Multiplexing of eMBB/URLLC Services Assisted by Reconfigurable Intelligent Surfaces," *IEEE Commun. Lett.*, pp. 1–1, 2024.
- [21] L. Liu and W. Yu, "Uplink multiplexing of embb/urllc services assisted by reconfigurable intelligent surfaces," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 7, pp. 1931–1945, 2021.
- [22] M. Soleymani, A. Zappone, E. Jorswieck, M. Di Renzo, and I. Santamaria, "Rate region of RIS-aided URLLC broadcast channels: Diagonal versus beyond diagonal globally passive RIS," *IEEE Wirel. Commun. Lett.*, 2024.
- [23] M. Soleymani, I. Santamaria, E. Jorswieck, and B. Clerckx, "Rate splitting multiple access for ris-aided urllc mimo broadcast channels," *arXiv preprint arXiv:2411.11028*, 2024.
- [24] W. Jiang and H. D. Schotten, "A Simple Multiple-Access Design for Reconfigurable Intelligent Surface-Aided Systems," in *IEEE Globecom*, 2023, pp. 2414–2419.
- [25] N. Abu-Alli, A.-E. M. Taha, M. Salah, and H. Hassanein, "Uplink Scheduling in LTE and LTE-Advanced: Tutorial, Survey and Evaluation Framework," *IEEE Commun. Surv. Tutor.*, vol. 16, no. 3, pp. 1239–1265, 2014.
- [26] 3GPP, "NR; Physical layer procedures for data (Release 18)," Technical Specification (TS) V18.3.0, June 2024.
- [27] J. He, Z. Tang, Z. Ding, and D. Wu, "Successive Interference Cancellation and Fractional Frequency Reuse for LTE Uplink Communications," *IEEE Trans. Veh. Technol.*, vol. 67, no. 11, pp. 10 528–10 542, 2018.
- [28] H. Li, P. Zhiwen, W. Bin, L. Nan, and Y. Xiaohu, "Channel Estimation for Reconfigurable-Intelligent-Surface-Aided Multiuser Communication

Systems Exploiting Statistical CSI of Correlated RIS–User Channels,” *IEEE Internet Things J.*, vol. 11, no. 5, pp. 8871–8881, 2024.

- [29] H. Cheng, P. Johari, M. A. Arfaoui, F. Periard, P. Pietraski, G. Zhang, and T. Melodia, “Real-Time AI-Enabled CSI Feedback Experimentation with Open RAN,” in *WONS*, 2024, pp. 121–124.
- [30] J. Sang, J. Lan, M. Zhou, B. Gao, W. Tang, X. Li, M. Matthaiou, S. Jin, and M. D. Renzo, “Measurement-based small-scale channel model for sub-6 ghz ris-assisted communications,” *IEEE Transactions on Vehicular Technology*, vol. 73, no. 8, pp. 12 178–12 183, 2024.
- [31] P. Xu, G. Chen, Z. Yang, and M. D. Renzo, “Reconfigurable Intelligent Surfaces-Assisted Communications With Discrete Phase Shifts: How Many Quantization Levels Are Required to Achieve Full Diversity?” *IEEE Wireless Communications Letters*, vol. 10, no. 2, pp. 358–362, 2021.
- [32] M. A. Badiu and J. P. Coon, “Communication through a large reflecting surface with phase errors,” *IEEE Wireless Communications Letters*, vol. 9, no. 2, pp. 184–188, 2019.
- [33] J.-Y. Le Boudec and P. Thiran, *Network Calculus: A Theory of Deterministic Queuing Systems for the Internet*. Berlin, Germany: Springer, 2001, vol. 2050.
- [34] Y. Jiang *et al.*, *Stochastic Network Calculus*. London, U.K.: Springer, 2008, vol. 1.
- [35] S. Ross, *A First Course in Probability*. Pearson, 2014.
- [36] O. Adamuz-Hinojosa, L. Zanzi, V. Sciancalepore, A. Garcia-Saavedra, and X. Costa-Pérez, “ORANUS: Latency-tailored Orchestration via Stochastic Network Calculus in 6G O-RAN,” in *IEEE INFOCOM*, 2024.
- [37] R. Hemmecke, M. Köppe, J. Lee, and R. Weismantel, *Nonlinear Integer Programming*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 561–618. [Online]. Available: https://doi.org/10.1007/978-3-540-68279-0_15
- [38] 3GPP, “Study on channel model for frequencies from 0.5 to 100 GHz (Release 18),” Technical Report (TR) V18.0.0, March 2024.
- [39] M. Rossanese, P. Mursia, A. Garcia-Saavedra, V. Sciancalepore, A. Asadi, and X. Costa-Perez, “Open Experimental Measurements of Sub-6GHz Reconfigurable Intelligent Surfaces,” *IEEE Internet Comput.*, vol. 28, no. 2, pp. 19–28, 2024.
- [40] F. Bai, N. Sadagopan, and A. Helmy, “IMPORTANT: a framework to systematically analyze the Impact of Mobility on Performance of Routing Protocols for Adhoc Networks,” in *IEEE INFOCOM*, vol. 2, 2003, pp. 825–835 vol.2.
- [41] NEC Corporation, “NEC Smart Surface v2022,” <https://www.nec-enterprise.com/it/projects>, 2023.
- [42] R. Falkenberg and C. Wietfeld, “FALCON: An Accurate Real-Time Monitor for Client-Based Mobile Network Data Analytics,” in *2019 IEEE GLOBECOM*, 2019, pp. 1–7.



Oscar Adamuz-Hinojosa received the B.Sc., M.Sc., and Ph.D. degrees in telecommunications engineering from the University of Granada, Granada, Spain, in 2015, 2017, and 2022, respectively. He was awarded a Ph.D. fellowship by the Spanish Ministry of Education in September 2018. He is currently an Interim Assistant Professor with the Department of Signal Theory, Telematics, and Communications, University of Granada (Spain), and has been a Visiting Researcher at NEC Laboratories Europe (Germany) on several occasions. His research interests

include network dimensioning in 5G/6G RAN, 5G/6G integration with TSN, and quantum communications and their integration with classical networks, with a focus on mathematical modeling and architecture design across these domains.



Lanfranco Zanzi (Member, IEEE) received the B.Sc. and M.Sc. degrees in telecommunication engineering from the Polytechnic of Milan, Italy, in 2014 and 2017, respectively, and the Ph.D. degree from the Technical University of Kaiserslautern, Germany, in 2022. He works as a Senior Research Scientist with NEC Laboratories Europe. His research interests include network virtualization, machine learning, blockchain, and their applicability to 5G and 6G mobile networks.



Vincenzo Sciancalepore (Senior Member, IEEE) is a Principal Researcher with NEC Laboratories Europe GmbH, Germany. He is currently focusing his activity in the area of reconfigurable intelligent surfaces (RIS) and integrated sensing and communications. He is the Standard Delegate of NEC actively contributing to the standard ETSI RIS ISG. He currently holds the Italian habilitation as a Full Professor of Telecommunications issued by MIUR. He is currently the Innovation Project Manager of an H2020-funded project called 6G-GOALS that will analyze and design new innovation for semantic communications on 6G systems. He is an Editor of *IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS* and *IEEE TRANSACTIONS ON COMMUNICATIONS*. He is currently a member of the IEEE Emerging Technologies Standing Committee leading the initiatives on RISs.



Marco Di Renzo (Fellow, IEEE) received the Laurea (cum laude) and Ph.D. degrees in electrical engineering from the University of L'Aquila, Italy, in 2003 and 2007, respectively, and the Habilitation à Diriger des Recherches (Doctor of Science) degree from University Paris-Sud (currently Paris-Saclay University), France, in 2013. Currently, he is Chair Professor of Telecommunications Engineering, the Director of the Centre for Telecommunications Research, and the Head of the Telecommunications Group, Department of Engineering, King's College London, London, United Kingdom. He is also a CNRS Research Director (Professor) with the Laboratory of Signals and Systems at CNRS-CentraleSupélec, Paris-Saclay University, Paris, France. He was a France-Nokia Chair of Excellence in ICT at the University of Oulu (Finland), a Tan Chin Tuan Exchange Fellow in Engineering at Nanyang Technological University (Singapore), a Fulbright Fellow at The City University of New York (USA), a Nokia Foundation Visiting Professor at Aalto University (Finland), and a Royal Academy of Engineering Distinguished Visiting Fellow at Queen's University Belfast (U.K.). He is a Fellow of the IEEE, IET, EURASIP, and AAIA; an Academician of AIIA; an Ordinary Member of the European Academy of Sciences and Arts, an Ordinary Member of the Academia Europaea, and an Ordinary Member of the Italian Academy of Technology and Engineering; an Ambassador of the European Association on Antennas and Propagation; and a Highly Cited Researcher. He has received several distinctions, including the Michel Monpetit Prize conferred by the French Academy of Sciences, the IEEE Communications Society Heinrich Hertz Award, and the IEEE Communications Society Marconi Prize Paper Award in Wireless Communications. Also, he is a principal investigator of an ERC Synergy grant on metasurface-based information processing. He served as the Editor-in-Chief of *IEEE Communications Letters* from 2019 to 2023, and as the Director of Journals and Chair of the Publications Misconduct Ad Hoc Committee of the IEEE Communications Society from 2024 to 2025. Currently, he sits on the IEEE-COMSOC Fellow Evaluation Standing Committee and on the Editorial Board of the *Proceedings of the IEEE*.



Xavier Costa-Pérez (Senior Member, IEEE) received the Ph.D. degree in telecommunications from the Polytechnic University of Catalonia, Barcelona. He is a Research Professor with ICREA, a Scientific Director with i2Cat Research Center, and the Head of 5G/6G Networks R&D with NEC Laboratories Europe. He has served on the Organizing Committees of several conferences, published papers of high impact, and holds tens of granted patents. He was a recipient of the National Award for his Ph.D. thesis.