

RESEARCH ARTICLE

COMIX: Generalized Conflict Management in O-RAN xApps—Architecture, Workflow, and a Power Control Case

ANASTASIOS E. GIANNOPOULOS¹, (Member, IEEE), SOTIRIOS T. SPANTIDEAS¹,
GEORGE LEVIS¹, ALEXANDROS S. KALAFATELIS, (Graduate Student Member, IEEE),
AND PANAGIOTIS TRAKADAS

Research and Development Department, Four Dot Infinity (FDI), 16777 Athens, Greece

Corresponding author: Anastasios E. Giannopoulos (angianno@fourdotinfinity.com)

This work was supported in part by UNITY-6G Project, funded by EU HORIZON-JU-SNS-2024 Program, under Grant 101192650; and in part by “REACT-6G” Project, funded by HORIZON-JU-SNS-2022, 2nd 6G-SANDBOX Open Call, under Grant GA101096328.

ABSTRACT Open Radio Access Network (O-RAN) is transforming the telecommunications landscape by enabling flexible, intelligent, and multi-vendor networks. Central to its architecture are xApps hosted on the Near-Real-Time RAN Intelligent Controller (Near-RT RIC), which optimize network functions in real time. However, the concurrent operation of multiple xApps with conflicting objectives can lead to suboptimal performance. This paper introduces a generalized Conflict Management scheme for Multi-Channel Power Control in O-RAN xApps (COMIX), designed to detect and resolve conflicts between xApps. To demonstrate COMIX, we focus on two Deep Reinforcement Learning (DRL)-based xApps for power control: one maximizes the data rate across UEs, and the other optimizes system-level energy efficiency. COMIX employs a standardized Conflict Mitigation Framework (CMF) for conflict detection and resolution and leverages the Network Digital Twin (NDT) to evaluate the impact of conflicting actions before applying them to the live network. We validate the framework using a realistic multi-channel power control scenario under various conflict resolution policies, demonstrating its effectiveness in balancing antagonistic objectives. Evaluation results show that COMIX achieves up to 60% energy savings across different Service-Level Agreement (SLA) policies compared to a baseline conflict-unaware system, with negligible impact (around 3%) on system throughput. While this study considers power control xApps, the COMIX framework is generalizable and can be applied to any xApp conflict scenario involving resource contention or KPI interdependence.

INDEX TERMS 6G, conflict detection, conflict management, deep reinforcement learning, energy efficiency, O-RAN, power control, resource management, xApp.

I. INTRODUCTION

A. OPEN RAN OVERVIEW

The Open Radio Access Network (O-RAN) architecture represents a transformative approach to the design and operation of radio access networks, emphasizing openness, intelligence, and interoperability [1]. At its core, O-RAN fosters a disaggregated architecture where components such as the Radio Unit (RU), Distributed Unit (DU), and Central Unit (CU) can be provided by different vendors, promoting

The associate editor coordinating the review of this manuscript and approving it for publication was Ding Xu¹.

a multi-vendor ecosystem. This vendor-agnostic framework enables operators to deploy a flexible and cost-efficient network by selecting the preferred components [2]. Central to this architecture is the RAN Intelligent Controller (RIC), which is divided into the Non-Real-Time RIC (Non-RT RIC) and the Near-Real-Time RIC (Near-RT RIC). These components host specialized applications—rApps on the Non-RT RIC and xApps on the Near-RT RIC—which drive network intelligence and optimization. rApps, with their long timescales, provide strategic guidance, while xApps focus on fine-grained adaptations to network dynamics in real time. By enabling seamless integration of vendor-specific

or third-party applications, O-RAN not only accelerates innovation but also enhances adaptability in addressing diverse network scenarios [3].

The Near-RT RIC plays a pivotal role in achieving intelligent and dynamic control of the RAN through its support for closed-loop optimization [4]. By operating on timescales ranging from 10 milliseconds to 1 second, the Near-RT RIC manages critical tasks such as resource allocation, interference management, and power control in near real-time. This level of responsiveness enables xApps to directly influence the behavior of RAN components like RUs, DUs, and CUs, ensuring adaptability to rapidly changing network conditions [5]. The Near-RT RIC’s tight control loop complements the slower, policy-driven adjustments of the Non-RT RIC, enabling a hierarchical and time-sensitive optimization framework. This separation of control loops ensures scalability and efficiency while allowing the integration of diverse xApps targeting distinct objectives, such as enhancing data rates or improving energy efficiency.

B. CONFLICT MANAGEMENT IN O-RAN xApps

The dynamic and decentralized nature of O-RAN introduces significant challenges in ensuring harmonious operation among multiple xApps and/or rApps running concurrently on the RICs [6]. A conflict arises when the actions or objectives of one application negatively impact the performance or objectives of another [7]. These conflicts can degrade the overall efficiency, stability, and fairness of the network, leading to contradicting decisions. For example, unresolved conflicts between xApps managing power control or resource allocation can lead to suboptimal utilization of network resources, increased interference, and inconsistent user experiences. Consequently, robust conflict detection mechanisms are required to identify potential issues proactively, and effective conflict resolution strategies are essential to restore balance and ensure that the system functions optimally without favoring specific applications unjustly. To address this challenge, O-RAN Alliance specifies a Conflict Management Framework (CMF) within the RIC architecture [8]. The CMF is tasked with detecting, classifying, and resolving conflicts among xApps and rApps to ensure optimal coordination and resource utilization.

Conflicts in O-RAN can be categorized as shown in Fig. 1. Vertical conflicts occur between applications at different layers, such as an xApp contradicting the policies enforced by an rApp. In contrast, horizontal conflicts emerge between xApps operating at the same layer, often competing for shared resources. Intra-component conflicts occur between xApps at the same Near-RT RIC, while inter-component conflicts concern xApps at different Near-RT RICs that control neighboring RAN nodes. Intra-component conflicts are further classified as direct, indirect, or implicit. A direct conflict arises when two xApps take different actions on the same parameter; for instance, one xApp increases transmission power to improve throughput, while another reduces it

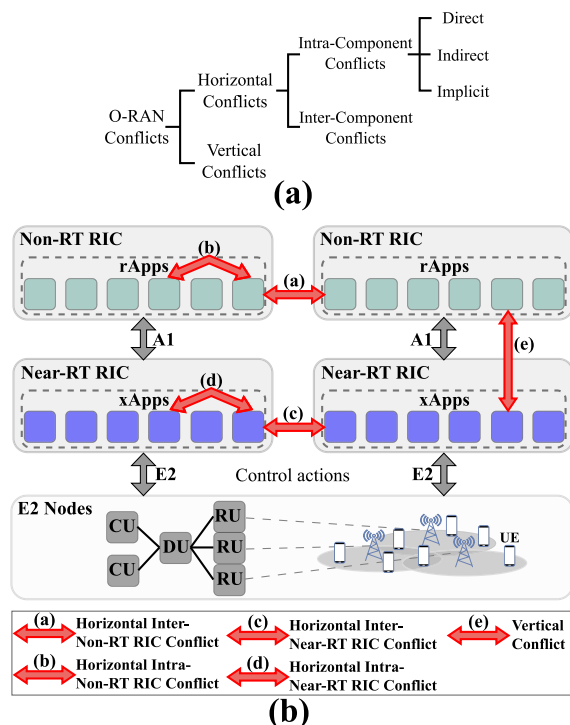


FIGURE 1. Conflicts in O-RAN. (a) Conflict categories in O-RAN. (b) Architectural view of conflict categories.

to save energy. Indirect conflicts occur when one xApp’s actions inadvertently affect another’s objectives; for example, an xApp optimizing handovers might unintentionally reduce available resources for a neighboring cell managed by another xApp. Implicit conflicts between xApps’ decisions are harder to detect because, while individually aligned with their specific objectives, they result in an undesirable overall network state. For instance, an xApp prioritizing latency implicitly violates another xApp aimed at load balancing. Implicit conflicts usually cannot be detected in advance; rather, they become apparent during long-term network operation.

C. RELATED WORK

Conflict detection and resolution have gathered significant attention in the context of O-RAN, with many studies addressing the issue from architectural and algorithmic perspectives [9], [10]. A recent study proposed an architectural conflict mitigation framework embedded within the Near-RT RIC, which identifies and resolves conflicts using predefined rules and real-time monitoring, aligning with O-RAN’s modular architecture principles [7]. Another study introduced a pre-emptive conflict detection architecture within the Service Management and Orchestration (SMO) framework, focusing on early identification of potential conflicts before they escalate to performance degradation. Their solution emphasizes proactive monitoring of multi-xApp interactions [11].

Algorithmic approaches have also been extensively studied for conflict mitigation in O-RAN. The authors in [12] proposed a mobility-driven energy-saving approach to manage xApp-level conflicts, particularly those arising from overlapping objectives, such as energy efficiency and mobility management. However, this energy-saving framework focuses on a single optimization objective without addressing real-time conflict resolution across heterogeneous xApps. Another study in [13] employed a game-theoretic framework within the Near-RT RIC to model and resolve conflicts among xApps competing for shared network resources. A limitation of this study is the assumption for fully cooperative agents and static system parameters, which may not reflect dynamic O-RAN conditions. Another study proposed a learning-based and data-driven method for reconstructing and labeling conflict graphs in O-RAN, aiming to dynamically learn the hidden relationships between xApps, parameters, and Key Performance Indicators (KPIs) [14]. Additionally, another study presented a team-based learning framework to promote collaborative decision-making among xApps, reducing implicit and indirect conflicts during resource allocation [15]. This work indirectly addresses conflicts by promoting cooperative learning among xApps for efficient resource distribution. However, collaboration among xApps presents high complexity in real scenarios where multiple vendor-specific xApps are considered. While these studies provide valuable insights into conflict detection and resolution mechanisms to achieve conflict-free O-RAN operation, there is a notable lack of specific use case-based validations or proof-of-concept scenarios that quantitatively compare different conflict resolution policies under realistic network conditions.

D. PAPER SUMMARY AND CONTRIBUTIONS

This work presents a conflict resolution management scheme for multi-channel power control in O-RAN xApps (COMIX), designed to address conflicting objectives among RU-targeted decisions. Unlike prior CMF implementations that rely solely on reactive or rule-based conflict resolution, COMIX uniquely integrates a Network Digital Twin (NDT) to proactively simulate the outcomes of conflicting xApp actions. This predictive capability enables more informed and reliable conflict resolution decisions, minimizing risks to live network operations. Aligned with the standardized CMF defined by the O-RAN Alliance, we adopt a generalizable approach to classify and manage conflicts. Our focus is on resolving direct conflicts between two Deep Reinforcement Learning (DRL)-based xApps for multi-channel power control at the RU level in practical O-RAN scenarios. The main contributions of this work may be summarized as follows:

- We formulate the power control problem for RUs in multi-channel, multi-cell, and mobility-aware multi-user O-RAN networks as a non-convex optimization problem. This formulation considers the impact of inter-channel and inter-cell interference on both experienced

TABLE 1. Acronyms.

Acronym	Meaning
API	Application Programming Interface
CD	Conflict Detector
CMF	Conflict Mitigation Framework
COMIX	Conflict Management for Multi-Channel Power Control in O-RAN xApps
CP	Control Parameter
CQI	Channel Quality Indicator
CU	Central Unit
DC	Direct Conflict
DMPC	Downlink Multi-channel Power Control
DRM	Data Rate Maximization
DRL	Deep Reinforcement Learning
DU	Distributed Unit
E2SM	E2 Service Model
EE	Energy Efficiency
EES	Energy Efficiency-based Selection
EEVS	EE Violation-based Selection
FH	Front-haul
IC	Indirect Conflict
KPI	Key Performance Indicator
MaxTS	Maximum Throughput-based Selection
MinPS	Minimum Power-based Selection
MNO	Mobile Network Operator
NDT	Network Digital Twin
Near-RT RIC	Near-Real-Time RAN Intelligent Controller
Non-RT RIC	Non-Real-Time RAN Intelligent Controller
O-RAN	Open Radio Access Network
QoS	Quality of Service
RB	Resource Block
RIC	RAN Intelligent Controller
RU	Radio Unit
SLA	Service-level Agreement
SINR	Signal-to-Interference-plus-Noise Ratio
SMO	Service, Management and Orchestration
TVS	Throughput Violation-based Selection
UE	User Equipment
xApp	eXtended Application

users' throughput and system energy efficiency (EE). To address this, we propose DRL algorithms to solve two variants of the problem, adhering to RU power budget and resource block (RB) constraints while ensuring Quality of Service (QoS) requirements.

- Two DRL-based xApps are developed for intelligent power regulation at the RU level. Leveraging E2 O-RAN metrics, the first xApp aims to maximize total system throughput, while the second focuses on optimizing overall EE. Both xApps dynamically adjust RU power levels under multi-channel transmissions, highlighting the inherent conflict in their objectives. We adopt and compare multiple decision-making policies applied by the CMF to resolve antagonistic conflicts by jointly considering different policies such as Service-Level Agreement (SLA) violations, power budget constraints, and QoS satisfaction.
- In the COMIX framework, we provide a structured methodology for handling competitive suggestions from xApps with conflicting objectives. To improve the precision of the decision, the proposed framework leverages the NDT [16], [17] to evaluate the impact of candidate actions before final decision, ensuring minimal disruption of live network operations. An analytical

and general-purpose workflow is demonstrated to guide the application of COMIX in realistic O-RAN scenarios.

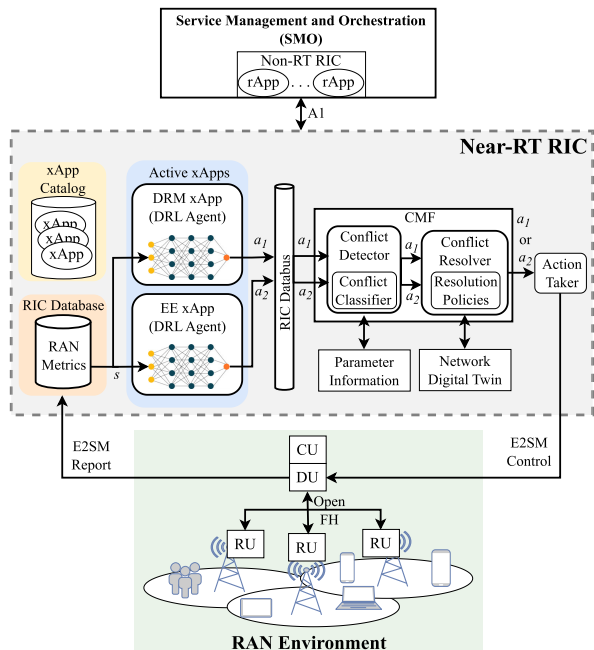


FIGURE 2. General O-RAN architecture considered for COMIX.

For ease of exposure, Table 1 summarizes all the acronyms used throughout this article.

II. ARCHITECTURAL SYSTEM MODEL

This section outlines the overall architecture for realizing an end-to-end and O-RAN-compliant optimization scenario, including O-RAN metrics collection, multi-xApp models' inference, CMF-driven conflict detection and resolution, and final action taking [18]. The architecture considered in COMIX is shown in Fig. 2. The key components of the architecture and their role are described below.

A. COMPONENTS AND THEIR ROLES

Firstly, the SMO hosts the Non-RT RIC, which operates at long timescales to provide high-level guidance and configuration to the Near-RT RIC through the A1 interface. This includes hosting rApps for policy generation, performance monitoring, and non-real-time optimization. The definition and dynamic adaptation of KPI-based policies is the responsibility of the Mobile Network Operator (MNO). This is facilitated through the SMO framework and, more specifically, the Policy Management Function (PMF) of the Non-RT RIC. The PMF enables the MNO to formulate intent-based policies that define the optimization targets (e.g., throughput maximization, energy minimization, SLA fulfillment) and the relative importance of different KPIs. These policies are transmitted to the Near-RT RIC via the A1 interface, where they are used to guide the CMF during the resolution of conflicting xApp actions. In this context, the use of the standardized A1-AP interface enables the semantic

abstraction of intent-based policies in the PMF enables adaptive policy updates to reflect changing network objectives, such as transitioning from peak-performance operation to energy-aware operation under traffic fluctuations.

Thus, the A1 interface enables the exchange of policy and configuration data between the Non-RT RIC and Near-RT RIC, ensuring alignment between strategic and real-time control actions. For instance, vendor-specific or network-defined policies may be enforced from the Non-RT RIC to the Near-RT RIC through A1, imposing xApps and/or CMF component to comply with these policies (e.g. the system-level energy consumption must be below a certain threshold, or the experienced data rate by certain users should exceed a threshold) [19]. Such policies are integrated within Conflict Resolver (noted as Resolution Policies) to be taken into account when potential conflicts are evaluated.

The Near-RT RIC lies at the core of the architecture and enables real-time optimization and conflict management through its key components:

- 1) **xApp Catalog:** Contains a repository of all available xApps originated by different vendors. xApps can be activated or instantiated based on network requirements. For example, a Data Rate Maximization (DRM) xApp and an EE xApp may be active at a certain period for purposes of maximizing the users' throughput and system EE, respectively [20].
- 2) **Active xApps:** All xApps that are active to provide control decisions (e.g. DRM and EE xApps). These xApps may leverage supervised, unsupervised, or DRL algorithms to optimize RAN resources (e.g., regulate RU power levels for DRM). In the COMIX scenario, the DRM xApp focuses on maximizing system throughput, while the EE xApp optimizes energy efficiency. Both xApps receive an O-RAN state vector s and provide action suggestions (a_1 and a_2) or control decisions, which are further evaluated by the CMF component for potential conflicts. Note that the CMF may be fed with multiple candidate actions when more xApps are active in order to assess conflicting xApp groups.
- 3) **RIC Database:** Maintains RAN metrics and potential historical performance data gathered from the network through the E2 Service Model (E2SM) interface. These data are accessible as E2SM Reports to both active xApps and CMF for informed decision making.
- 4) **RIC Databus:** The messaging layer for data provisioning within Near-RT RIC. Usually, CMF is subscribed to RIC Databus for purposes of receiving or consuming new data from the xApps output. xApps often act as data providers for the RIC Databus, bridging communication with the CMF.
- 5) **CMF Component:** The CMF is responsible for detecting, classifying, and resolving conflicts between xApps. It consists of:

- *Conflict Detector* which identifies and classifies conflicts (e.g. direct, indirect, or implicit) between the xApps' proposed actions based on predefined rules and historical data. This component uses *Parameter Information* data to define conflicts such as the recently changed parameters and the parameters groups (a parameters group is defined as a cluster of configuration parameters that affect the same KPI). The internal functionality of the Conflict Detector, called *Conflict Classifier*, is further described in Section II-B.

- *Conflict Resolver* which employs various Resolution Policies, such as priority-based, SLA-based, data rate-targeted, or fairness-driven approaches, to resolve conflicts. The resolution process also integrates the NDT, which simulates the impact of candidate actions on network performance before selecting the optimal action to execute [21]. NDT-driven conflict resolution is advantageous against any other reactive mechanism, which evaluates the decisions after their application on the live network, allowing pro-active knowledge about each decision's impact.

- 6) **NDT:** The NDT is a virtualized model of the live network that predicts the effects of xApp actions by applying the all candidate decisions on the simulated environment. In the case of DRL models, the NDT can simply use the same environment as used during the training to evaluate different candidate corrective actions. By evaluating potential actions through KPIs in a simulated environment (e.g., overall EE and users' data rate derived upon a_1 and a_2 actions), the NDT improves decision accuracy pre-action, minimizing the risk of adverse impacts on the live network. Note that the absence of NDT capabilities forces the Conflict Resolver to implement non-guided actions (based on predefined rules or patterns), potentially leading to network reconfiguration when the action is degradative, further imposing network disruption.
- 7) **RAN Environment:** Consists of the E2 RAN nodes, namely the RUs, DUs and CUs. RUs are managed through the Open Front Haul (Open FH) interface, enabling flexible, multi-vendor configurations. The RUs serve UEs and are directly influenced by the actions taken by the Near-RT RIC. Note that the UE traffic pattern, UE metrics, UE demands, UE mobility, UE handovers, inter-cell interferences, channel losses, and antenna system parameters (transmitting power levels, UE/RU association, physical resource blocks) are considered part of the RAN environment. Actions imposed by the Near-RT RIC's *Action Taker* are conveyed over E2 interface as E2SM Control messages, whereas the RAN metrics acknowledged to the RIC Database are collected as E2SM Report messages over E2.

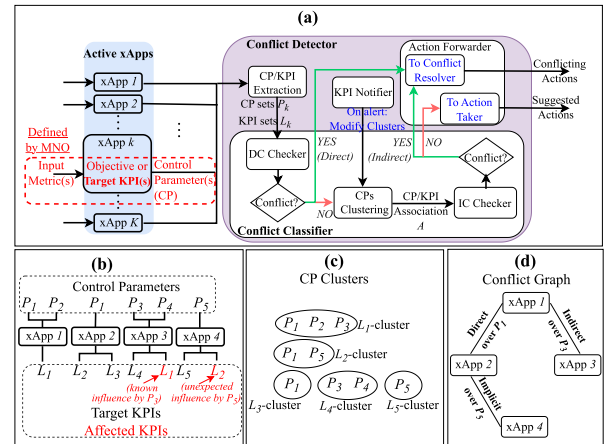


FIGURE 3. Generalized Direct/Indirect Conflict Detector. (a) Internal Architecture; (b) Example of CPs/KPIs associated with 4 xApps; (c) CP Clusters per KPI; (d) Conflict Graph.

B. GENERALIZED xAPP CONFLICT CLASSIFICATION

This section outlines a general-purpose *Conflict Detector* (CD) used in the COMIX framework. For the ease of exposure, Table 2 lists all the mathematical notations used in this work. The role of the CD is to identify pro-actively (i.e. beforehand) direct conflicts (DCs) and indirect conflicts (ICs). Since implicit conflicts can only be detected upon unexpected KPI degradation occurrences, CD supports the identification of implicit conflicts re-actively (i.e. after their negative impact). As shown in Fig. 3a, CD contains several internal logical modules to analyze, classify and forward potential contradicting decisions made by different xApps. Let $\mathcal{K} = \{1, 2, \dots, K\}$ denote the set of active xApps identifiers at a given time slot. Each xApp $k \in \mathcal{K}$ is characterized by its input metrics which are used for inferring the xApp model to obtain output Control Parameters (CPs), referred to as xApp decisions. Decisions about CPs configuration targets at optimizing a specific objective or a target KPI (e.g. power consumption below a threshold), according to which the xApp model has been trained. In some cases, an xApp might provide CPs that affect not only its target KPIs, but also KPIs of other xApps (e.g. EE-targeted power control of RUs by one xApp affects the data rate of UEs, which might be the objective of another xApp). In this sense, we assume that MNO -which serves as the xApp owner- is responsible for predefining the xApp inputs, CPs, and all *a priori* known KPIs that are affected by the CPs [14]. Once the active xApps complete their model inference, all the CP decisions enter the CD (via the RIC Databus) for potential identification of contradicting proposed actions. The CD internal architecture is structured as follows:

- 1) **CP/KPI Extraction:** This module extracts the basic information (i.e. CPs and KPIs) from each active xApp $k \in \mathcal{K}$. Firstly, it identifies the different parameters and KPIs received from all xApps, denoting $\mathcal{P} = \{P_1, P_2, \dots, P_X\}$ as the CP set and

$\mathcal{L} = \{L_1, L_2, \dots, L_Y\}$ as the KPI set. As output, it simply provides the CP set $\mathcal{P}_k \subseteq \mathcal{P}$ (i.e. CPs controlled by xApp k) and the KPI set $\mathcal{L}_k \subseteq \mathcal{L}$ (i.e. KPIs affected by xApp k).

- 2) **DC Checker:** Occurrences of direct conflicts between xApps are checked by the DC Checker. The tests involved in this module search for combinations of xApps with at least one common parameter. Thus, the DC Checker loops over all CPs and returns the xApp groups with common CPs, indicating direct conflicts over specific control parameters. The xApp decisions with direct conflicts are passed to the *Action Forwarder*, whereas CP/KPI sets are also forwarded to the *CPs Clustering* module for further indirect conflict investigation.
- 3) **CPs Clustering:** This module creates a cluster for each KPI $L_y \in \mathcal{L}$ which contains all the CPs that affect L_y ($y = 1, 2, \dots, Y$). Clusters are formed based on the Association Matrix A (dimensions $X \times Y$) with elements $A_{i,j} = 1$, when CP P_i affects KPI L_j , or $A_{i,j} = 0$, otherwise.
- 4) **IC Checker:** Indirect conflicts are detected by the IC Checker, which searches for xApp groups with CPs belonging in the same cluster. In this way, IC Checker is able to detect which xApps influence common KPIs and, in case of IC, xApps decisions will be sent to the Conflict Resolver, otherwise they will be directly passed to Action Taker due to conflict absence.
- 5) **Action Forwarder:** This module establishes the proper end-points or interfaces (e.g. by exposing APIs) to forward the xApp decisions towards either the Conflict Resolver (for contradicting decisions) or the Action Taker (for conflict-free decisions). Hence, if DCs or ICs are raised by the DC or IC Checkers, xApp decisions are passed to the component called *To Conflict Resolver*, otherwise they are passed to the component called *To Action Taker*.
- 6) **KPI Notifier:** To partially support the detection of implicit conflicts, the proposed CD includes a KPI Notifier. This module keeps track of KPI monitoring data (e.g. stored in RIC Database) and raises alarms when an unexpected KPI degradation is observed due to the control of a specific CP. Thus, if an alert is triggered, KPI Notifier enforces a manual modification of the CP clusters by extending the CP group affecting the degraded KPI.

Note that, although the above descriptions concern a single time instance, the CP clusters are updated dynamically under the following conditions: (i) When new xApps are instantiated and registered, their reported CPs and target KPIs are integrated into the existing clusters; (ii) When a new control action is proposed by an xApp, the CP/KPI relationships are re-evaluated to detect any potential conflicts, triggering an update of the relevant clusters; (iii) When the KPI Notifier detects an unexpected degradation in a KPI (e.g., a drop in throughput or energy efficiency due to a specific

TABLE 2. Mathematical symbols.

Symbol	Meaning	Symbol	Meaning
\mathcal{K}	Set of active xApps	K	Number of active xApps
\mathcal{P}	Set of all CPs	P_i	Control parameter i
\mathcal{L}	Set of all KPIs	L_j	KPI j
X	Total number of CPs	Y	Total number of KPIs
\mathcal{P}_k	Set of CPs controlled by xApp k	\mathcal{L}_k	Set of KPIs affected by xApp k
A	Association matrix	$A_{i,j}$	Impact index (1 if CP P_i affects KPI L_j , else 0)
\mathcal{N}	Set of RUs	N	Number of RUs
\mathcal{M}	Set of RBs	M	Number of RBs
W_m	Bandwidth of RB m	s	Service identifier
\mathcal{U}	Set of UEs	U	Total number of UEs
D	Demand vector of all UEs	d_u	Demand of UE u
$A_{n,m,u}$	Association index (1 if UE u is served by RB m of RU n , else 0)	$p_{n,m}$	Power level allocated to RB m of RU n
\bar{P}	Maximum power RUs	\bar{p}	Minimum power of RBs
$\gamma_{n,m,u}$	SINR of UE u over RB m of RU n	n_0	Received noise power at the UE
$L_{n,m,u}$	Path loss between UE u and RU n over RB m	F_1	System data rate
$r_{n,m,u}$	Data rate of UE u over RB m of RU n	F_2	System EE
\mathcal{T}	Set of time slots	T	Number of time slots
\mathbf{p}	Power vector of all RUs and RBs	ξ	Throughput coefficient in EE xApp
$CQIn,m$	CQI of RB m of RU n	$\mathbf{s}(t)$	State vector at time t
$\mathbf{V}(t)$	Variation vector at time t	P_s	Power step
$\mathbf{r}_1(t)$	Reward received at time t by DRM xApp	$\mathbf{a}_1(t)$	Action proposed at time t by DRM xApp
$\mathbf{r}_2(t)$	Reward received at time t by EE xApp	$\mathbf{a}_2(t)$	Action proposed at time t by EE xApp

CP's influence), it raises an alert, triggering a manual update of the CP clusters by extending the CP group associated with the degraded KPI.

To concretely explain the CD information flow and components' functionality, let us consider the example of Fig. 3b, where 4 active xApps are considered. The set of all CPs is $\mathcal{P} = \{P_1, P_2, P_3, P_4, P_5\}$ and the set of all KPIs is $\mathcal{L} = \{L_1, L_2, L_3, L_4, L_5\}$. Initially, the CP/KPI Extraction component defines the CP/KPI sets per xApp as: $\mathcal{P}_1 = \{P_1, P_2\}$, $\mathcal{L}_1 = \{L_1\}$, $\mathcal{P}_2 = \{P_1\}$, $\mathcal{L}_2 = \{L_2, L_3\}$, $\mathcal{P}_3 = \{P_3, P_4\}$, $\mathcal{L}_3 = \{L_4, L_1\}$, $\mathcal{P}_4 = \{P_5\}$, $\mathcal{L}_4 = \{L_5\}$. For the ease of exposure, we assume that, initially, it is not known that P_5 affects KPI L_2 . Noteworthy, KPI L_1 is a target KPI of xApp 1, but is unavoidably affected by the operation of xApp 3 due to the variation of P_3 . Then, the DC Checker tests for common elements (i.e. same CPs regulated by different xApps) by computing the intersections across the CP sets. In this example, there is a DC between xApp 1 and xApp 2 since $\mathcal{P}_1 \cap \mathcal{P}_2 = \{P_1\}$, meaning that both control P_1 . Upon DC detection, the decisions of xApp 1 and xApp 2 are passed to the Conflict Resolver, whereas all CP/KPI sets are further forwarded to the CPs Clustering. The latter forms the following clusters: L_1 -cluster = $\{P_1, P_2, P_3\}$, L_2 -cluster = $\{P_1\}$, L_3 -cluster = $\{P_1\}$, L_4 -cluster = $\{P_3, P_4\}$, L_5 -cluster = $\{P_5\}$. Based on these clusters (see Fig. 3c), IC Checker identifies that P_3 is a member of L_1 -cluster, so it affects the operation of xApp 1 while being only a direct CP of xApp 3. Thus, the decisions of xApp 1 and xApp 3 form an IC and are passed to the Conflict Resolver, whereas

xApp 4 decision is passed directly to the Action Taker. At a later time slot, KPI Notifier can trigger an alarm due to an unexpected degradation of KPI L_2 attributed to P_5 . This in turn enforces a modification of L_2 -cluster by adding P_5 . Including P_5 in L_2 -cluster leads the IC Checker to detect an IC between xApp 2 and xApp 4. Thus, KPI Notifier allows for detecting implicit conflicts over the long-term network operation, which are finally identified as ICs. Fig. 3d depicts the overall conflict graph, including the implicit conflict detected later in time (relative to the proactively detected DC/IC conflicts), after observing an unexpected degradation of L_2 .

III. DESIGN OF xApps

This section describes the mathematical background for designing and training two DRL-based algorithms used as Near-RT RIC xApps. Both xApps operate in the same multi-cell RAN environment, aiming to provide multi-channel power control policies for O-RAN RUs.

A. DOWNLINK MULTI-CHANNEL POWER CONTROL

Solving the optimization problem of Downlink Multi-channel Power Control (DMPC) involves finding the proper power-level configuration of all RUs so as to optimize a specific objective, under power constraints, user demands, mobile traffic, inter-channel and inter-cell interferences. We consider an O-RAN wireless network consisting of N RUs from the set $\mathcal{N} = \{1, 2, \dots, N\}$, where each RU serves a coverage area. In this setup, the operational frequency band is segmented into multiple RBs based on a given channel segmentation. Each RU supports M available RBs (from the set $\mathcal{M} = \{1, 2, \dots, M\}$) for downlink physical-layer transmissions, where each RB $m \in \mathcal{M}$ has a specific bandwidth W_m . A complete frequency reuse scheme is assumed across the O-RAN transmissions, meaning all RUs operate on the same frequency band, introducing inter-cell interferences over common-RB transmissions.

Each UE in the network, denoted as $u \in \mathcal{U}$ (where $\mathcal{U} = \{1, 2, \dots, U\}$), requests a specific service s based on its SLA profile in order to meet its QoS requirements. Each service s corresponds to a particular throughput demand, expressed in Mbps. To represent these demands, a demand vector D with elements d_u is used to notify the requested service class of each user u . In the proposed system model, each UE u is associated with a single RB m from a specific RU n , determined by the binary association matrix A , where each element $A_{n,m,u} = 1$ if user u is associated with RB m of RU n , or 0 otherwise. The power level allocated to RB m of RU n is represented as $p_{n,m}$. To ensure power efficiency, the system imposes a sum-power constraint on each RU n , as given by $\sum_{m=1}^M p_{n,m} \leq \bar{P}$, where \bar{P} is the maximum total transmitting power allowed for each RU. Additionally, a minimum power level constraint $p_{n,m} \geq \bar{p}$ is enforced to account for signaling and sleeping operations of each RB. Also, each RU can cover multiple users.

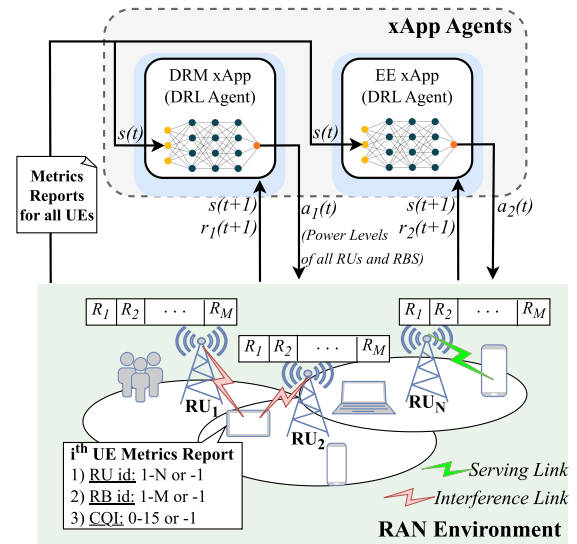


FIGURE 4. DRL cycles between xApps and O-RAN environment.

The Signal-to-Interference-plus-Noise Ratio (SINR) experienced by a user u , associated with RB m of RU n , is given by:

$$\gamma_{n,m,u} = \frac{p_{n,m} L_{n,m,u}}{\sum_{t \neq k} p_{t,m} L_{t,m,u} + n_0} \quad (1)$$

where $L_{n,m,u}$ denotes the channel losses between RU n and UE u over RB m , and n_0 represents the received noise power at the UE. The achievable data rate $r_{n,m,u}$ for a user $u \in \mathcal{U}$ associated with RB m of RU n is computed using the Shannon formula:

$$r_{n,m,u} = W_m \cdot \log_2(1 + \gamma_{n,m,u}) \quad (2)$$

The user association is determined so as to ensure that each UE receives the highest achievable rate. Note that the system-level data rate is defined as:

$$F_1 = \sum_{n=1}^N \sum_{m=1}^M \sum_{u=1}^U A_{n,m,u} \cdot r_{n,m,u} \quad (3)$$

Similarly, the system-level EE is defined as:

$$F_2 = \frac{\sum_{n=1}^N \sum_{m=1}^M \sum_{u=1}^U A_{n,m,u} \cdot r_{n,m,u}}{\sum_{n=1}^N \sum_{m=1}^M p_{n,m}} \quad (4)$$

B. OPTIMIZATION PROBLEM VIA DRL

To regulate power control across RUs intelligently, the DMPC problem is solved using a DRL framework, implemented as two distinct xApps with different objectives: (i) the DRM xApp aims to optimize power allocation to maximize the total system throughput for UEs [22]; (ii) the EE xApp selects the power allocation to minimize energy consumption while maintaining QoS requirements [23].

The intelligent agent running on the Near-RT RIC adheres to the following constraints: (i) respecting the total power budget \bar{P} of RUs and the minimum power limitations \bar{p}

of RBs; (ii) considering the available RBs for each RU; (iii) satisfying the QoS demands specified by each UE's SLA profile. Each UE reports the following metrics to the Near-RT RIC's Database: Associated RU ID, Associated RB ID, Current data rate in Mbps. Fig. 4 shows the DRL interaction cycles between the DRM/EE xApps and the RAN environment.

The simultaneous operation of DRM and EE xApps introduce conflicts, so the goal of the COMIX scheme is to ensure that RUs dynamically adjust their power levels while balancing the conflicting objectives based on predefined resolution policies. These policies are enforced by the Non-RT RIC and are applied by the Conflict Resolver. Below, we formulate the optimization problem defined for each xApp, and we determine the DRL elements (state, action, reward) considered in the solution [24].

1) DATA RATE MAXIMIZATION xApp

The DRL-based DRM xApp focuses on solving the following optimization problem:

$$\begin{aligned}
& \arg \max_{\mathbf{p}} \left(\frac{1}{T} \sum_{t=1}^T F_1(\mathbf{p}) \right) \\
\text{s.t. (C1)} \quad & \sum_{m=1}^M p_{n,m} \leq \bar{P}, \quad \forall n \in \mathcal{N} \\
\text{(C2)} \quad & p_{n,m} \geq \bar{p}, \quad \forall n \in \mathcal{N}, m \in \mathcal{M} \\
\text{(C3)} \quad & \sum_{n=1}^N \sum_{m=1}^M A_{n,m,u} \leq 1, \quad \forall u \in \mathcal{U} \\
\text{(C4)} \quad & r_{n,m,u} \geq d_u, \quad \forall u \in \mathcal{U} \quad (5)
\end{aligned}$$

where it is implied that DRM xApp looks for a power vector $\mathbf{p} = \{p_{n,m}\}$ (with dimensions $N \times M$) such that to maximize the expected (long-term) system-level data rate after a series of T training episodes [25].

The following DRL elements are considered by the DRM xApp during the training process:

State: Considering the set of time slots $\mathcal{T} = \{1, 2, \dots, T\}$, the DRL agent observes a state vector $\mathbf{s}(t)$ at each time $t \in \mathcal{T}$. The state vector is defined as $\mathbf{s}(t) = \{CQI_{n,m}\}$ (dimensions $N \times M$), where each element $CQI_{n,m}$ denotes the channel quality indicator (CQI) of the UE that is served by RB m of RU n . CQI values vary between 0 – 15 and represent the quantized version of the SINR values $\gamma_{n,m,u}$ into 16 discrete levels (i.e. zero CQI corresponds to poor SINR levels, CQI at 15 means perfect propagation conditions), as defined in [24]. When no UE is associated with a certain RB m of RU n , we note $CQI_{n,m} = -1$ by convention.

Action: After observing $\mathbf{s}(t)$, DRL agent provides a control decision on the RUs power configuration by assigning a new power vector $\mathbf{a}_1(t) = \mathbf{p}(t) = \mathbf{p}(t-1) + \mathbf{V}(t)$. Specifically, the new power vector is equal to previous-slot power vector changed by the variation vector $\mathbf{V}(t)$. The latter is formed by selecting one RB per RU and either increase or decrease its transmitting power by a predefined power step P_s . For

instance, the variation vector:

$$\mathbf{V}(t) = \begin{bmatrix} 0 & +P_s & \cdots & 0 \\ -P_s & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & +P_s \end{bmatrix}_{N \times M} \quad (6)$$

reflects that RB 2 of RU 1 and RB M of RU N will be increased by P_s (in Watts), whereas the RB 1 of RU 2 will be decreased by P_s . The power level of the rest RBs will be the same as in the previous-slot power level. In this formulation, DRL agent is able to apply smooth changes in the power configuration across RUs to optimize objective F_1 .

Reward: After applying an action $\mathbf{a}_1(t)$, the DRL agent receives a scalar reward $\mathbf{r}_1(t+1)$ reflecting the extent to which the action is beneficial relative to the objective. In the case of DRM xApp, the reward is defined as $\mathbf{r}_1(t+1) = F_1(t) - F_1(t-1)$, where $F_1(t)$ and $F_1(t-1)$ are derived by applying (3) considering the power configuration $\mathbf{p}(t)$ and $\mathbf{p}(t-1)$, respectively. In this way, the DRL agent's reward reflects the data rate increment/decrement derived after applying $\mathbf{a}_1(t)$.

2) ENERGY EFFICIENCY MAXIMIZATION xApp

Similarly, the EE xApp aims to solve the following optimization problem:

$$\begin{aligned}
& \arg \max_{\mathbf{p}} \left(\frac{1}{T} \sum_{t=1}^T F_2(\mathbf{p}) \right) \\
\text{s.t. (C1)-(C3)} \quad & \\
\text{(C4')} \quad & r_{n,m,u} \geq \xi \cdot d_u, \quad \forall u \in \mathcal{U} \quad (7)
\end{aligned}$$

where EE xApp searches a power vector \mathbf{p} for all RUs/RBs to ensure a maximized data rate vs power consumption ratio [26]. Note that constraint (C4') implies that the allocated throughput of each UE can be decreased for purposes of ensuring power savings, given that $\xi \in [0, 1]$ denotes a constant value that allows for reduced throughput allocation at the UE (e.g., $r_{n,m,u} \geq 0.8 \cdot d_u$). In this manner, DRM xApp ensures maximized throughput and full QoS satisfaction (due to (C4)) while ignoring the power consumption, whereas EE xApp guarantees low power consumption while allowing for a slight degradation of the requested throughput (due to (C4')) by a factor $(1 - \xi)\%$. Overall, both xApps optimize the O-RAN from different perspectives, highlighting the full QoS satisfaction vs EE trade-off.

To provide a suboptimal solution, the EE xApp adopts the same state/action space as in DRM, but the rewarding function is given by $\mathbf{r}_2(t+1) = F_2(t) - F_2(t-1)$, representing the EE change due to the application of action $\mathbf{a}_2(t)$. In summary, both xApps observe the same environment state $\mathbf{s}(t)$ and propose different actions (i.e. $\mathbf{a}_1(t)$ and $\mathbf{a}_2(t)$) such that to optimize F_1 and F_2 , respectively. For convenience, a comparative summary of the state, action, and reward definitions for both xApps is presented in Table 3.

TABLE 3. Comparison of DRL elements for DRM and EE xApps.

Element	DRM xApp	EE xApp
State $s(t)$	CQI of each RU/RB pair	Same as DRM xApp
Action $a(t)$	Increase or decrease the power level of one RB per RU by a predefined power step.	Same as DRM xApp
Reward	$r_1(t+1) = F_1(t) - F_1(t-1)$ (reflects the total experienced throughput summed across UEs)	$r_2(t+1) = F_2(t) - F_2(t-1)$ (reflects the total energy efficiency summed across UEs)
Objective	Maximize long-term network-wide data rate (F_1)	Maximize long-term system-level EE ratio (F_2)

It is important to note that the optimization problems (5) and (7) are non-convex, primarily due to the interference-coupled SINR formulation and logarithmic utility functions across multiple RUs and RBs. As a result, the DRL-based xApps are designed to learn a locally optimal policy, rather than guarantee a globally optimal solution. Hence, the agents are trained in a realistic simulation environment to make near-optimal decisions under dynamic network conditions, with optimality defined in terms of its reward function and achieved performance during training.

C. NETWORK DIGITAL TWIN FUNCTIONALITY

By definition, NDT is a virtual representation of a real network and its associated objects and processes which mirrors physical counterparts, allowing for real-time monitoring, simulation, and analysis. The NDT is an integral component of the COMIX framework, providing the dynamic simulation environment necessary for computing RAN/UE metrics and training the DRL-based xApps. The NDT models the O-RAN network environment, enabling the DRL agents to interact with a realistic, yet simulated, representation of the physical-layer RAN conditions [21]. In the case of DRL algorithms, NDT should only coincide with the environment in which the agents perform actions, and is not necessarily a large-scale representation of the whole network. In this sense, NDT can be deployed as different instances depending on the DRL algorithm, with each instance abstracting only the environment required by a given DRL agent [27]. This ensures that the xApps can learn optimal policies (e.g., for multi-channel power control in RUs) before deployment in the live network.

In the COMIX framework, the NDT allows for flexible configuration of the simulation parameters to accurately represent the O-RAN environment, and can be parameterized on-demand. The key configurable parameters include:

- The number and topology of cells within the network area, along with the spatial location of the RUs.
- The channel and path loss models to estimate the received signal strength type-specific cells (e.g., macro, micro, or pico-cells; urban, suburban, or rural cells) [28], [29].

- The maximum power budget per RU and the minimum power level of each RB.
- The operating 5G frequency band, numerology (spectrum channelization), and the total bandwidth allocated for downlink transmissions [29].
- The initial number of UEs in the network area, their maximum velocities under a mobility model (e.g. random walk), and their potential movement across cell boundaries.
- The probability of new UEs entering the network area from the borders during a given time slot.

1) PARAMETERIZING NDT DURING xApp TRAINING

The NDT initializes the simulation by randomly placing UEs within the network and assigning initial power levels to each RB, respecting the constraints on maximum and minimum power budgets. It computes channel gains based on the relative distance between UEs and RUs, following the 3GPP specifications for 5G cells [29]. Using this information, the NDT calculates key metrics such as SINR and data rate for each UE across all RU and RB combinations. The core functionalities of the NDT include: (i) Interference calculation to evaluate the accumulated interference from non-serving RUs for each UE; (ii) Channel metrics production such as SINR and data rate for all potential RU/RB associations; (iii) Dynamic association to assign UEs to RBs that maximize their throughput; (iv) UE Mobility through a pre-defined model. This dynamic modeling supports time-varying network conditions, including potential handovers, ensuring that the DRL agents are trained in realistic and variable scenarios.

At each simulation time step, the NDT generates detailed UE metrics that can be used to form the DRL state vector, including the CQI and the IDs of the associated RU and RB. By simulating diverse network conditions during the DRL training, the NDT enables the xApps to refine their policies in a controlled environment, ensuring they are well-prepared for deployment in live O-RAN systems.

2) PARAMETERIZING NDT DURING xApp INFERENCE AND CONFLICT RESOLUTION

When the xApp model is inferred with real UE metrics collected via E2, NDT can be configured instantly using the current UE metrics (positions, serving RUs and RBs) and the O-RAN parameters (RU power levels, topology). Also, when potential xApp conflicts are addressed by the Conflict Resolver, xApps decisions can be directly applied on the current NDT instance to gather the new O-RAN metrics. The Conflict Resolver then uses these metrics to evaluate which xApp complies with the current resolution policy (e.g. which xApp leads to maximum EE).

It is important to note that, in COMIX, the evaluation of conflicting actions by the NDT occurs before deployment to the live network, enabling proactive and informed decision-making. Although such pre-deployment simulation incurs computational cost, the process should be designed to be

lightweight and fast. Usually, the NDT instance is virtualized, containerized, and runs on dedicated infrastructure separate from the live RAN control loop, allowing for parallel execution without introducing blocking overhead. The simulation environment leverages reduced-order models, compressed radio maps, and pre-trained predictors to approximate action impact efficiently. Moreover, the responsiveness of the NDT depends on several implementation-specific factors, such as CPU/GPU capabilities, simulation granularity, and iteration length. To ensure minimal latency (typically within milliseconds), the system can adopt lightweight evaluation strategies, including: (i) limiting the number of simulation iterations per candidate action, (ii) applying early-stopping heuristics when actions exhibit poor KPI trends during the initial steps; and (iii) selecting representative yet simplified network states for evaluation instead of full-scale replicas. As a result, the use of the NDT remains practical and scalable for near-real-time conflict resolution scenarios in operational O-RAN environments.

3) SYNCHRONIZATION AND DEPLOYMENT CONSIDERATIONS

In COMIX, the same NDT instance can serve multiple purposes throughout the xApp lifecycle: during training, it offers a controlled simulation of mobility-aware multi-user RAN scenarios, while during inference or conflict resolution, it rapidly emulates the impact of xApp actions based on current E2-reported metrics [1], [30]. This reuse of a consistent modeling engine helps ensure that the xApps' learned behavior aligns with inference evaluations. To maintain accuracy, the NDT periodically synchronizes with live RAN data collected from the SMO and Near-RT RIC, including updates on propagation environments, UE mobility trends, and evolving power configurations. These updates allow the simulation to remain relevant as the physical network changes over time [21].

In terms of deployment, the NDT is typically implemented as a containerized microservice (e.g., Docker-based) that is invoked by the RAN operator (or it is automatically deployed with Near-RT RIC), depending on implementation selections [21], [30]. Depending on the deployment architecture, it may run on the same platform as the Near-RT RIC or on edge/cloud servers within the RAN domain. Its modularity enables the support of multiple NDT instances, each fine-tuned for a specific xApp or policy evaluation scenario, running in parallel without mutual interference [21].

IV. SIMULATION RESULTS

This section presents the validation results of the COMIX scheme. First, we demonstrate the impact of the key learning hyperparameters to ensure optimal convergence of individual xApps. Then, we assess the network performance in terms of target KPIs (i.e. system-wide data rate and EE) under different conflict resolution policies using multiple simulation scenarios. Finally, considering the pre-trained DRM and EE xApps in inference mode, we analytically present the

end-to-end COMIX workflow in a general-purpose manner, explaining the information flow step-by-step and showcasing the role of CMF and NDT components in addressing xApp-level conflicts. The constant hyperparameters used in the simulations are tabulated in Table 4. It is important to note that the purpose of the following experiments is not to prove that DRM or EE xApps outperform state-of-the-art optimization algorithms, but rather to demonstrate the applicability of the proposed COMIX framework in resolving xApp-level conflicts through standardized workflows and configurable resolution policies.

TABLE 4. Simulated environment and DRL configuration.

Parameter	Value
Number Of RUs (N)	3
Number Of RBs (M)	8
Number Of UEs (U)	12
Type of cells	Urban
Maximum RU Power (\bar{P})	80 W
Minimum RB Power (\bar{p})	0.1 W
UE demands (d_{ii})	1 – 10 Mbps
Throughput coefficient in EE xApp (ξ)	0.8
Training Episodes (T)	2000
Time Steps Per Episode	200
UE Velocity	5 m/s
UE Move Probability	50%
Batch Size	32
Replay Memory Capacity	5000
Target Q -Network Update Frequency	6

A. TRAINING PERFORMANCE OF xApp

Extensive simulations were carried out to ensure convergence of the algorithms. Both DRM and EE xApps were trained in identical environment setup after 2000 training episodes. The xApps only differed in the rewarding function, as specified in (3) and (4), respectively. Although acting on the same environment with identical available actions, different rewards led the xApps to learn completely different policies.

To guarantee optimal convergence in DRL algorithms, the critical hyperparameters should be stabilized. Upon experimentation on different values of learning parameters, the most impactful hyperparameters were deemed to be the learning rate α (affects the intensity of weight update during backpropagation) and the discount factor γ (balances the trade-off between immediate and future rewards). Additionally, as a problem-specific hyperparameter, different values of the power step were tested, as they affect the granularity of RU power regulation. To assess the impact of each hyperparameter, we depict the learning curve of each DRL configuration, which reflects the reward time-course as a function of the training episodes. Fig. 5a-c shows the reward of different DRM xApp configurations in terms of the overall achieved data rate (in Mbps), whereas Fig. 5d-f depicts the same curves for EE xApp, representing the reward in terms of the system-wide EE (in Mbps/Watt). Due to the large number of overlapping learning curves in Fig. 5, confidence intervals were omitted for visual clarity. However, we report that the 95% confidence interval for the average reward values across

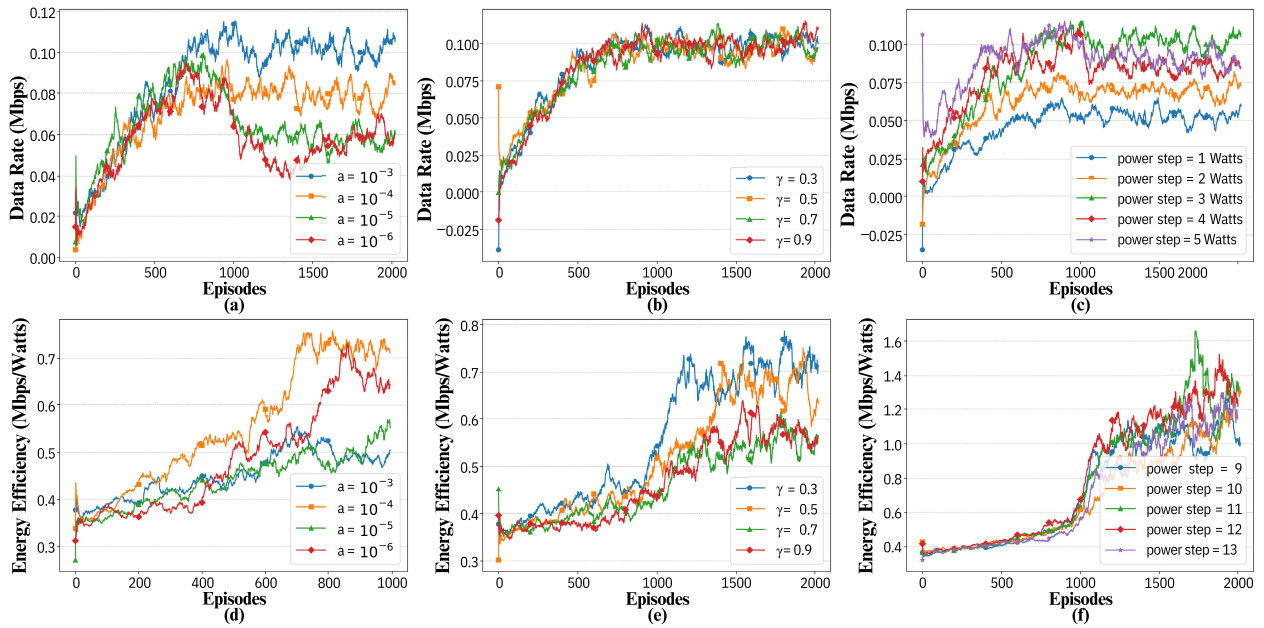


FIGURE 5. Learning curves of DRM (a-c) and EE (d-f) xApps for different learning rates (column 1), discount factors (column 2), and power steps (column 3).

multiple runs was within $\pm 5.3\%$ for DRM and $\pm 6.1\%$ for EE xApps under the optimal hyperparameter settings.

As observed, the DRM xApp reward converges to a similar value across all γ values due to the limited variance in achievable system throughput under the given setup. In contrast, the EE xApp reward diverges across γ values, with $\gamma = 0.3$ presenting the best EE, reflecting the higher sensitivity of energy efficiency optimization to long-term planning and discounting. The different impact of discount factor on different xApps further highlights the distinct objectives defined, with the EE xApp requiring low discount values for optimality. Also, the two xApps showed optimal convergence for different learning rates, namely $a = 10^{-3}$ for the DRM xApp and $a = 10^{-4}$ for the EE xApp.

Considering the optimal a and γ values, there was a noticeable impact of power granularity on long-term reward achievement. Specifically, DRM xApp presented optimal performance for a power step of $P_s = 3$ Watts, leading to an average data-rate improvement of 0.11 Mbps. This means that following the DRM decisions at a given time slot, the overall system throughput may be increased by 0.11 Mbps at the next time slot. Considering the continuous operation of DRM xApp, consecutive decisions may lead to considerable data-rate improvements. On the other hand, EE xApp presented performance improvements for higher values of P_s (ranging 9 to 13 Watts), with a $P_s = 11$ Watts showing the highest reward (about 1.3 Mbps/Watt). Note that, since EE xApp allows for slight throughput degradation (by a factor of $1 - \xi$) to achieve power reduction (see (C4') of problem (7)), it is capable of changing more rapidly (i.e. by larger power steps) the power level of the RUs.

B. COMPARISON OF DIFFERENT RESOLUTION POLICIES

Following the hyper-parameter tuning of each xApp individually, this subsection evaluates the CMF resolution performance under various resolution policies. Since both xApps propose actions concerning the power control of RUs, different resolution policies can be defined to account for handling the direct conflict. To this end, comparative analyses were conducted considering five different resolution policies to illustrate their impact on network KPIs such as system-level data rate and EE. A simulated DRL environment is assumed to be hosted within NDT for purposes of evaluating different candidate actions for power control. To ensure consistency, the NDT environment considered for the proactive action assessment is the same software package used to train both xApps. Under a given resolution policy i at time slot t , NDT returns an evaluation score $s_j^i(t)$ for each candidate action j , which quantifies the extent to which action j benefits the target KPI defined by policy i . Finally, the xApp action with the maximum evaluation score is chosen.

The resolution policies considered were the following:

- 1) **Maximum Throughput-based Selection (MaxTS):** At time slot t , this policy selects the action of xApp j ($j = 1$ for DRM xApp or $j = 2$ for EE xApp) that maximizes the system-wide throughput by computing the following score for each candidate action:

$$s_j^1(t) = F_1^j(t) \tag{8}$$

where $F_1^j(t)$ is given by (3) after applying action j .

- 2) **Minimum Power-based Selection (MinPS):** This policy resolves the direct conflicts by selecting the power vector that minimizes the total power consumed.

Thus, the score of each action j is computed by:

$$s_j^2(t) = - \sum_{n=1}^N \sum_{m=1}^M p_{n,m}^j \quad (9)$$

where $p_{n,m}^j$ is the power level assigned to RB m of RU n based on the action of xApp j .

- 3) **Energy Efficiency-based Selection (EES):** This policy selects the action that provides the best ratio of total network throughput divided by the total consumed power [31].

$$s_j^3(t) = F_2^j(t) \quad (10)$$

where $F_2^j(t)$ is given by (4) after applying action j and reflects the network EE.

- 4) **Throughput Violation-based Selection (TVS):** In this policy, specific SLA requirements are considered for the achieved throughput. We consider an SLA-defined throughput requirement of c_u (in Mbps) for each UE u , which reflects that each UE, upon action selection, should experience a throughput of at least c_u Mbps. Without loss of generality, we assumed $c_u = 2$ Mbps for each UE u . The evaluation score of this policy quantifies how many UEs exhibit SLA violation by counting the number of UEs (denoted as C_u) experiencing a throughput below c_u . In case of score ties among xApps (i.e. multiple actions show the same number of SLA violations), the minimum power-consuming action is chosen. The above-mentioned policy returns a score $s_j^4(t)$ for each action j , as reflected by the following formula:

$$s_j^4(t) = - \sum_{u=1}^U C_u^j(t) - \frac{1}{1 + e^{-p_{total}^j(t)}} \quad (11)$$

where $p_{total}^j(t)$ is the total power consumption (summed across RUs and RBs) after applying action j at time slot t , and $C_u^j(t)$ is the SLA violation (binary) index of UE u after applying action j , which is given by:

$$C_u^j(t) = \begin{cases} 1, & \text{if } r_u^j(t) < c_u \\ 0, & \text{otherwise} \end{cases} \quad (12)$$

where $r_u^j(t)$ is the experienced throughput (in Mbps) of UE u after applying action j . It is evident that this policy scores an action based on the resulting number of SLA violation occurrences (first term) and the total power consumption scaled between $[0,1]$ using the sigmoid function (second term). Given that the number of SLA violations C_u is an integer, the second (sigmoid) term is used to resolve the score ties, ensuring that the contribution of the power consumption on the final score will always be less than one. Under this definition, the second term influences the action selection only when the candidate actions present equality in SLA violations. For instance, if both DRM

and EE xApps' actions result in 5 SLA violations, but they require different total power consumption (e.g., $p_{total}^{DRM}(t) = 3W$ and $p_{total}^{EE}(t) = 2W$, then the returned scores are $s_{DRM}^4(t) = -5 - 0.95 = -5.95$ and $s_{EE}^4(t) = -5 - 0.88 = -5.88$. In this example, EE xApp will be finally applied because it produces the highest score (i.e. equal SLA violations compared to DRM xApp, with lower power consumption).

- 5) **EE Violation-based Selection (EEVS):** Similar to TVS policy, this resolution policy considers a UE-level SLA requirement concerning the lower-bounded EE ratio e_u (i.e. a throughput to consumed power ratio). Considering that each UE u should experience at least $e_u = 2$ Mbps/W after RU power allocation, this policy calculated and returns the following score:

$$s_j^5(t) = - \sum_{u=1}^U E_u^j(t) - \frac{1}{1 + e^{-p_{total}^j(t)}} \quad (13)$$

where $p_{total}^j(t)$ is the total power consumption (summed across RUs and RBs) after applying action j at time slot t , and $E_u^j(t)$ is the EE-based SLA violation (binary) index of UE u after applying action j , which is given by:

$$E_u^j(t) = \begin{cases} 1, & \text{if } \frac{r_u^j(t)}{p_u^j} < e_u \\ 0, & \text{otherwise} \end{cases} \quad (14)$$

where $r_u^j(t)$ and p_u^j is the experienced throughput (in Mbps) and the associated RU/RB power level of UE u , respectively, after applying action j . The EEVS scores quantify how many EE violations occurred and, if score ties between xApps are observed, the least power-consuming action is chosen.

Notably, the proposed framework is generalizable, meaning that additional resolution policies or modifications of the considered policies could be easily adopted without significant changes in the architectural setup.

To demonstrate the performance of each resolution policy, we conducted 200 validation episodes, each of which simulating different network instances with randomized initial state (UE position, RU power levels, RB-UE association). In each validation scenario, we run 100 consecutive steps with both DRM and EE xApps being activated, allowing UEs to randomly move within the network area. In each step, the environment state is acknowledged to both xApps for inference, both producing conflicting power allocation actions that are further detected and resolved by the CMF component.

After applying the five presented resolution policies hosted in the Conflict Resolver, we measure the Network Data Rate (in Mbps) and the resulted Total Power Consumption (in Watts), averaged across the 200 episodes and 100 steps. Both metrics were computed under two conditions: once with the CMF in place (considering a given resolution policy) and once without it. In the CMF-free scenario, the final action




FIGURE 6. Network Data Rate (a) and Total Power Consumption (b) resulted from different resolution policies (blue curves) after 200 validation episodes. Red curves reflect the same metrics without the presence of CMF.

is determined based on the timing of the xApp’s action arrival, effectively resulting in a random selection between the DRM and EE xApps. This baseline method represents a typical scenario where conflict detection is not performed. Consequently, the most recently received action at the Action Taker is directly applied to the network, overwriting any previously proposed action. Notably, to ensure statistical robustness, each of the 200 validation episodes was generated with different random seeds affecting UE initial locations, mobility patterns, and RB/UE associations. These variations were selected to emulate realistic deployment variability and allow multiple randomized comparisons.

In Fig. 6, we depict the total data rate and the power consumption across the five different resolution policies against CMF-free scenarios. Evidently, regardless of the resolution policy used, CMF inclusion does not significantly benefit the network data rate relative to the CMF-free method (around 3% difference). This is attributed to the fact that, at each step, CMF-free scheme leads to the implementation of the action proposed by either DRM or EE xApp, both targeting at the throughput maximization in their objective function (see (3) and (4)). This means that, even without CMF inclusion, both xApps produce optimized decisions

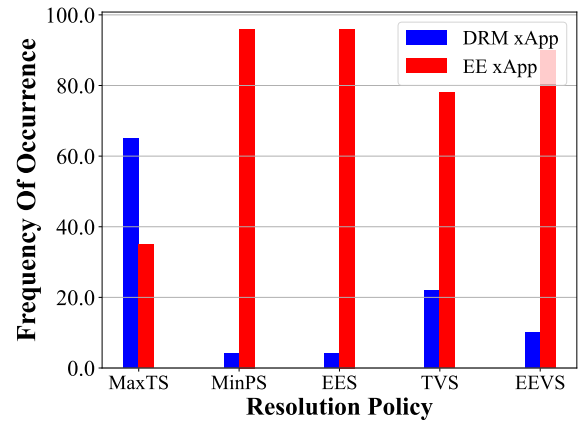


FIGURE 7. Frequency of occurrence for each xApp’s action (DRM and EE) across different resolution policies, showing how many times each xApp’s action was finally selected based on the NDT-oriented evaluation scores.

with regards to network data rates, further explaining the equal performance between the CMF-based and CMF-free schemes in terms of the overall achieved throughput. Noteworthy, the CMF-based scheme outperforms the CMF-free method only when the MaxTS policy is considered, with the latter constantly selecting the xApp that provides the highest network data rate. Also, the lower data rate observed in all resolution policies compared to the MaxTS is explained by the following key point: MaxTS selects the highest-throughput action, neglecting power consumption and whether UEs are over- or under-satisfied. For instance, an action selected by MaxTS may present extremely high network data rate, simply because some UEs experience extremely high data rate (due to their close proximity from the RUs), while many other UEs are under-satisfied. On the contrary, the rest of the methods either take care of the power consumption (MinPS and EES) or follow an SLA-driven selection (TVS and EEVS), which usually lead lower network-wide data rates.

Importantly, the CMF considerably increases the energy savings of the system, as clearly illustrated in Fig. 6b, with power consumption being reduced by up to 60% using CMF. Independently of the resolution policy used, CMF inclusion constantly outperforms the CMF-free scheme, highlighting the advantageous usage of CMF for improving system’s EE. Following MaxTS, the power consumption is increased compared to the rest of the policies, mainly because MaxTS exclusively selects based on UEs’ throughput. The resolution policies that take the power consumption into consideration before selecting the final action (i.e. MinPS, EES and EEVS) showed equally optimal performance in terms of power savings, suggesting that their use can significantly improve EE under the presence of power control-based direct conflicts. Note that, since Fig. 6 shows the average values across 200 validation episodes, the standard error of the mean (SEM) across runs was approximately 3.67 Mbps for network data rate and 5.9 W for total power consumption, indicating consistent trends across random environmental realizations.




FIGURE 8. Sequence diagram and information flow in the COMIX scheme for multi-xApp inference and conflict management.

Fig. 7 depicts how many times each xApp was finally selected during the validation scenarios. Interestingly, EE xApp is chosen in the vast majority of time across all resolution policies. This is attributed to the fact that EE xApp jointly considers the throughput and the power levels, achieving a balanced rate of both KPIs (data rate and power consumption). Note that DRM xApp is selected only when the objective is to maximize the overall data rate (as in MaxTS policy) and is rarely chosen when power control or SLA violations are considered. This behavior does not indicate a framework-level bias, but rather reflects the fact that the EE xApp aligns more closely with the scoring objectives of the considered policies, particularly those that

favor reduced power consumption or SLA satisfaction. The DRM xApp, by design, focuses solely on throughput, and thus performs optimally only under MaxTS. Notably, the framework remains generalizable, and the resolution criteria can be adapted to prioritize different KPI trade-offs based on operator-defined intents.

Note that, to ensure comprehensive conflict detection, the COMIX framework supports both syntactic (direct) and outcome-based (indirect) conflict analysis. First, the DC Checker identifies pairs of xApps that target overlapping CPs, such as transmit power for the same RU/RB pair. This module loops over all CPs and returns the xApp groups with conflicting parameter values, indicating direct conflicts.

For example, the DRM and EE xApps typically propose different power vectors for the same RUs, triggering direct conflict flags. These action pairs, along with their associated CPs and targeted KPIs, are forwarded to subsequent CMF components for resolution. Second, for cases where xApps suggest identical or near-identical actions but differ in predicted outcomes (due to distinct internal models or reward structures) the CMF leverages the NDT to evaluate the utility of each action. This dual mechanism allows COMIX to robustly handle both explicitly conflicting control actions and semantically divergent decisions in multi-objective scenarios.

C. END-TO-END WORKFLOW FOR PRE-TRAINED xApp

To concretely demonstrate the information flow in the COMIX framework, this subsection presents the stepwise workflow to integrate a hierarchical and iterative decision-making under the presence of conflicts [32]. The workflow is depicted as sequence diagram in Fig. 8, highlighting the interactions among the key components of the O-RAN ecosystem, namely the SMO (Non-RT RIC), Near-RT RIC, DRM and EE xApps, CMF, and the RAN environment (DU, RUs, UEs). The main phases and steps are the following:

Phase 1 - Activation of xApps: Upon operator request towards the Near-RT RIC's xApp Catalog, DRM and EE xApps are activated for purposes of DRL model inference. Once activated, these xApps begin functioning as the primary DRL agents in the workflow.

Phase 2 - Policy Enforcement by SMO: Then, SMO (Non-RT RIC) enforces the conflict resolution policies onto the Conflict Resolver component. These policies define how potential conflicts should be resolved based on operator-defined objectives, such as prioritizing throughput or energy efficiency (see Section IV-B). Conflict Detector component subscribes to RIC Database to collect real-time RAN metrics.

Phase 3 - Initial State Fetching: The DRM and EE xApps retrieve the initial RAN metrics, such as the associated RUs and RBs of the UEs and CQI values, from the RIC Database. This ensures that xApps have the necessary information about the current state of the network to make informed decisions.

Phase 4 - xApp Inference and Action Proposals: Both xApps independently generate power control actions (a_1 for DRM and a_2 for EE) based on their DRL policies, as defined in (3) and (4). These proposed actions are then sent to the Conflict Detector for conflict assessment.

Phase 5 - Conditional Conflict Resolution: Conflict Detector assesses the proposed actions as described in Fig. 3. If no conflict is detected, both actions are approved and sent to the RAN environment for direct implementation. If a conflict is detected, the Conflict Detector informs the Conflict Resolver, which interacts with NDT to obtain the evaluation KPI score of each proposed action. The NDT simulates the impact of both actions on the network, considering metrics such as those presented in Section IV-B. Based on these evaluations and the enforced policies, the Conflict Resolver selects the optimal action and applies it to the RAN environment through the Action Taker.

Phase 6 - Action Application in the RAN: The final action, resolved or approved, is applied in the RAN environment (DU or RUs). This adjusts the power levels of the RUs according to the chosen policy, ensuring efficient and conflict-free operation of the network.

Phase 7 - KPI Reporting: Using RIC Database, the RAN environment continuously reports KPI metrics, such as QoS performance, throughput, and energy consumption, back to the SMO. These reports allow the SMO to monitor network performance and refine its policies over time if needed.

Noteworthy, while the validation in this study focused on two conflicting xApps, the COMIX framework is inherently generalizable to more than two active xApps. Specifically, the Conflict Detector operates over the full set of input control parameters and affected KPIs for each xApp, forming clusters of overlapping influences. This enables the detection of direct and indirect conflicts across multiple xApps simultaneously. Additionally, the Conflict Resolver is also agnostic to the number of candidate actions and applies policy-based evaluation to select the most appropriate one based on the specified objective (e.g., throughput, energy efficiency, SLA satisfaction). The use of the NDT further supports this generalizability, as it can simulate and score any number of candidate actions in parallel before deployment.

V. CONCLUSION

A. WORK SUMMARY

This work presents COMIX, a comprehensive Conflict Management framework for Multi-Channel Power Control in O-RAN xApps. The framework addresses the critical challenge of managing conflicting objectives among xApps operating on the Near-RT RIC. It also integrates a CMF component for detecting and resolving conflicts and leverages the NDT to simulate the impact of xApp actions before their deployment in the live network. The framework is demonstrated using two DRL-based xApps, namely the DRM and EE xApps. Realistic multi-channel power control scenarios are used for validation process, where various conflict resolution policies are evaluated. The results showcase the effectiveness of power consumption-aware policies in achieving significant energy savings and balanced optimization objectives, outperforming baseline throughput-centric resolution policies and CMF-free methods.

B. KEY CHALLENGES AND LIMITATIONS

While COMIX demonstrates promising results in resolving xApp conflicts through the integration of CMF and NDT, several limitations should be underlined. First, the centralized nature of the CMF may introduce a processing bottleneck as the number of conflicting xApps increases, potentially impacting scalability and real-time responsiveness. This may be resolved by considering multiple COMIX frameworks across distributed Near-RT RICs and RAN domains, while allowing for inter-domain coordination [33]. Second, the framework assumes an accurate NDT implementation, which may not perfectly replicate the dynamic behavior of a live

RAN under all conditions. This limitation could be mitigated by coupling COMIX with many recent frameworks proposing adaptive and accurate NDT platforms, as in [27], [30]. Third, future work should focus on deploying COMIX in real O-RAN testbeds to assess hardware-specific performance constraints and communication over real and open interfaces. This can be facilitated through collaborations with open testbed initiatives and integration with open-source O-RAN platforms such as those presented in [30], [34], and [35].

C. WORK EXTENSIONS

Future extensions of this work could further enhance the applicability and robustness of the COMIX framework. First, although architectural principles are generalizable, the incorporation of additional xApps introducing direct, indirect and implicit conflicts and targeting diverse objectives, such as latency minimization or fairness, could broaden the scope of the proposed framework. Second, the introduction of adaptive conflict resolution policies, informed by real-time network dynamics and powered by advanced machine learning techniques, could further improve decision-making. Third, COMIX can incorporate resolution policy selection mechanisms that dynamically choose policies based on real-time conflict patterns or historical xApp performance. This would allow the framework to better respond to evolving network conditions and optimize conflict resolution accordingly. Finally, given that conflicts are resolved based on NDT-oriented metrics, elaborating on how NDT could continuously provide accurate representations of the real O-RAN system remains of paramount importance.

ACKNOWLEDGMENT

The authors warmly thank Ilias Paralikas for his help in constructing the simulations.

REFERENCES

- [1] M. Polese, L. Bonati, S. D'Oro, S. Basagni, and T. Melodia, "Understanding O-RAN: Architecture, interfaces, algorithms, security, and research challenges," *IEEE Commun. Surveys Tuts.*, vol. 25, no. 2, pp. 1376–1411, 2nd Quart., 2023.
- [2] A. S. Abdalla, P. S. Upadhyaya, V. K. Shah, and V. Marojevic, "Toward next generation open radio access networks: What O-RAN can and cannot do!" *IEEE Netw.*, vol. 36, no. 6, pp. 206–213, Nov. 2022.
- [3] S. D'Oro, L. Bonati, M. Polese, and T. Melodia, "OrchestRAN: Network automation through orchestrated intelligence in the open RAN," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, May 2022, pp. 270–279.
- [4] A. Giannopoulos, S. Spantideas, N. Kapsalis, P. Gkonis, L. Sarakis, C. Capsalis, M. Vecchio, and P. Trakadas, "Supporting intelligence in disaggregated open radio access networks: Architectural principles, AI/ML workflow, and use cases," *IEEE Access*, vol. 10, pp. 39580–39595, 2022.
- [5] S. Spantideas, A. Giannopoulos, M. A. Cambeiro, O. Trullols-Cruces, E. Atxutegi, and P. Trakadas, "Intelligent mission critical services over beyond 5G networks: Control loop and proactive overload detection," in *Proc. Int. Conf. Smart Appl., Commun. Netw. (SmartNets)*, Jul. 2023, pp. 1–6.
- [6] N. M. Yungaicela-Naula, V. Sharma, and S. Scott-Hayward, "Misconfiguration in O-RAN: Analysis of the impact of AI/ML," *Comput. Netw.*, vol. 247, Jun. 2024, Art. no. 110455.
- [7] C. Adamczyk and A. Kliks, "Conflict mitigation framework and conflict detection in O-RAN near-RT RIC," *IEEE Commun. Mag.*, vol. 61, no. 12, pp. 199–205, Dec. 2023.
- [8] *O-RAN Working Group 3 (Near-Real-Time RAN Intelligent Controller and E2 Interface Workgroup)(O-RAN. WG3. Ricarch-r003-v04.00)*, O-RAN Alliance, Alfter, Germany, 2023.
- [9] S. Marinova and A. Leon-Garcia, "Intelligent O-RAN beyond 5G: Architecture, use cases, challenges, and opportunities," *IEEE Access*, vol. 12, pp. 27088–27114, 2024.
- [10] P. B. D. Prever, S. D'Oro, L. Bonati, M. Polese, M. Tsampazi, H. Lehmann, and T. Melodia, "PACIFISTA: Conflict evaluation and management in open RAN," *IEEE Trans. Mobile Comput.*, early access, May 20, 2025, doi: 10.1109/TMC.2025.3570632.
- [11] J. Armstrong, E. Fallon, and S. Fallon, "Pre-emptive conflict detection architecture for O-RAN service management and orchestration," in *Proc. IEEE Int. Conf. Ind. 4.0, Artif. Intell., Commun. Technol. (IAICT)*, Jul. 2024, pp. 335–340.
- [12] A. Wadud, F. Golpayegani, and N. Afraz, "XApp-level conflict mitigation in O-RAN, a mobility driven energy saving case," 2024, *arXiv:2411.03326*.
- [13] A. Wadud, F. Golpayegani, and N. Afraz, "Conflict management in the near-RT-RIC of open RAN: A game theoretic approach," in *Proc. IEEE Int. Conf. Internet Things (iThings) IEEE Green Comput. & Commun. (GreenCom) IEEE Cyber, Phys. & Social Comput. (CPSCom) IEEE Smart Data (SmartData) IEEE Congr. Cybermatics (Cybermatics)*, Dec. 2023, pp. 479–486.
- [14] A. Zolghadr, J. F. Santos, L. A. DaSilva, and J. Kibilda, "Learning and reconstructing conflicts in O-RAN: A graph neural network approach," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Mar. 2025, pp. 01–06.
- [15] H. Zhang, H. Zhou, and M. Erol-Kantarci, "Team learning-based resource allocation for open radio access network (O-RAN)," in *Proc. IEEE Int. Conf. Commun. (ICC)*, May 2022, pp. 4938–4943.
- [16] T. Karamplias, S. T. Spantideas, A. E. Giannopoulos, P. Gkonis, N. Kapsalis, and P. Trakadas, "Towards closed-loop automation in 5G open RAN: Coupling an open-source simulator with xApps," in *Proc. Joint Eur. Conf. Netw. Commun. 6G Summit (EuCNC/6G Summit)*, Jun. 2022, pp. 232–237.
- [17] K. Skianis, A. Giannopoulos, A. Kalafatis, and P. Trakadas, "Digital twin for automated industrial optimization: Intelligent machine selection via process modelling," in *Proc. IEEE Int. Conf. Ind. 4.0, Artif. Intell., Commun. Technol. (IAICT)*, Jul. 2023, pp. 84–90.
- [18] A. Garcia-Saavedra and X. Costa-Perez, "O-RAN: Disrupting the virtualized RAN ecosystem," *IEEE Commun. Standards Mag.*, vol. 5, no. 4, pp. 96–103, Dec. 2021.
- [19] J. Dai, L. Li, R. Safavinejad, S. Mahboob, H. Chen, V. V. Ratnam, H. Wang, J. Zhang, and L. Liu, "O-RAN-Enabled intelligent network slicing to meet service-level agreement (SLA)," *IEEE Trans. Mobile Comput.*, vol. 24, no. 2, pp. 890–906, Feb. 2025.
- [20] P. Dinh, M. Ghoshal, Y. Han, Y. Feng, D. Koutsonikolas, and J. Widmer, "Demystifying resource allocation policies in operational 5G mmWave networks," *IEEE Trans. Netw.*, vol. 33, no. 2, pp. 745–760, Apr. 2025.
- [21] B. Salehi, U. Demir, D. Roy, S. Pradhan, J. Dy, S. Ioannidis, and K. Chowdhury, "Multiverse at the edge: Interacting real world and digital twins for wireless beamforming," *IEEE/ACM Trans. Netw.*, vol. 32, no. 4, pp. 3092–3110, Aug. 2024.
- [22] A. Giannopoulos, S. Spantideas, N. Nomikos, A. Kalafatis, and P. Trakadas, "Learning to fulfill the user demands in 5G-enabled wireless networks through power allocation: A reinforcement learning approach," in *Proc. 19th Int. Conf. Design Reliable Commun. Netw. (DRCN)*, Apr. 2023, pp. 1–7.
- [23] S. T. Spantideas, A. E. Giannopoulos, N. C. Kapsalis, A. Kalafatis, C. N. Capsalis, and P. Trakadas, "Joint energy-efficient and throughput-sufficient transmissions in 5G cells with deep Q-learning," in *Proc. IEEE Int. Medit. Conf. Commun. Netw. (MeditCom)*, Sep. 2021, pp. 265–270.
- [24] A. Giannopoulos, S. Spantideas, N. Kapsalis, P. Karkazis, and P. Trakadas, "Deep reinforcement learning for energy-efficient multi-channel transmissions in 5G cognitive HetNets: Centralized, decentralized and transfer learning based solutions," *IEEE Access*, vol. 9, pp. 129358–129374, 2021.
- [25] A. Ahmadian, W. Shin, and H. Park, "Long-term throughput maximization in wireless powered communication networks: A multitask DRL approach," *IEEE Internet Things J.*, vol. 11, no. 11, pp. 19616–19631, Jun. 2024.

- [26] S. A. Ullah, M. A. Sohail, H. Jung, M. O. B. Saeed, and S. A. Hassan, "Sum rate maximization in IoT networks with diversity-enhanced energy harvesting: A DRL-guided approach," *IEEE Internet Things J.*, vol. 11, no. 18, pp. 30309–30322, Sep. 2024.
- [27] H. X. Nguyen, K. Sun, D. To, Q.-T. Vien, and T. A. Le, "Digital twin for O-RAN toward 6G," *IEEE Commun. Mag.*, vol. 63, no. 3, pp. 174–181, Mar. 2025.
- [28] T. Jiang, J. Zhang, P. Tang, L. Tian, Y. Zheng, J. Dou, H. Asplund, L. Raschkowski, R. D'Errico, and T. Jämsä, "3GPP standardized 5G channel model for IIoT scenarios: A survey," *IEEE Internet Things J.*, vol. 8, no. 11, pp. 8799–8815, Jun. 2021.
- [29] *Study on Channel Model for Frequencies from 0.5 to 100 GHz*, document 38.901, 3GPP, 2022.
- [30] M. Polese, L. Bonati, S. D'Oro, P. Johari, D. Villa, S. Velumani, R. Gangula, M. Tsampazi, C. P. Robinson, G. Gemmi, A. Lacava, S. Maxenti, H. Cheng, and T. Melodia, "Colosseum: The open RAN digital twin," *IEEE Open J. Commun. Soc.*, vol. 5, pp. 5452–5466, 2024.
- [31] S. S. Malik, S. Sil, K. Ranjan, M. S. Duhan, and A. Singh, "Achieving energy efficiency in open radio access networks (oran) using xapps," *Quantum Sci. Academic Res. Consortium*, vol. 1, no. 1, pp. 1–10, 2024.
- [32] M. Corici, R. Modroui, F. Eichhorn, E. Troudt, and T. Magedanz, "Towards efficient conflict mitigation in the converged 6G open RAN control plane," *Ann. Telecommun.*, vol. 79, nos. 9–10, pp. 1–11, Oct. 2024.
- [33] A. Amrallah, Y. Tsukamoto, T. Murakami, A. Ikami, H. Shinbo, and Y. Amano, "Disaggregated and distributed near-real-time RIC for large scale user-centric RAN," *IEEE Open J. Commun. Soc.*, vol. 6, pp. 1591–1609, 2025.
- [34] P. S. Upadhyaya, N. Tripathi, J. Gaeddert, and J. H. Reed, "Open AI cellular (OAIC): An open source 5G O-RAN testbed for design and testing of AI-based RAN management algorithms," *IEEE Netw.*, vol. 37, no. 5, pp. 7–15, Sep. 2023.
- [35] A. Da Silva, M. R. Chowdhury, A. Sathish, A. Tripathi, S. F. Midkiff, and L. A. Da Silva, "CCI xG testbed: An O-RAN based platform for future wireless network experimentation," *IEEE Commun. Mag.*, vol. 63, no. 2, pp. 62–68, Feb. 2025.



ANASTASIOS E. GIANNOPOULOS (Member, IEEE) received the M.Eng. degree from the National Technical University of Athens (NTUA), in 2018, the Ph.D. degree from the Wireless & Long Distance Communications Laboratory of NTUA, in 2024, and the Diploma degree in electrical and computer engineering from NTUA. He is currently a Research Associate at Four Dot Infinity (FDI), and at the Department of Ports Management & Shipping of the National

& Kapodistrian University of Athens. He has authored more than 60 scientific publications in the fields of Wireless Network Optimization, Maritime Communications, ML and Multi-dimensional Analysis. His research interests include optimization techniques for wireless systems, ML-aided resource allocation, maritime communications, and multi-dimensional data analysis.



SOTIRIOS T. SPANTIDEAS received the Diploma degree in electrical and computer engineering from the University of Patras in 2010, the M.Sc. degree from the Royal Institute of Technology in Stockholm (KTH), in 2013, and the Ph.D. degree from the National Technical University of Athens (NTUA), in 2018. From 2014, he has worked as a Research Associate with NTUA, the National & Kapodistrian University of Athens (Department of Ports Management & Shipping).

He is currently a Research Associate at Four Dot Infinity (FDI). He has published over 40 papers in scientific journals and conferences in the fields of Electromagnetic Compatibility, Optimization Methods for Wireless Networks and ML for Resource Allocation problems. His research interests include electromagnetic compatibility, ML for wireless networks, magnetic cleanliness for space missions and computational electromagnetics.



GEORGE LEVIS received the Diploma degree in applied mathematics and theoretical computer science from the National Technical University of Athens, in 2025, and the M.Eng. degree. Specializing in deep reinforcement learning, his work pioneers adaptive decision-making systems for robotics and real-time strategy optimization. He is also interested in explainable AI (XAI), interpretable models, complexity with transparency, and healthcare diagnostics. His research integrates

rigorous mathematical frameworks with novel AI architectures, including self-interpreting neural networks and ethical reinforcement learning agents designed to audit their own biases. He is currently with the Research and Development Department, Four Dot Infinity, developing scalable AI solutions where trust meets innovation. Recent projects include explainable autonomous systems for enterprise logistics and multi-agent reinforcement learning platforms for dynamic resource allocation.



ALEXANDROS S. KALAFATELIS (Graduate Student Member, IEEE) received the B.Sc. degree (Hons.) in biomedical sciences from the University of East London, and the B.Eng. degree in electrical engineering and the M.Sc. degree in intelligent management of renewable energy systems from NKUA. He is currently pursuing the Ph.D. degree with the Department of Ports Management and Shipping, National and Kapodistrian University of Athens (NKUA). He is also a Senior Research

Scientist with Four Dot Infinity. Part of his research has been conducted in the framework of several European funded research and development projects, focusing on predictive maintenance applications leveraging federated learning, with an emphasis on developing secure aggregation methods tailored to the maritime industry. He also serves as a Committee Member for the Ship Maintenance, Repair, and Safety Special Interest Group, Institute of Marine Engineering, Science and Technology (IMarEST), and was a recipient of the Stanley Gray Fellowship.



PANAGIOTIS TRAKADAS received the Diploma and Master of Engineering (M.Eng.) degrees in electrical and computer engineering and the Ph.D. degree from the National Technical University of Athens. He is currently an Associate Professor with the Department of Ports Management and Shipping, National and Kapodistrian University of Athens, and has been actively involved in many national and EU-funded research projects. He has published more than 170 papers in magazines,

journals, books, and conferences, and is a reviewer in several journals and a TPC in conferences. His research interests include 5G/6G technologies, such as O-RAN, NFV, and NOMA, wireless sensor networks, ML-based optimization for wireless systems, and maritime communications.

...