



dApps: Enabling real-time AI-based Open RAN control

Andrea Lacava ^{a,b}*, Leonardo Bonati ^a, Niloofar Mohamadi ^a, Rajeev Gangula ^a, Florian Kaltenberger ^{a,c}, Pedram Johari ^a, Salvatore D’Oro ^a, Francesca Cuomo ^b, Michele Polese ^a, Tommaso Melodia ^a

^a Institute for the Wireless Internet of Things, Northeastern University, Boston, MA, USA

^b Sapienza University of Rome, 00184 Rome, Italy

^c EURECOM, Sophia-Antipolis, France

ARTICLE INFO

Keywords:
 Open RAN
 dApps
 Real-time control loops
 Radio Resource Management (RRM)
 Spectrum sharing
 Positioning
 Integrated Sensing and Communication (ISAC)

ABSTRACT

Open Radio Access Networks (RANs) leverage disaggregated and programmable RAN functions and open interfaces to enable closed-loop, data-driven radio resource management. This is performed through custom intelligent applications on the RAN Intelligent Controllers (RICs), optimizing RAN policy scheduling, network slicing, user session management, and medium access control, among others. In this context, we have proposed dApps as a key extension of the O-RAN architecture into the real-time and user-plane domains. Deployed directly on RAN nodes, dApps access data otherwise unavailable to RICs due to privacy or timing constraints, enabling the execution of control actions within shorter time intervals. In this paper, we propose for the first time a reference architecture for dApps, defining their life cycle from deployment by the Service Management and Orchestration (SMO) to real-time control loop interactions with the RAN nodes where they are hosted. We introduce a new dApp interface, E3, along with an Application Protocol (AP) that supports structured message exchanges and extensible communication for various service models. By bridging E3 with the existing O-RAN E2 interface, we enable dApps, xApps, and rApps to coexist and coordinate. These applications can then collaborate on complex use cases and employ hierarchical control to resolve shared resource conflicts. Finally, we present and open-source a dApp framework based on OpenAirInterface (OAI). We benchmark its performance in two real-time control use cases, i.e., spectrum sharing and positioning in a 5th generation (5G) Next Generation Node Base (gNB) scenario. Our experimental results show that standardized real-time control loops via dApps are feasible, achieving average control latency below 450 microseconds and allowing optimal use of shared spectral resources.

1. Introduction

The Open Radio Access Network (RAN) architecture promotes open, multi-vendor, software-driven, and programmable cellular networks. Formalized in the O-RAN ALLIANCE technical specifications, this novel network architecture introduces the RAN Intelligent Controller (RIC), a software component that leverages open interfaces to gather RAN data, run inference, and enact control, adapting the network to current demands, conditions, and requirements. O-RAN specifications discuss two versions of the RIC. The Near-RT RIC oversees the RAN operations via xApps that operate in the 10 ms to 1 s timescale, while the Non-RT RIC hosts rApps that operate at timescales higher than 1 s.

xApps and rApps enable a variety of use cases where the RAN can be dynamically configured to optimally handle and adapt to varying network conditions. Wireless systems are characterized by rapidly changing channel conditions, dynamic traffic patterns, user mobility, and periodic or seasonal patterns in network utilization, to name a few. Examples of RIC-enabled control include network slicing [1,2], traffic steering [3,4], beamforming and mobility management [5], advanced sleep modes [6], anomaly detection [7], and spectrum and radio resource allocation [8].

The current O-RAN architecture, though, comes with two key limitations: (i) xApps and rApps are primarily designed to handle control-plane data and operations, thus not considering user-plane data, such as I/Q samples and packets for inference and optimization; and (ii) does

* Corresponding author at: Institute for the Wireless Internet of Things, Northeastern University, Boston, MA, USA.

E-mail addresses: lacava.a@northeastern.edu (A. Lacava), l.bonati@northeastern.edu (L. Bonati), n.mohamadi@northeastern.edu (N. Mohamadi), r.gangula@northeastern.edu (R. Gangula), f.kaltenberger@northeastern.edu (F. Kaltenberger), p.johari@northeastern.edu (P. Johari), s.doro@northeastern.edu (S. D’Oro), francesca.cuomo@uniroma1.it (F. Cuomo), m.polese@northeastern.edu (M. Polese), t.melodia@northeastern.edu (T. Melodia).

not enable control loops at timescales below the 10 ms one enabled by xApps. Indeed, as we will discuss in later sections of this paper, being unable to process user-plane data and perform inference and control below the 10 ms timescale limits the application of O-RAN technologies at the lower levels of the protocol stack, as well as the introduction of novel use cases and applications (e.g., Radio Frequency (RF) fingerprinting [9], one-shot beam steering, anomaly and attack detection [10], spectrum sensing and incumbent detection [11], joint sensing and communications [12,13], to name a few) which will be at the center of 6G [14].

To overcome these limitations, we proposed the concept of dApps [15]. Similarly to xApps and rApps, dApps are software components that can execute as microservices, designed to be co-located with Central Units (CUs) and Distributed Units (DUs), where user-plane data is readily available. Currently, the adoption of dApps is under investigation by the O-RAN next Generation Research Group (nGRG) as a way to bring below 10 ms Artificial Intelligence (AI)/Machine Learning (ML) routines to the RAN. The advantages introduced by dApps are manifold: (i) they can execute real-time operations at the DU/CU directly to achieve real-time control and monitoring of the RAN without the need to involve the RICs; (ii) they have a lightweight lifecycle management, and can be instantiated and deleted seamlessly to provide functionalities as-a-service based on operator requirements; and (iii) they have access to user-plane data that cannot leave the DU/CU due to privacy concerns (e.g., user-related data) or impracticality (e.g., I/Q streams require high bandwidth in the order of Gbps, which would generate congestion over the other O-RAN interfaces [15]).

How to bring AI/ML to O-RAN and push inference toward the real-time domain has received increasing interest in the last few years. Apart from dApps, the literature offers different proposals that include the concepts of real-time RIC and μ Apps [16], zApps [17], tApps [18], to name a few. In this direction, we recently edited an O-RAN nGRG research report [19], with input and collaborations across academia and industry aimed at exploring use cases and applications that would benefit from dApps. These include real-time scheduler reconfiguration, spectrum sensing, compute resource scaling for energy savings, channel equalization, beam management and many others. Despite the well-defined use cases, an architecture for implementing dApps in cellular networks is still lacking. In particular, a comprehensive description of the procedures, interfaces, modules, and their interactions with existing O-RAN components – such as xApps, rApps, RICs, DUs, CUs, and open interfaces (e.g., O1, O2, and E2) – is still missing.

Novelty and Contributions. In this paper, we fill this gap and propose a reference architecture for dApps with the goal of fostering design and prototyping of dApp-based use cases and applications for O-RAN systems. Specifically, we propose an architecture that can be seamlessly integrated with the existing O-RAN architecture with minimal procedural changes. This minimizes the impact that dApps have on O-RAN standardization, while still making their use feasible and practical. We also propose a Lifecycle Management (LCM) for the deployment and management of dApps in the O-RAN architecture. Once deployed, we show that dApps can leverage already existing interfaces and procedures to exchange data with the Near-RT RIC over the E2 interface and perform monitoring and control tasks using a custom E2Service Model (SM) model, i.e., E2SM-DAPP. We introduce a novel E3 interface to allow dApps to interact in real time with DUs and CU, for data and control exchange.

We develop and release as open-source¹ a framework for dApps integrated with the popular OpenAirInterface (OAI) 5G RAN framework. We evaluate the framework on Colosseum [20] and Arena [21] platforms, conducting an extensive performance analysis to benchmark

dApp execution and feasibility. Our results demonstrate that dApps operate within real-time control intervals, efficiently processing both vector data (e.g., I/Q arrays extracted from the RAN) and scalar values in loops taking less than 450 microseconds—well below the 10 ms threshold required for real-time operations.

Additionally, we implement two distinct dApp use cases using the proposed framework for Integrated Sensing and Communication (ISAC): positioning and spectrum sharing. The results reveal that the positioning dApp is able to compute the distance between the User Equipment (UE) and the Next Generation Node Base (gNB) using the UE Uplink (UL) Channel Impulse Response (CIR) collected by the gNB in real-time, and with the advantage of plug-and-play, customizable processing routines. Finally, the spectrum-sharing dApp effectively detects incumbents in the 5th generation (5G) gNB RF context, enabling more efficient spectrum utilization through real-time analysis.

If compared to [15], this paper provides a significantly more detailed and practical architecture for dApps integrated with O-RAN systems. Specifically, we provide details on procedures, messages, protocols, and data structures necessary for dApp operation. Additionally, we provide a prototype implementation of dApps in OAI, demonstrating real-time interactions between dApps and the RAN protocol stack. Ultimately, we provide an extensive performance evaluation focusing on real-time feasibility and practical benefits. Our results show that dApps can perform inference within 10 ms, classify incumbents within 450 μ s, and enable real-time UE positioning.

We believe that the combination of architectural framework, open-source reference library for dApps, and the thorough performance evaluation will inspire further research and development efforts in the real-time control domain.

The remainder of this paper is organized as follows. Section 2 presents the current state of the art for the real-time control loops in O-RAN, highlighting the differences between our work and the real-time RICs proposed in the literature. Section 3 describes the role of the dApps within the O-RAN architecture and the use cases that can benefit from their implementation. Section 4 presents the integration of the dApps in the O-RAN architecture, introducing the E3 logic interface and describing the message exchange between the dApps and the other O-RAN components. Section 5 describes the lifecycle of the dApps, from their onboarding through the Service Management and Orchestration (SMO) to their deployment and in the RAN unit. In Section 6, we provide a reference implementation of the architecture proposed based on OAI and a custom Python framework to measure the real-time control loop communication. We also discuss benchmarking results that profile the performance of the control loop, demonstrating that our real-time control loop can execute in less than 450 μ s. In Section 7, we use our framework to develop two dApp-powered use cases: Spectrum Sharing and Sensing and Positioning. We evaluate their performance and impact on the RAN, and show how dApps can detect incumbents in a 5G network and enable spectrum sharing with them. Section 8 concludes the paper and envisions future developments for dApps. Further clarifications on the dApp deployment process are discussed in Appendix.

2. Related work and comparison with real-time RIC

The need for the introduction of real-time control loops in O-RAN has been advocated in recent works to control lower-layer functionalities of the gNB protocol stack [22,23]. In our previous work [15], we introduced the concept of dApps to process data locally available at the gNB, and the need for sub-10 ms control loops. However, our prior work primarily focused on advocating the benefits of such tighter control loops for generic use cases. In this work, instead, we propose a complete reference architecture and design for dApps, including interfaces for interoperability of dApps with the rest of the O-RAN components. Then, we showcase the capabilities and potential of dApps for use cases of interest to Open RAN networks.

¹ The framework is available at <https://github.com/wineslab/dApp-framework>. A tutorial on how to use the framework is available at <https://openrangym.com/tutorials/dapps-oai>.

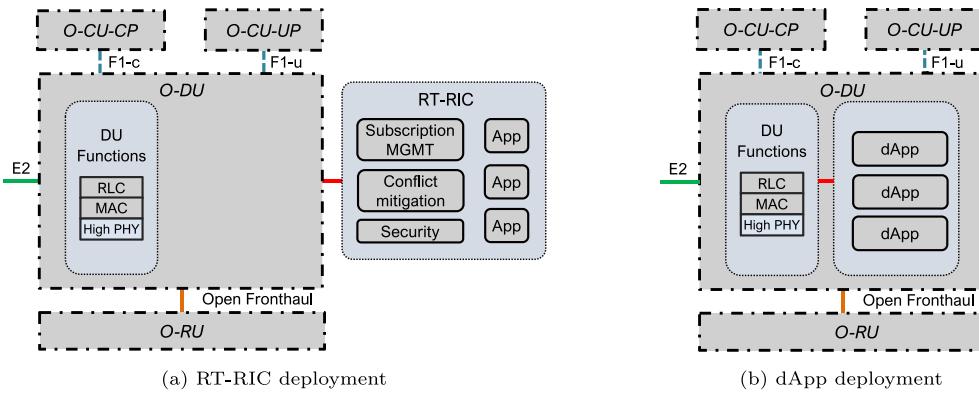


Fig. 1. Comparison between RT RIC and dApp architectures.
Source: From [19].

The need for real-time control loops and inference in O-RAN has also been highlighted by the broader community. For example, the concept of Real-time RICs and its applications have been proposed in [16–18]. Specifically, [17] proposes to embed CUs and DUs inside the real-time RIC, as well as to further disaggregate their functionalities, where each one (e.g., Medium Access Control (MAC) scheduler, channel estimators, and beam shapers) is carried out by a dedicated application, namely zApp. [18] proposes TinyRIC, which runs close to the CU/DU, and so-called tApps hosted therein. The latter aid the CU/DU by performing functionalities spanning from the management of non-3rd Generation Partnership Project (3GPP) interfaces (e.g., O1 and E2), to energy management, to providing primitives for data collection. Similarly, [16] proposed EdgeRIC and its applications, μ Apps, that interface with the RAN to make decisions at Transmission Time Interval (TTI) level. Although these approaches propose a viable way to introduce real-time control in the RAN, they do so by replicating functionalities typical of non- and Near-RT RICs, e.g., intelligence orchestration, in a new RT RIC platform, which inherently increases the network complexity, as well as resource utilization and energy consumption. Our approach, instead, leverages existing RIC and well-established routines to deploy and operate applications on the RAN elements, while avoiding the additional complexity and resource utilization that the additional RT controllers would require. Indeed, the main difference between the above-mentioned solutions and dApps lies in the way intelligent applications are hosted and executed. A high-level architectural comparison of the two approaches is shown in Fig. 1, which highlights the placement of intelligent applications in the each of them. While the above-mentioned applications require a dedicated RT RIC entity, as well as additional layers of network abstractions, dApps execute at the CU/DU directly. This minimizes the impact on the E2 nodes and takes advantage of interfaces and functionalities already defined in the O-RAN specifications.

Finally, [24] proposes Janus, a monitoring and control framework that can load and execute custom real-time “codelets” on the CU and DU. Even though this work—possibly the most similar to ours—extends O-RAN to real-time control, Janus’s codelets require direct access to the CU/DU protocol stack rather than interfacing with it, as dApps instead do. In some cases, this can be a limitation as vendors might need to host third-party codelets rather than exposing parameters via interfaces to regulate access to CU/DU functionalities. Moreover, this approach also requires the definition and deployment of a dedicated controller on the RIC, instead of relying on standardized components for the management of real-time RAN functions, as we instead do in our dApp architecture.

Table 1 summarizes the key differences between dApps and other real-time control solutions and the advantages of our proposal.

3. The role of dApps in the hierarchical O-RAN control architecture

Achieving programmability of the cellular protocol stack via closed-loop control is a fundamental innovation introduced by the O-RAN architecture and, more in general, by the Open RAN vision [22]. With closed-loop control, data-driven logic units executed in xApp and rApps access RAN telemetry and process it to identify RAN configurations that satisfy specific intents and objectives of the operator. The configuration is then applied to the RAN nodes either in form of control (e.g., update a certain RAN parameter) or as a policy, allowing the system to reach the desired state.

Fig. 2 reports a simplified view of the O-RAN logical architecture and control loops. This includes the RAN, with the Open CUs-Control Plane (CP) and User Plane (UP), the DU, and the Radio Unit (RU).² The F1 interface connects CU and DU, while the fronthaul interface bridges the DU and the RU. The E1 interface is in between the CU-CP and UP network functions. The O-RAN architecture leverages the Non-RT RIC and the SMO framework for orchestration and policy definition, through the O1 interface to all RAN nodes and the A1 interface to the Near-RT RICs. The latter performs radio resource management with direct control of the RAN nodes through the E2 interface. These elements are deployed over a set of physical and virtualized resources on a heterogeneous infrastructure (i.e., the O-Cloud). The O2 interface from the SMO interacts with the O-Cloud to configure the compute layer.

The current O-RAN architecture features two control loops, at the non-real-time and the near-real-time scales. The first, exercised by the Non-RT RIC, identifies policies that apply to the network considering Key Performance Measurements (KPMs) reported on a scale of thousands of end devices. It operates at a scale of 1 s or more. The second loop runs between 10 ms and 1 s, and is exercised by the Near-RT RIC. It can perform control or deploy policies, potentially based on the ones defined by non-real-time loops. Due to the tighter requirements on the completion of the control loop, the Near-RT RIC usually targets tens of base stations with hundreds of end devices [25].

These loops, however, share some limitations, as discussed in the right portion of Fig. 2, such as the lack of control loops faster than 10 ms, and the lack of interaction and programmability on the user-plane data. Real-time interactions open up many fine-grained and new customizable and programmable inference and control loop capabilities, including beam management, scheduling profile selection, packet tagging, dynamic spectrum access, and Quality of Service (QoS) enforcement, among others. Further, access to user-plane data would

² For the purposes of this paper, we consistently refer to the Open RAN units while omitting the prefix “O-”.

Table 1

Comparison of dApps with RT RIC-based solutions and Janus.

Feature	dApps (this work)	RT RIC-based Solutions (zApps [17], tApps [18], μ Apps [16])	Janus [24]
Connectivity to RAN stack	Disaggregated	Disaggregated	Integrated
Timescale	< 1 ms	from < 1 ms (μ App, tApp) to < 10 ms (zApp)	< 1 ms
Deployment Location	Directly on CU/DU	Require dedicated RT RIC platform	Directly modifies CU/DU protocol stack
Additional Network Complexity	Minimal, leverages existing RIC components	High, add new RIC layers	High, modifies CU/DU stack
Resource Utilization	Optimized, avoids additional overhead	Increased due to extra control layers	High, requires additional controllers
Integration with O-RAN	Use standardized interfaces and components	Require additional network abstractions	Limited, modifies protocol stack instead of interfacing with it
Scalability	High, built on existing RAN elements	Limited, require new infrastructure	Limited, protocol stack modifications reduce compatibility
Modularity	High, they interface with CU/DU without modifying stack	Medium, dependent on RT RIC architecture	Low, tightly coupled with CU/DU stack
Flexibility	Adapt to different CU/DU implementations	Constrained by RT RIC architecture	Low, as it relies on direct modifications to protocol stack
Use of Standard O-RAN APIs	Yes, integrates with O-RAN standard components	Partially, introduce custom abstractions	No, modifies stack implementation directly

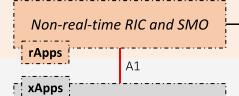
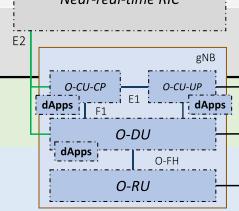
Control and learning objective	Input data	Timescale	Architecture	Challenges and limitations	Supported by O-RAN
Policies, models, slicing	Infrastructure KPMs	Non-real-time > 1 s		<ul style="list-style-type: none"> No real-time interactions with the RAN protocol stack (limits control use cases) No access to the user plane of the RAN (security, privacy, data rates) 	
User Session Management e.g., load balancing, handover	CU KPMs e.g., number of sessions, PDCP traffic	Near-real-time 10-1000 ms			
Medium Access Management e.g., scheduling policy, RAN slicing	MAC KPMs e.g., PRB utilization, buffering	Near-real-time 10-1000 ms			
Radio Management e.g., scheduling, beamforming	MAC/PHY KPMs e.g., PRB utilization, channel estimation	Real-time < 10 ms		<ul style="list-style-type: none"> dApp architecture and integration with stack (this paper) Efficient and fast AI/ML models for closed-loop inference 	dApp Extension
Device DL/UL Management e.g., modulation	I/Q samples	Real-time < 1 ms			

Fig. 2. O-RAN control loops, limitations, and extensions to the real-time and user-plane domains.
Source: Adapted from [22].

enable direct interaction with waveforms and Packet Data Units (PDUs) at different layer of the stack, as well as inference and control based on the rich information that these elements carry. There is a significant body of research demonstrating the benefit of inference and classification based on raw I/Q samples for anomaly detection, spectrum sensing, fingerprinting, and beam management, among others [9, 11, 26–28]. The same applies to user-plane units at higher layers, i.e., transport blocks, Radio Link Control (RLC) and Packet Data Convergence Protocol (PDCP) packets, e.g., for traffic classification and slice identification [1, 29].

However, O-RAN xApp/rApp-based control loops are not suitable for real-time control. For example, transferring I/Q samples or, in general, user data out of the RAN, is generally unfeasible due to security, privacy, and timing/bandwidth constraints, as also mentioned in Section 1. We discussed this in [15, 19], estimating that transferring I/Q samples through a rate-limited E2 interface would take seconds, which is incompatible with real-time interactions, and can limit the capabilities of the E2SM for the Lower Layer Control (LLC) [30]. Bringing programmability and observability to RAN nodes is, thus, paramount to address both limitations. The bottom part of Fig. 2 illustrates the

role that dApps can play in future iterations of the O-RAN architecture. dApps are lightweight and plug-and-play services that provide secure and real-time access to the stack without introducing the overhead that an additional real-time RIC would bring to network-oriented RAN nodes.

Use cases for dApps and real-time control loops. In [19], recently published as contributed research report by the O-RAN ALLIANCE, we reviewed how dApps can complete the hierarchical structure of control loops within the O-RAN architecture. The use cases can be classified into four main groups, described next.

- **Direct Processing of Waveforms and PDUs.** Access to I/Q samples can be used for physical layer security, e.g., through anomaly detection to detect spoofing of base stations or legitimate users [10], and with RF fingerprinting for secure authentication. Similarly, information carried by I/Q samples on scheduled or unscheduled symbols can be used to detect spectrum holes or incumbents. This enables spectrum sharing but also remote interference detection, including from other base stations in the operator network. In both cases, information processed at the

dApp level can be shared with the stack or xApps to coordinate responses across multiple RAN nodes, paving the way for the RAN to function as an ISAC system.

- **Real-Time Scheduling and Beam Management.** This includes real-time scheduling reconfiguration to embed ad hoc policies and Service Level Agreements (SLAs) for specific traffic patterns and profiles (e.g., new slices), scheduling acceleration for efficient massive Multiple Input, Multiple Output (MIMO) and multi-cell coordination [31], and real-time sharing coordination with arbitration across different entities [32]. Similarly, beam management procedures can be enhanced and aided by direct access and manipulation of reference and synchronization signals in the DU [33]. In this context, dApps can extract real-time measurements, perform data-driven inference at near-real-time periodicity, and send the resulting data to an xApp, which uses it to control beamforming via the existing E2SM LLC [30].
- **RAN Nodes and Fronthaul Configuration.** dApps can be used to dynamically reconfigure and tune the RAN nodes underlying compute and the fronthaul interface between the DU and RU. For example, dApps can bridge real-time RAN telemetry from the stack with the compute node configuration to dynamically modify CPU pinning or CPU energy states for energy-efficiency optimization. Similarly, the real-time telemetry can dynamically affect how the fronthaul interface is configured, e.g., adapting the compression level to save resources when users have good Signal-to-Noise-Ratio (SNR) and providing uncompressed streaming for users at the edge of the cell.
- **Augmented Sensing and Channel Estimation.** dApps can enable dynamic Channel State Information (CSI) compression in coordination with UEs [34], as well as custom AI/ML models for channel estimation that can be tailored to specific user conditions and requirements. dApps can also be used to process reference and synchronization signals for use cases not traditionally considered within RAN nodes, e.g., sensing and positioning [35].

4. dApp service-based architecture and integration with RAN nodes

Based on the use cases and related requirements discussed in Section 3, in the following paragraphs we discuss how the O-RAN architecture can be extended to support dApps as custom, pluggable RAN microservices. We focus on three key elements: (i) the integration of dApps with RAN nodes, extended via a service-based E3 interface; (ii) the proposed E3 procedures that enable seamless message exchange between the RAN and the dApps; and (iii) the interactions between dApps and xApps, facilitated through our E3 interface and the O-RAN E2 interface to the Near-RT RIC. Moreover, the E3 interface requires two main components, the E3 API and E3 Agent. The E3 Application Programming Interface (API) is the software implementation of the E3 logical interface, providing the necessary functions for interaction. Meanwhile, the E3 Agent is a entity that implements the E3 interface using the E3 API to facilitate and manage communication. Further discussion on the dApps lifecycle and the role of the O1 interface is provided in Section 5. Finally, Section 6 describes a prototype of E3 Agent based on this architecture and the open-source OAI stack as a reference.

4.1. dApps data, telemetry, and control flows

In Fig. 3, dApps are shown as standalone pluggable microservices with a southbound E3 interface to the RAN node (either DU and CU) that is also bridged with the RAN E2 termination for cooperation with the xApps. The southbound interface, which we design as a service-based API, discussed below, can be used to expose a variety of data and telemetry from the RAN. Data can be accessed through streams, by polling, and with different periodicity and granularity.

The DU RAN function can provide the dApp with access to a variety of information elements in the data plane. I/Q samples can be accessed pre- or post-equalization, depending on the dApp requirements, at different levels of granularity and periodicity. Such data can be collected from various channels such as Physical Uplink Control Channel (PUCCH) or Physical Uplink Shared Channel (PUSCH), and can include all subcarriers or specific subsets. The Buffer Status Reports (BSRs), exchanged by MAC and RLC instances, provide real-time information about buffer conditions and can be streamed at different intervals (e.g., per slot or subframe) or accessed on-demand. Channel Quality Information (CQI) values are generated based on O-DU processing and can be streamed immediately or polled on demand to support dynamic resource allocation. CSI can be similarly streamed in real-time or accessed on-demand for adaptive modulation and coding decisions. Finally, uplink Sounding Reference Signal (SRS), transmitted by the UE and processed by the O-DU, can be made available either continuously or on request, providing insights for uplink resource configuration.

The DU and CU-UP RAN functions can also expose transport blocks or PDUs at the MAC, RLC, PDCP, and Service Data Adaptation Protocol (SDAP) layers according to configurable policies, which may involve streaming a subset of packets at certain intervals or polling them on demand. In the control plane, the dApps can access MAC Downlink Control Information (DCI) and Uplink Control Information (UCI) in the DU. These indications can be streamed when scheduling decisions are made, or retrieved on request to inform dApps of the current state of resource allocation. Additional information is represented by compute telemetry, i.e., information on CPU, RAM, and accelerator utilization that can be streamed at high frequency (e.g., hundreds of microseconds) to monitor system performance, or polled on demand for less frequent assessments, and fronthaul configuration, to be polled as needed to ensure that the network setup aligns with the requirements of various dApps. The CU-CP can expose information on slice and bearers configuration, user sessions, among others. Finally, KPMs from CU-CP, CU-UP, and DU, which are available over E2 [36–38], can also be provided to the dApps.

In the opposite direction, dApps can provide the RAN functions with control actions, which must be applied within a short interval (e.g., 0.5 ms) from the time a configuration is selected in the dApp. The areas identified for control in the RAN nodes include elements to enable the use cases in Section 3. The following examples illustrate the capabilities of dApps but are not exhaustive and can be extended as needed. Within the DU, dApps can configure various aspects of the MAC scheduler, including prioritization parameters, scheduling policies, and acceleration configurations. They can also manage beamforming weights, codebooks, Synchronization Signal (SS) blocks and bursts, and Channel State Information - Reference Signal (CSI-RS), as well as update mapping tables for Signal to Interference plus Noise Ratio (SINR)-to-Modulation and Coding Scheme (MCS), apply Physical Resource Block (PRB) masking or nulling, and control DCIs and UCIs. At the CU-UP and CU-CP levels, dApps enable interactions with slice and cell configurations, allowing dynamic adjustments to network slices and modifications to parameters such as transmit power and carrier frequency at the cell level.

4.2. Service-based dApps E3 API endpoint

As depicted in Fig. 3, dApps are deployed as services coexisting with RAN nodes. An E3 API endpoint (or E3 agent) bridges the RAN functions with the dApps. This agent provides a standardized API layer that supports functionalities such as dApp setup, configuration, data streaming or polling, and control. By encapsulating RAN function capabilities, the E3 agent ensures dApp independence from specific RAN implementations. In addition to interfacing with dApps, the E3 agent coordinates with the E2 termination within the RAN node, enabling seamless communication between dApps and xApps. This facilitates the synchronization of control and adaptation processes across the Near-RT

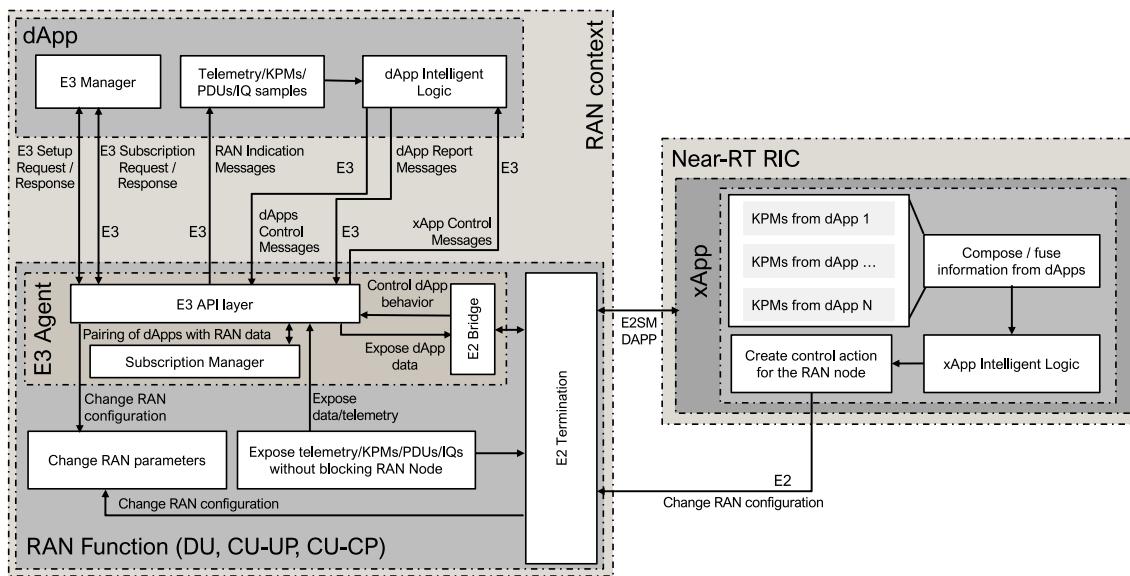


Fig. 3. O-RAN Near-RT RIC integration with dApps and a generic RAN node.

RIC and dApps at varying timescales, ensuring efficient and dynamic network management.

The architecture supports multiple dApps interacting with the same RAN function, provided there are no conflicts. Conflict resolution is addressed holistically within the network as part of the dApp lifecycle (Section 5). Moreover, the E3 agent, in collaboration with the O-Cloud and SMO, monitors available compute resources on the O-Cloud physical node. It ensures that RAN workloads are prioritized to meet SLAs while maintaining resource allocation for dApps.

To streamline dApp operations, the E3 agent consolidates access to RAN data, such as telemetry or data-plane units (e.g., I/Q samples). A centralized subscription manager within the agent handles the dispatching of the management messages between the dApps and the RAN, i.e., the initial pairing, the subscription request and response, and the eventual de-registration of the dApp. This unified access mechanism eliminates redundant protocol stack calls and ensures efficient coordination by allowing dApps to register callbacks for querying control and data functionalities. The design simplifies interactions and maintains low latency, supporting seamless integration with RAN operations without introducing overhead. With The E3 APIs enable dApps to extract necessary RAN data and deliver computed control actions or inferences without disrupting RAN operations. The API design ensures concurrent functionality across multiple dApps while maintaining operational integrity. Given the co-location of CUs/DUs and dApps on the same host, the E3 API leverages Inter-Process Communication (IPC) mechanisms – such as shared memory, logical First In First Out (FIFO) queues, function calls, or system-domain sockets – instead of network PDUs or sockets. This approach is demonstrated in our prototype implementation (Section 6).

4.3. dApp and RAN interactions over E3

An example of the interaction flow between a dApp and the RAN E3 agent is shown in Fig. 4. Similar to the O-RAN approach for the E2 interface, we introduce a procedural API, called E3 Application Protocol (AP) (E3AP) consisting of two distinct logical steps to establish the association between a single dApp and a RAN node with its data.

When deployed for the first time, the dApp triggers the *E3 Setup* procedure to perform the authentication and pairing with the RAN node. Then, the *E3 Subscription* procedure enables the dApp to query and subscribe to the data and control functionalities supported by the RAN function, as defined through the E3 SMs APIs.

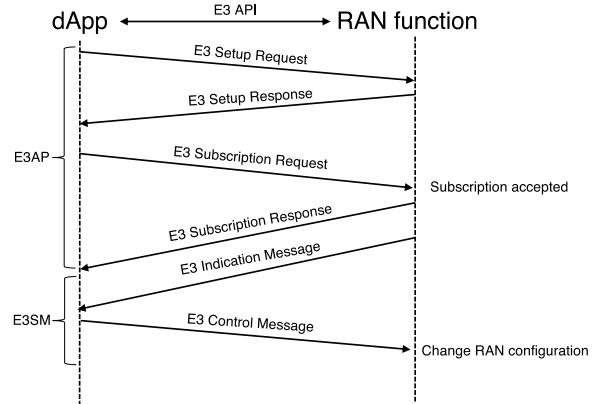


Fig. 4. Message exchanges for a real time control loop between a dApp and the RAN node through the E3 interface.

In order to enable the exposure of data or telemetry and the control of specific RAN functionalities, the E3SM are introduced as modular abstractions implemented through two dedicated APIs. The first API is for the *E3 Indication Message*, to let the RAN function expose data to the dApp. Access to this API is coordinated by the subscription manager in the E3 agent, as discussed above. Similarly, the control generated by the dApp is delivered to the RAN function for processing in the stack using the *E3 Control Message API*. In Section 7.1, we will provide a practical use case where control messages from the dApp are used to avoid interfering signals via PRB nulling.

4.4. E2SM-DAPP: a service model for managing interactions between dApps and xApps

Apart from control actions that actively reconfigure RAN parameters based on input data, dApps can also perform forecasting and prediction tasks. The outcome of these tasks is important because it builds contextual awareness on RAN operations and conditions and, for this reason, it might be useful to make it available to higher layers of the protocol stack, e.g., xApps executing at the Near-RT RIC. This can be achieved by forwarding dApp inference outcomes to xApps in the form of enrichment information transmitted over the RAN E2 interface (e.g., similarly to how the A1 interface is used to forward

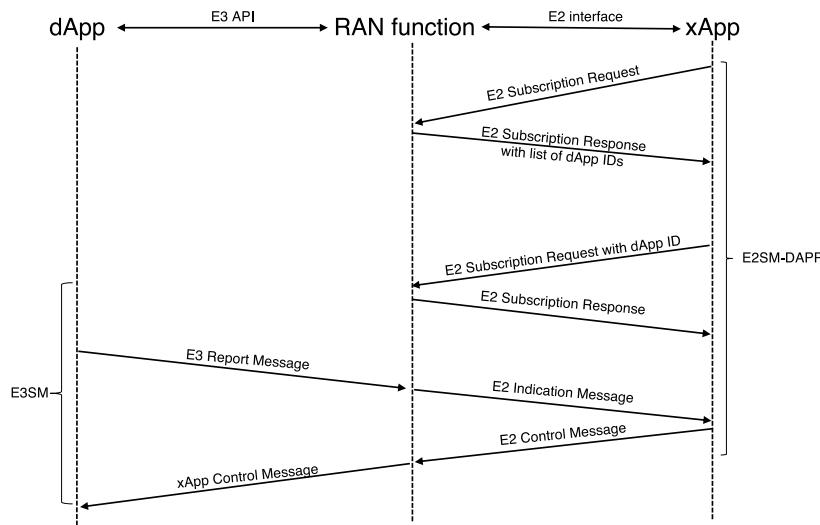


Fig. 5. Interactions between a dApp, the E2 RAN node, and an xApp through the E3 interface and the E2SM-DAPP custom Service Model.

useful information from the Non-RT RIC to the Near-RT RIC). From the E2 perspective, the data exposed by dApps is managed through a novel, custom E2 SM [38], which we call E2SM-DAPP, identified by a unique RAN function ID during the E2 Setup Request.

Fig. 5 summarizes the interactions between the dApps, the E2 RAN node, and the xApps. To initiate these procedures, the xApps onboarded on the Near-RT RIC can send an E2 Subscription Request to the RAN node, which responds via an E2 Subscription Response with a message that includes the IDs of the dApps currently executing on the node. These discovery procedures are executed continuously throughout the xApp lifecycle, enabling the discovery of new dApps as they become available.

xApps that are willing to receive inference data (e.g., spectrum sensing outcomes, real-time channel estimations, scheduling information) from the dApp can send an additional E2 Subscription Request with the same dApp ID and a payload specifying the data they want to collect from the dApp.

Once this pairing is established, the dApp periodically collects real-time data from the RAN using the real-time control loop procedures described in Section 4.3, performs near-real-time inference (e.g., spectrum classification), and transmits these inferences to the xApp using the *E3 Report API*, which connects the dApp to xApps leveraging an E2 Indication Message via the RAN E2 interface.

It is also worth mentioning that xApps can configure how dApps take decisions and perform inference by specifying high-level parameters and policies. These policies influence how the dApp operates and can be used by xApps to influence decisions. For example, in the case of a dApp performing scheduling operations, an xApp can specify which scheduling profile to use (e.g., round-robin, proportional fairness) as well as defining priority weights for scheduling UEs and prioritizing certain types of traffic. To enable this, the xApp sends an E2 Control Message to the RAN node, which is then relayed over the E3 interface using the *xApp Control Message* to the target dApp.

5. dApp lifecycle and interaction with Near-RT RIC, Non-RT RIC, and SMO

In this section, we discuss how the LCM of dApps can be integrated into LCM practices defined within the O-RAN architecture and processes. We follow and extend the specifications outlined in [39, 40], which describe the standard procedures for the LCM of O-RAN applications, cloudification, and orchestration use cases for deploying virtualized O-RAN solutions. The LCM process follows a structured 7-stage model for the Software Development Lifecycle (SDLC) to ensure

that applications are developed, onboarded, and operated in a standardized way. It begins with *Need* to identify requirements, followed by *Ideation* to brainstorm potential solutions, *Analysis* to assess feasibility, *Develop* to build the application, *Deliver* to deploy the solution, *Validate* to test for quality, and finally *Operate* to maintain the application in production.

This structured LCM approach is divided into three main phases: (i) Development, in which applications are designed and built; (ii) Onboarding, in which applications are integrated into the target environment; and (iii) Operations, in which applications are managed and maintained in production [39]. Fig. 6 provides a high-level representation of dApp LCM, with particular focus on the onboarding and operations phases. While the development process is common to all types of applications (i.e., xApps, rApps, and dApps), Fig. 6 focuses specifically on unique aspects related to onboarding and operation of dApps. While the SMO is responsible for deploying the dApp, once deployed, the dApp can operate independently or be managed by other entities, such as an xApp, through the E2SM-DAPP. In the following, we provide a detailed overview of how these phases can be extended and adapted to dApp LCM.

5.1. Development

The development phase of the dApp LCM aligns with the Solution App Lifecycle and Solution AppPackage Lifecycle outlined in the LCM model in [39]. The Solution App Lifecycle covers the end-to-end process of designing, developing, deploying, and managing the dApp, ensuring its smooth operation across its entire lifecycle. The Solution AppPackage Lifecycle focuses specifically on the packaging, onboarding, and management of the application package, including the creation of deployment and management descriptors to standardize the process. The process begins with defining the use case, requirements, and features, followed by the integration of components such as intelligent logic, telemetry, and Key Performance Indicators (KPIs) of the dApp. The dApp is then containerized using tools like Docker for efficient deployment in cloud environments such as Kubernetes or OpenShift. Once containerized, the Solution AppPackage is created to ensure proper onboarding, containing key descriptors like the DeploymentDescriptor and ManagementDescriptor, which standardize deployment and management processes for seamless orchestration.

To ensure secure onboarding, digital signatures and checksums are added to verify the package source, preventing unauthorized modifications and ensuring the application integrity as it transitions through different environments.

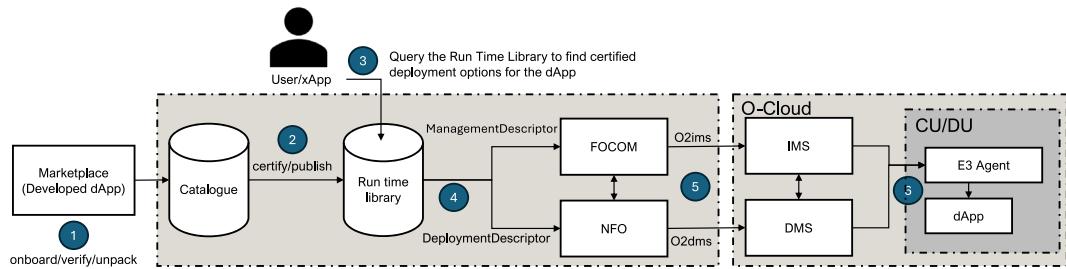


Fig. 6. dApps LCM diagram.

5.2. Onboarding

The onboarding phase begins with the transfer of the dApp App-Package, generated during the development phase, from the marketplace where it was stored. This package undergoes verification and validation to ensure it is authorized and secure. Once validated, the AppPackage is unpacked, and its components are stored in a catalog within the SMO (Step 1, Fig. 6). Following this, each recommended configuration of the dApp is certified and published to a runtime library (Step 2).

The initiation of the onboarding process can typically be triggered either by the SMO or directly by an operator, who queries the runtime library to locate certified deployment options for the dApp. In certain scenarios, this decision may also be delegated to xApps or rApps if they are equipped with the necessary intelligence and reliability to orchestrate and control dApp execution. Indeed, it can be advantageous for an xApp to be able to initiate the deployment of dApps. This allows the xApp to use the available dApp for tasks requiring real-time control. The process flow involves the xApp querying the available dApps in the Run-Time Library through the O1 interface, and requesting the dApp deployment via the O1 interface (Step 3).

However, in Fig. 6, the process for xApp-initiated onboarding is not explicitly shown as this procedure requires ad-hoc service models, procedures and functionalities that are not yet available in the O-RAN specifications, and we hope to cover in future research.

5.3. Deployment

In contrast to the deployment of xApps and rApps, hosted in the Near-RT and Non-RT RICs, dApps require a different deployment procedure, as they are directly hosted in the DU/CU. dApps are deployed as Cloud-Native Functions (CNF), adhering to the standard Network Functions (NF) deployment procedures set out by the O-RAN ALLIANCE [40].

Using the SMO framework and the O-Cloud, the dApp is managed through the O2 interface, which facilitates communication between the SMO and the Deployment Management Services (DMS) of the CU/DU. This direct deployment model allows the dApp to utilize the real-time processing capabilities of the CU/DU while also leveraging the flexibility and scalability of the O-Cloud infrastructure. However, this deployment diverges significantly from that of xApps and rApps, as it is executed outside of the RIC environments, focusing instead on efficient resource allocation and real-time operations directly at the CU/DU level. This approach integrates dApps seamlessly into the CU/DU environment while adhering to the cloud-native principles and ensuring optimal network function performance. The deployment steps of dApps in the LCM workflow are shown in Steps 4, 5, and 6 in Fig. 6.

The information flow for deploying the dApp follows a structured sequence, beginning with a service request initiated by the Network Function Install Project Manager to the SMO for deploying a new dApp instance on the CU/DU. The SMO processes this request by decomposing it, identifying the required dApps and their deployment

order, and determining the deployment parameters based on policies or explicit input. The SMO retrieves the CloudNativeDescriptor for the dApp from the runtime library and directs the O-Cloud Deployment Management Service (DMS) to create the dApp deployment.

The DMS allocates the necessary compute, storage, and network resources on the CU/DU, deploys the dApp container(s), and notifies the SMO upon successful instantiation, providing a Deployment ID. The SMO then updates its inventory with the deployment status, and the dApp instance reads its configuration and begins operation. Following deployment, the SMO continuously monitors the dApp's health, performance, and connectivity, managing its lifecycle, scaling, and updates as required. Finally, the SMO informs the Network Function Install Project Manager of the overall success or failure of the deployment request. For a detailed breakdown of these steps, please refer to Table A.4 in Appendix.

Finally, we also consider the case where xApps can also request the SMO to deploy a new dApp, and connect to it, whenever there is a need to perform real-time control. Deployment is performed following the procedures described above.

6. A reference open-source dApp framework

In this section, we present a reference implementation and prototype of the proposed dApp architecture (Section 4), based on OAI and a Python framework. We also present an exhaustive benchmarking analysis with results that confirm RAN control loop latencies well below 1 ms (i.e., below the target threshold of 10 ms). The reference implementation supports real-time control loops. The open-source library can thus be used for real-time Open RAN control prototypes. The use-cases developed as a proof of concept for our framework and their evaluation are discussed in Section 7.

6.1. dApp framework

We define a reference dApp framework written in Python, implementing the E3AP as described in previous sections. Fig. 7 illustrates the Unified Modeling Language (UML) diagram of our framework. We design three main classes composing the dApp framework to implement diverse functions: the DApp, the E3Interface, and the E3Connector classes.

The DApp Class – First, we define the DApp abstract class for the management of the E3SM operations. For each use case, this class will be extended to include functionalities, parameters and operations specific to the use case. This class wraps all the E3SM functionalities into a single entity to simplify and streamline the generation of new children classes and enable the extension to new use cases. The SpectrumSharingDApp and the PositioningDApp reported in Fig. 7 are two examples of children classes implementing the operations reported in the spectrum sharing use case of Section 7.1 and the positioning use case of Section 7.2. The abstract DApp class uses the E3Interface class for the communication between with the RAN node.

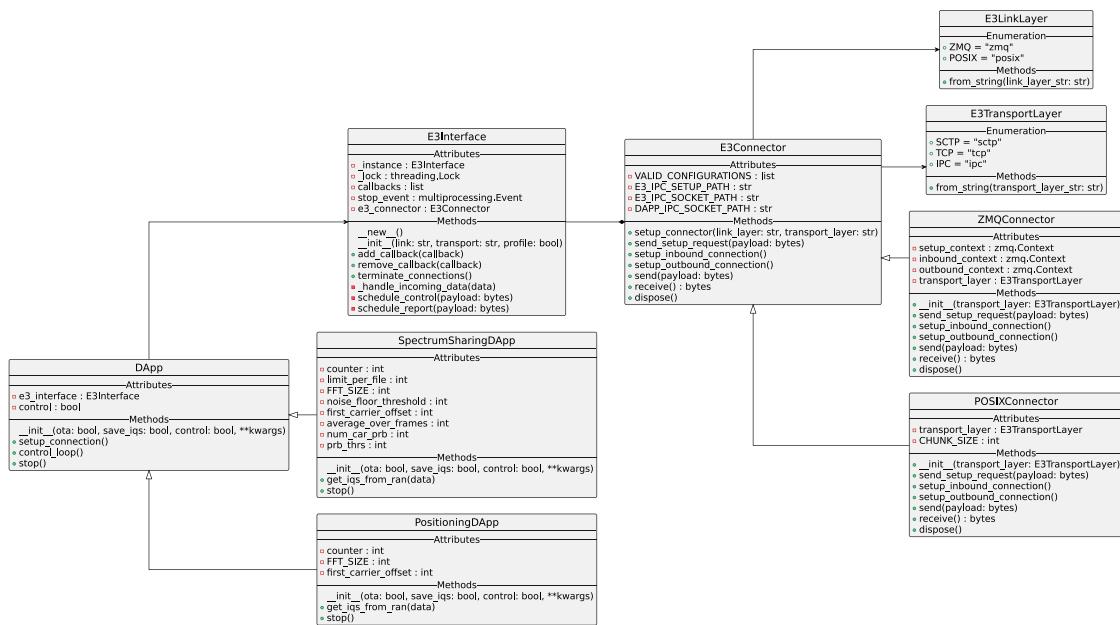


Fig. 7. Unified Modeling Language (UML) diagram of the dApp framework with its components and the two dApps presented in this work used as example.

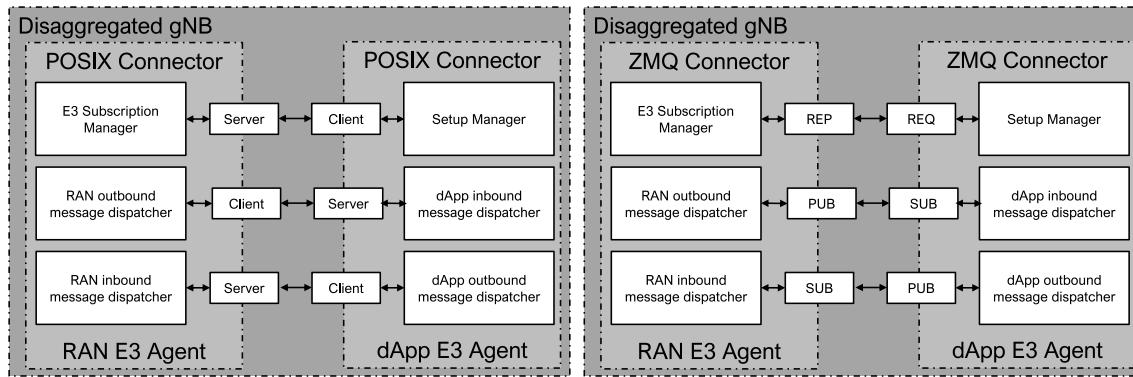


Fig. 8. Design of the E3Connector sockets.

The abstract DApp class leverages the E3Interface class to facilitate communication with the RAN node. The E3AP operations and interactions between the RAN and dApps utilize a publish-and-subscribe mechanism, where dApps register callbacks to access RAN data through specific RAN function IDs. These data are published via the E3Interface, which is a private variable of the abstract class, as illustrated in Fig. 7. The `setup_connection()` class method initializes the E3 interface and the dApp callback, while the `control_loop()` method is designed to implement the data polling and periodically trigger the core logic of the dApp.

The E3Interface Class – The E3Interface class provides standardized APIs for accessing the E3 interface. It manages the E3AP and facilitates the dispatch of data from the RAN node to dApps, enabling both analysis and control operations via function calls. The class implements a singleton pattern, ensuring that all dApps deployed on the same RAN node share the same E3 interface.

The DApp class interacts with the E3Interface by invoking methods such as `add_callback()` and `remove_callback()` to manage subscriptions to RAN data. Additionally, it uses `schedule_control()` to schedule the delivery of E3 Control Actions to the RAN and `schedule_report()` to send E3 Report Messages to xApps. Once a dApp has initiated its control loop, the E3Interface exposes RAN data to the dApps through the `handle_incoming_data()` method.

The E3Connector Class – Ensuring fast and reliable connectivity between dApps, the RAN and xApps is crucial to enable dApp operations and real-time monitoring and control of the network. For this reason, to understand what are the trade-offs and the best protocols to achieve real-time control loops, we have implemented and compared several connectivity solutions and protocols. These efforts have been synthesized in the E3Connector class, which is an abstract class that the E3Interface uses to create the socket connections between the dApps and the RAN node. The E3Connector exposes a set of methods that are used to perform the initial pairing between the dApp and the RAN node, receive the messages from the RAN unit and send the messages to it. As shown in Fig. 8, which represents the different logical pairings between the configurations of the connectors according to the protocols used, the E3Connector creates and manages three different sockets: the first one is for the initial E3 Setup Request–Response and the E3 Subscription Request–Response exchanges; the second one for receiving the inbound messages from the RAN unit, i.e., the E3 Indication Message and the xApp Control Message, and the last one for delivering to the RAN node the outbound messages of the dApps, i.e., the dApp Control Messages and the dApp Report Messages.

In this work, we have created two mutual children classes of the E3Connector, the ZMQConnector and the POSIXConnector, each one implementing a different data link layer and several transport layers. The first one is the POSIXConnector, which implements the message

exchange using the classic POSIX system functions based on the Linux kernel socket implementation. Such functions expose low-level APIs for IPC like shared memory, pipes, message queues, and sockets and they require detailed handling for synchronization, data serialization, and transport, focusing on one system or tightly coupled systems. All the sockets implemented by the `POSIXConnector` follow a client–server pattern, having the receiving entity set as a server and the delivering entity acting as a client.

The `ZMQConnector` class uses ZeroMQ as the link layer, an high-performance asynchronous messaging library, aimed at use in distributed or concurrent applications [41]. ZeroMQ abstracts the complexity of operations such as connection management and data serialization via simple APIs calls, and it is designed for both intra- and inter-machine communication with flexible built-in patterns, as discusses below. As illustrated in [Fig. 8](#), our approach leverages ZeroMQ. Specifically, the initial E3 setup employs a request–reply (REQ–REP) pattern – where the dApp issues the request (REQ) and the RAN side provides the reply (REP) – since this procedure occurs only once during the dApp’s lifecycle. In contrast, subsequent outbound and inbound data exchanges adopt a publish–subscribe (PUB–SUB) pattern, with the data-generating entity acting as the publisher (PUB) and the data-processing entity as the subscriber (SUB). This choice reflects the need to support continuous data transfer between the RAN and the dApp, in contrast to the one-time setup operation.

Both classes implement three transport-layer protocols – Transmission Control Protocol (TCP), IPC (using Unix Domain sockets), and Stream Control Transmission Protocol (SCTP) – to facilitate communication between local and remote hosts. While the implementation can support other transport protocols, in our design we focus on those that offer reliable message delivery (even if this introduces overhead) and excludes those, e.g., User Datagram Protocol (UDP), that do not guarantee data delivery and are not reliable enough for RAN operations. Furthermore, to maintain consistency, the `ZMQConnector` class includes SCTP support, despite the ZeroMQ link layer not yet offering full SCTP compatibility [41,42]. It is also worth mentioning that we decided to maintain the possibility to have a connection for the E3 to a remote virtual host, not co-located with the RAN node, since in the O-RAN vision the disaggregation of the network functions may also happen across diverse physical nodes in the network.

Data received through the `E3Connector` socket, regardless of the underlying link and transport layers, is processed by a callback mechanism operated by the subscription manager. This mechanism associates newly available data with the registered dApps that require it, thus enabling real-time data handling and avoiding data duplication. For the PDUs used for establishing communications between the RAN node and dApps, we employ Abstract Syntax Notation One (ASN.1) definitions for message formatting in the E3AP and E3SM, adhering to established standard 3GPP and O-RAN message exchange protocols. ASN.1 enables efficient, rapid, and compact serialization of data. Each class extending our `DApp` abstract class should implement its own SM and intelligent logic according to its use case, as shown in [Fig. 7](#).

In this work, we propose two different use cases leveraging dApps based on this framework with further implementation details discussed in Section 7.

6.2. OpenAirInterface and T-tracer

OAI is an open-source project that implements a 3GPP-compliant full 5G New Radio (NR) stack on general-purpose computing hardware and off-the-shelf Software-defined Radios (SDRs) [43,44]. It supports both CU-DU split and monolithic deployments, enabling flexibility in testing and development of RAN functions. Additionally, OAI provides implementations for key RAN components, including gNB and UE, offering a robust platform for validating 5G network functionalities and integrating advanced features such as slicing and AI/ML-based optimizations.

The project implements Open RAN-compliant functionalities, including code functions to connect with the Near-RT RIC (both the O-RAN Software Community (OSC) [45] version and custom versions such as FlexRIC [46]) and the Non-RT RIC.

We have extended the original OAI codebase by introducing a new E3 Agent module that facilitates exchanges based on the architecture proposed in this paper, with a particular attention to the message exchange between the dApp and the xApps over the E3AP, and the RAN control functionalities employed by the proposed use cases.

It is worth mentioning that OAI integrates T-tracer [47], a programmable toll that provides an external connection via TCP sockets, enabling developers to extract and collect metrics during RAN operations for debugging and analysis purposes. Our module integrates with the T-tracer tool, enabling real-time data sensing and extraction, and making this data accessible through the E3 interface. We extended the T-tracer socket utilities to support IPC via Unix domain sockets, enhancing the communication and reducing the overhead generated in the data exchange and, consequently, the latency. Moreover, we developed a C-language implementation of the `E3Connector` to integrate it into the OAI codebase through a merge request, avoiding external dependencies between the dApp framework and cellular stack. This component works in the very same way as its Python counterpart, enabling different link- and transport-layer tests over different virtual hosts by design. We will propose the integration of this connector in the OAI main codebase upon acceptance of this manuscript. Finally, we extended the OAI codebase to implement the data extraction and the RAN control of the use cases discussed in Section 7, supporting real-time control loops.

6.3. Evaluating real-time capabilities of dApps

We ran more than two hundreds experiments on two different testbeds to assess and evaluate the capabilities of our framework executing the `SpectrumSharingDapp`. A detailed discussion of the use case is provided in Section 7.1. In the following paragraphs we focus specifically on the benchmarking of the framework in terms of real-time control capabilities.

The first test bench for this evaluation is the Colosseum testbed, a wireless network emulator [20] consisting of SDR-equipped servers and a channel emulator capable of reproducing real-world wireless channel effects, such as path loss and fading. The second is Arena [21], a publicly available over-the-air indoor testbed deployed in a laboratory environment. Both systems used the same version of OAI and our dApp prototype, running on Dell PowerEdge R730 servers and Universal Software Radio Peripheral (USRP) x310 SDRs. On Arena, the UE is a One Plus Nord 5G AC2003, while on Colosseum we use the OAI softUE. We use the OAI configuration for the Frequency Range 1 (FR1) band n78, with a center frequency of 3.6192 GHz and a bandwidth of 40 MHz.

Each experiment has a duration of six minutes. However, for the scope of the measurements reported, we only consider the five-minute window where UE downlink transmissions are active (i.e., we exclude the network bootstrap, UE attachment and setup phases). We present the results related to the performance of the real-time control loop.

One column for the table We evaluated different options for minimizing overhead, focusing on the link layer and the transport layer. Specifically, we considered the two `E3Connector` implementation options based on raw POSIX functions and ZeroMQ. Moreover, we have implemented and evaluated three different options for the transport layer built on top of these two links, i.e., the TCP protocol, the SCTP protocol and the IPC based on Unix Domain sockets.

[Table 2](#) presents a summary of the overheads introduced by the different transport layers protocols, i.e., the number of bytes introduced by the protocol stack to guarantee message delivery. The first protocol we consider is TCP, which is commonly used for inter-process communication of application that can scale over different virtual hosts. The

Table 2
Overhead introduced in the E3AP message exchange using the E3Connector.

Protocol	Overhead percentage (no ASN)	Overhead percentage (with ASN)
TCP	20%	20.05%
SCTP	42%	42.05%
IPC	0%	0.05%

additional overhead caused by TCP is 20%, making it a good candidate in the case where the dApp is connected across virtual hosts over the same network. As a second candidate, we evaluated the SCTP protocol since it is an affirmed transport-layer protocol for cellular networks. Results show that SCTP introduces significant overhead, which reduces goodput, increases latency, and can potentially slow down the control loop process and result in late control action enforcement. While this overhead is essential for other message exchange procedures in O-RAN, such as E2AP, dApps can instead leverage IPC. This approach eliminates overhead entirely, as IPC allows direct communication between processes on the same host. Unlike network-based protocols, IPC avoids the need for data encapsulation and transmission across multiple network hosts, significantly reducing latency and processing requirements. Indeed, SCTP brings many benefits such as multi-homing support, path Maximum Transmission Unit (MTU) discovery, redundant transmission for reliability. However, these come with additional overhead, e.g., due to encapsulation and decapsulation of link and transport layers, that can slow down the flow of the dApp, which make IPC a better candidate to deliver real-time control over the same host.

Finally, we report the overhead generated by the ASN.1 implementation of the E3AP-SM PDUs application layer. In the design of such definitions, we opted for a minimalist approach that ensures the correct implementation of the architectural framework by providing only the necessary PDUs (i.e., Setup, Subscription, Indication, and Control) to enable a closed control loop, while using an architecture that can be extended with more PDUs in the future if needed. The ASN.1 implementation offers a structured messaging infrastructure, but ASN.1 structures and its padding generates overhead. This overhead is the necessary trade-off for representing the advanced logical interactions between the dApps and the RAN through the E3SMs. To reduce overhead, in our implementation we use the Packet Encoding Rule (PER), which is a more compact representation of ASN.1 that minimizes the overhead by omitting extra information such as tags and lengths of packet fields. For instance, if fields have a limited number of possible values, the encoding will use just enough bits to represent those values, making the encoded data smaller compared to other encoding rules. In our experiments, the ASN.1 PDUs generate a constant 0.05% overhead across all protocols, which adds to the overhead generated by each protocol.

Fig. 9 reports latency measurements for all operations required to perform dApp-based control using our prototype. Specifically, we measure the latency required to perform the four fundamental operations of dApp control: *Collect Data*, *Process Data*, *Create Control*, and *Deliver Control*. We also report the *Cumulative* latency, defined as the sum of all the previous operations. Initially, upon the generation of the RAN measurements, the dApp must *Collect Data* from the RAN using the T-tracer and the *E3 Indication Message*. Subsequently, it moves to *Process Data*, where the measurements are analyzed through either a heuristic or data-driven approach. Following this, the dApp proceeds to *Create Control*, which involves generating the control action and formulating the *E3 Control Message*. Finally, in the *Deliver Control* stage, this message is transmitted back to the RAN function.

To evaluate our prototype in a practical use case that relies on real-time operations and inference, we provide results obtained by executing the spectrum sharing dApp (described in Section 7.1). As we will discuss later, the input data to be retrieved consists of I/Q samples, and the output of the dApp corresponds to the list of PRBs that the MAC scheduler should not allocate due to the presence of external interference (e.g., incumbents or jammers). We also consider

the implementation that provided the best performance which is the case of an ZMQConnector-based E3 interface serving data through IPC and ASN.1, and provide results for varying dApp input and output sizes. Specifically, we considered 16 different configurations by varying two parameters: (i) the size of the payload in the *E3 Indication Messages* (i.e., input of the dApp); and (ii) the size of the *E3 Control Messages* (i.e., output of the dApp). For the Indication Messages, each I/Q sample is encoded using 4 bytes, and we evaluated four payload sizes of 1536, 3072, 6144, and 8192 bytes, which correspond to 384, 768, 1536, and 2048 I/Q samples, respectively. For the Control Messages, we defined payload lengths in bytes based on the number of PRBs blocked by the dApp of the aforementioned use case: 0 (no PRBs), 4 PRBs (16 bytes and about 1 MHz of blocked bandwidth), and 8 PRBs (32 bytes and about 2 MHz). It is worth mentioning that the list of PRBs to be blocked is of variable size and depends on the resolution of the Indication Message (i.e., how many I/Q samples in the frequency domain are fed to the dApp). For this reason, we also consider the maximum number of PRBs that can be blocked based on the amount of I/Qs being included in the Indication message size. Specifically, for Indication Message size of 384, 768, 1536, and 2048 we consider Control Messages of maximum size 128, 256, 208, and 380 bytes, respectively.

Results in Fig. 9 demonstrate that our dApp prototype achieves real-time control loops with an average aggregated latency of 400 μ s, consistently remaining below 450 μ s. As expected, the usage of IPC through ZeroMQ substantially reduces the impact of the socket operations on the total control loop time, making the operation of *Data Collection* the indication message and *Control Delivery* the fastest operations in the loop. The major contribution to the elapsing time of the loop is represented by the *Elaboration* of the Indication Message, i.e., the ASN.1 decode and the intelligent control logic implemented by the dApp to analyze the data. This behavior is also highlighted in the case of Control Messages with size 64 and 95 bytes, shown in Figs. 9(b) and 9(d), respectively, where the presence of padding in the E3AP PDU causes additional memory allocation, thus increasing the elaboration time. The average total elapsed time for each of the 16 cases is consistently below the 1 ms threshold as shown in Fig. 9, which is defined as the limit for real-time use cases involving the manipulation of I/Q samples (see Fig. 2, Section 3).

7. Empowering RAN control through dApps: use cases

After demonstrating that dApps can indeed enable real-time inference in O-RAN systems, in this section, we showcase two relevant RAN control use cases that benefit from dApps. These examples illustrate how the data extraction and dApp-driven control capabilities provided by our approach can enhance network performance and make network management more agile.

7.1. Use case: Spectrum sharing

Dynamic spectrum sharing among multiple wireless access technologies is expected to play a major role in the next-generation wireless system design. Currently, a popular spectrum sharing example in the U.S. is the case of the Citizens Broadband Radio Service (CBRS) band, where mobile network operators may coexist with high-priority incumbent federal radar and satellite systems users in the 3.55–3.7 GHz frequency range. When considering 5G NR systems, the CBRS band is a subset of the bands n77 and n78.

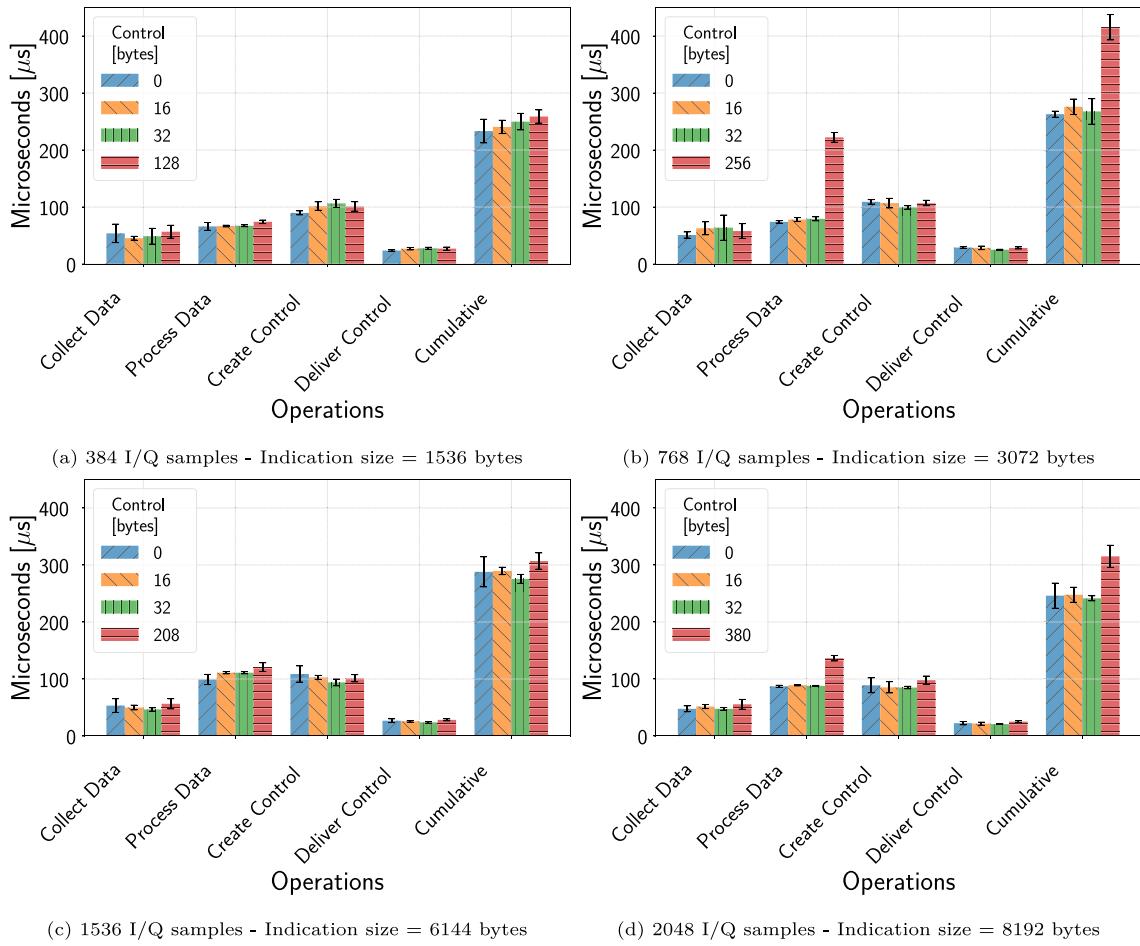


Fig. 9. Latency measurements for the real-time control loops using dApps with E3AP over ZeroMQ and IPC.

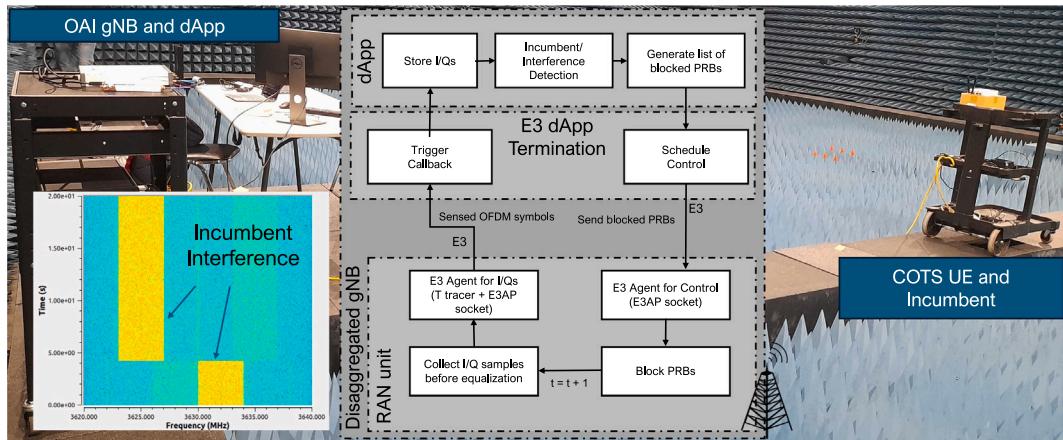


Fig. 10. Intelligent real-time control loop for the PRB blacklisting. We omit in this figure the initial E3AP procedures.

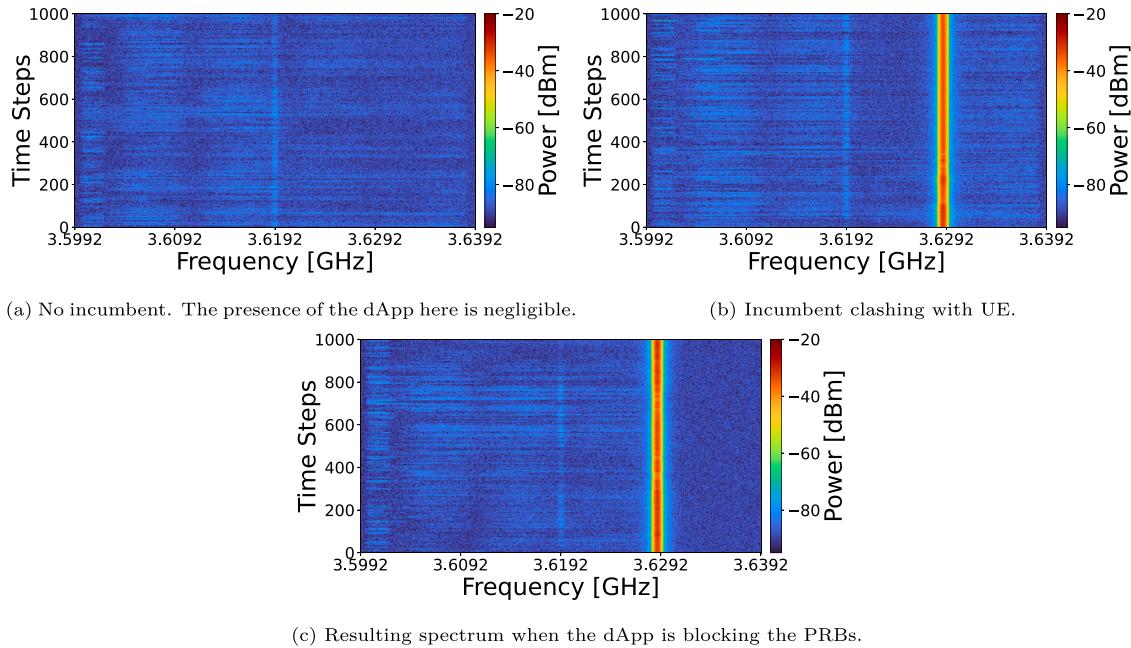


Fig. 11. Spectrum sharing dApp cases evaluated in the experiments: Sensed Spectrum of a 5G gNB-UE communication with the incumbent invading the spectrum of the gNB with and without the dApp.

In our recent work [48], we leveraged our dApp framework to implement a real-time RAN-driven spectrum sharing system. In this system, the gNB can communicate while simultaneously performing the spectrum sensing task in the CBRS band. Moreover, the gNB can adapt its communication parameters (e.g., frequencies or the list of PRBs to not be used for scheduling purposes due to presence of high interference) whenever a primary incumbent user is detected. The spectrum sharing dApp that we developed follows the prototype mentioned in the previous section and illustrated in Fig. 10. The dApp connects with the gNB using the E3 and performs the following tasks:

1. Extract I/Q samples from dedicated symbols reserved for spectrum sensing at the gNB through a callback registered with the E3Interface, enabling real-time sensing and data extraction;
2. Leverage an inference algorithm to process the I/Q samples and detect incumbent users by computing the magnitude of the samples and comparing each magnitude with a fixed threshold previously calibrated;
3. If an incumbent is detected, create a list of the PRBs affected by the incumbent and that should not be used by the gNB to schedule transmissions;
4. Deliver such list to the RAN node as a control action through the E3 interface.

We analyzed three different configurations of the spectrum in Fig. 11 using a Keysight EXA Spectrum Analyzer N9010B as an external observer. The first one, shown in Fig. 11(a), depicts an unbounded TCP downlink transmission of a 5G gNB to one UE.

In this case, the MAC scheduler of the gNB is allowed to use all possible PRBs. Fig. 11(b) introduces a narrowband incumbent in the spectrum of the gNB, making the 5G RF signal clash with the incumbent RF signal and degrading the gNB-UE performance. When the spectrum sharing dApp interacts with the gNB, its control logic analyzes the sensed I/Qs and is able to detect the presence of the incumbent and to create a control action to block the clashing PRBs, making them unavailable to the scheduler. Finally, the dApp delivers the control action to the RAN, vacating the affected portion of the spectrum, as

shown in Fig. 11(c). Under these conditions, 5G communication can continue to coexist with the incumbent, albeit with certain limitations.

Specifically, we would like to mention that the currently available OAI scheduler implements a type 1 resource allocation scheme as defined in the 3GPP Technical Specification (TS) 38.214 Section 6.1.2.2 [49], which requires that the PRBs assigned to a single UE form a contiguous, non-interleaved sequence of virtual resource blocks. As a consequence, when we mute a set of PRBs already occupied by the incumbent, we notice that the scheduler also blocks all scheduling activities in the PRBs occupying the rightmost part of the spectrum. Although this is a limitation that does not pertain our spectrum sensing dApp, it is also fair to point out that this specific implementation based on type 1 scheduler might reduce unnecessarily the list of PRBs that can be used for data transmission. However, we also point out that in the same specifications, precisely in Section 6.1.2.1 [49], the standard defines resource allocation of type 0, which allows for the allocation of noncontiguous PRBs to the same UE. Unfortunately, this latter type is not implemented in OAI at the time of this writing, and the evaluation of our dApp under this configuration is left for future studies.

As we will discuss in the next section, the results reported in Table 3 show that, although OAI's type 1 scheduler results in an unused portion of the spectrum (right portion of the spectrum in Fig. 11(c)), the usage of a reduced portion of the spectrum does not significantly reduce UE downlink performance.

Performance Evaluation – Table 3 summarizes the experiment results of the Spectrum Sharing use case for the two testbeds and the different test configurations considered. A 95% confidence interval, using a Z-value of 1.96, was applied to the results and it is reported along with the averages of each measurements. Given the radio configuration of the experiments presented in Section 6.3, the incumbent may interfere with the gNB spectrum if transmitting between 3.5992 GHz and 3.6392 GHz. We present the throughput results across various configurations: with and without the incumbent, and with and without the spectrum sharing dApp. These results highlight the scenarios where spectrum is shared and where it is not. For these experiments, the incumbent signal has been created using the `uhd_siggen` software set up to generate an uniform noise transmitted at 3.63 GHz with a

Table 3

Summary of the performance of OAI and spectrum sharing dApp. Average throughput reported with the 95% confidence interval.

Testbed	dApp	Incumbent	Throughput (Mbps)	Shared spectrum?
Colosseum	N	N	71.34 ± 1.28	N/A (gNB only)
Colosseum	Y	N	71.53 ± 0.76	N/A (gNB only)
Colosseum	N	Y	49.52 ± 3.37	No coordination
Colosseum	Y	Y	53.78 ± 1.55	Shared spectrum
Arena	N	N	77.98 ± 1.31	N/A (gNB only)
Arena	Y	N	76.37 ± 1.87	N/A (gNB only)
Arena	N	Y	38.97 ± 2.95	No coordination
Arena	Y	Y	43.86 ± 1.52	Shared spectrum

sampling rate of 1 MHz, an amplitude of 0.5 and gain of 90 dB on Arena and 60 dB on Colosseum. A dBm threshold used to determine whether to block the PRBs was calibrated specifically for each experimental environment. It is worth mentioning that each testbed is characterized by different radios with diverse sensitivity, and different RF conditions. For this reason, we set the noise floor for the two testbeds considered in our analysis to two different value. Specifically, we set the noise floor to 20 dBm and 53 dBm for Arena and Colosseum, respectively.

The choice of these values is conservative. Indeed, a too high of a threshold value might fail in detecting incumbent activities in low SINR regimes, and might result in gNB transmissions that overlap with the incumbent and cause severe interference. Similarly, a too low of a threshold might misinterpret noise as incumbent signals, and block all PRBs. In our case, we have evaluated the thresholds heuristically via experiments and data analytics. Note that, in general, a conservative threshold is to be preferred as incumbents have always priority and gNBs must vacated promptly the spectrum.

Results in Table 3 show that the execution of the dApp (cases where dApp = Y and Incumbent = N) does not introduce any significant impact on throughput if compared to the case where the dApp is not enabled (cases where dApp = N and Incumbent = N). Specifically, in the Colosseum testbed, the average throughput measured without the dApp is 71.34 Mbps, while we measure an average throughput of 71.53 Mbps when the dApp is enabled. Similarly, in the Arena testbed, the throughput decreases slightly from 77.98 Mbps without the dApp to 76.37 Mbps with the dApp. Although the average values may suggest performance degradation, the confidence intervals in Table 3 indicate that throughput takes values within the same confidence range, demonstrating that the dApp introduces minimal overhead and maintains performance parity.

In the case where the Incumbent is present (i.e., when Incumbent = Y), the RF signals of gNB and incumbent interfere with each other. The lack of spectrum sharing (cases where dApp = N and Incumbent = Y) causes the gNB to use PRBs affected by interference, which significantly degrades the UE throughput by nearly 50% across both testbeds. We instead notice that the presence of the spectrum sharing dApp (cases where dApp = Y and Incumbent = Y) is beneficial, as the blocking of PRBs affected by interference allows the sharing of a portion of the spectrum with the incumbent while preventing disruptive interference. This brings a slight improvement of the UE throughput because preventing the use of interfered PRBs results in less interference, lower errors and reduced need for retransmissions.

7.2. Use case: Sensing and positioning

Cellular operations at higher frequencies enable the use of wider bandwidths (e.g., Frequency Range 2 (FR2), Frequency Range 3 (FR3), and THz) and massive antenna arrays. Thanks to these technologies, beyond 5G systems are expected to offer not only higher data rates but also high-resolution positioning and environment sensing capabilities [50]. In current UL-based positioning methods, the gNB derives specific positioning-related measurements (defined in 3GPP standards) from the UL channel estimates and forwards them to the Location

Management Function (LMF) entity within the 5G core network. These measurements can be encoded into a few bits and are generally derived from low-complexity signal processing algorithms to comply with the real-time nature of the gNB.

On the other hand, super-resolution and ML-based algorithms offer precise sensing and positioning capabilities that outperform conventional algorithms [51,52]. However, they usually require data that is not accessible from outside the RAN, such as the CIR. Therefore, dApps are a natural tool to address this problem.

In our recent work [53], we have used the dApp framework in an UL ranging/positioning system, where the distance between the UE and gNB is computed based on the UE's UL CIR available at the gNB and exposed to the dApp. To demonstrate this method, we have developed a testbed consisting of an OAI gNB (enhanced with the E3 agent), a ranging dApp, and OAI UE. The dApp connects with the OAI gNB using the E3 agent and performs the following tasks: (i) extracts multiple wideband UL CIR measurements of a UE from the gNB; and (ii) runs a super-resolution algorithm [53] on the collected CIR measurements to find the distance between the UE and the gNB.

The above described ranging dApp framework is tested and validated with real-world experiments. The experimental setup consists of a single antenna gNB and a UE communicate over a line-of-sight channel. The gNB and the UE rely on the OAI protocol stack and USRP B210 software-defined boards. Furthermore, the SC2430 NR signal conditioning module is used as an external RF front-end at the gNB [54]. The system operates at band n78 with a carrier frequency of 3.69 GHz and a bandwidth of 40 MHz.

Performance Evaluation – The performance of the proposed ranging scheme is evaluated in terms of empirical Cumulative Distribution Function (CDF) of the distance estimation error at an UL SNR of -20 dB. The CDF is obtained from 48,000 CIR measurements by fixing the position of the gNB and moving the UE in a straight line between 3 to 10 meters with a 1-meter increment. At every point, a total of 6000 measurements are collected. The *ranging dApp* is used to extract these measurements as shown in Fig. 12. The CDF of the range estimation error with a super resolution Multiple Signal Classification (MUSIC) algorithm with $M=20$ to 60 number of CIR measurements is shown in Fig. 13. For details on ranging algorithms and prototype design, we refer the readers to [53]. While the *ranging dApp* here is used to collect measurements during the experiment, which are then processed offline to obtain the range error CDF, it is important to highlight that it can also be used for real-time CIR collection and distance evaluation. The measurement collection latency performance is similar to that described in Section 6.3.

8. Conclusions

In this paper, we introduced an extension of the O-RAN architecture focused on enabling (i) real-time control in the RAN, and (ii) the interaction of the O-RAN stack with the user-plane of the network. This is achieved through dApps, lightweight microservices co-located with DUs and CUs, and holistically managed through components such as the RICs and SMO. We first described the role that dApps have in enabling

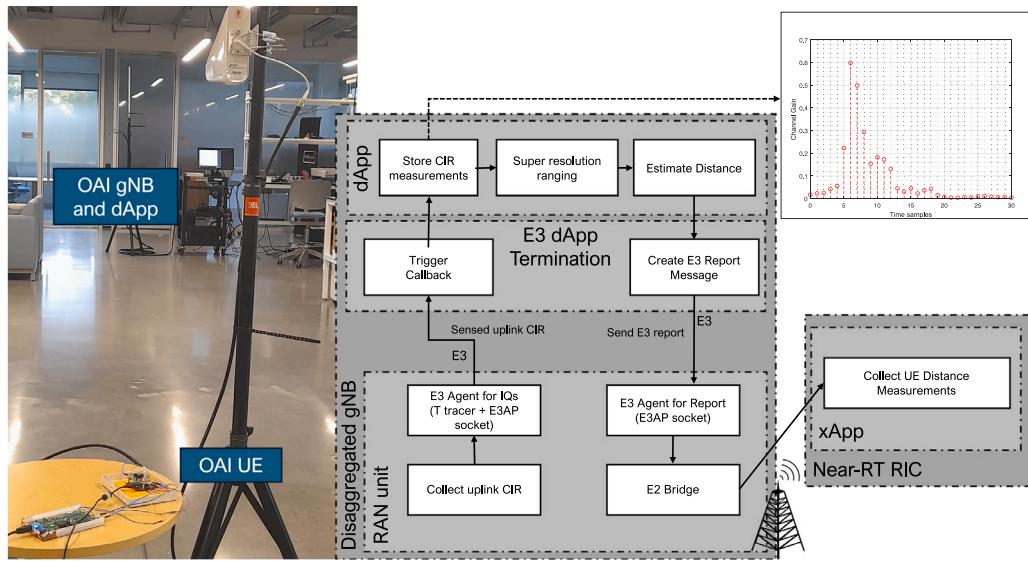


Fig. 12. Setup for the Sensing and Positioning dApp and UL CIR extracted by the dApp.

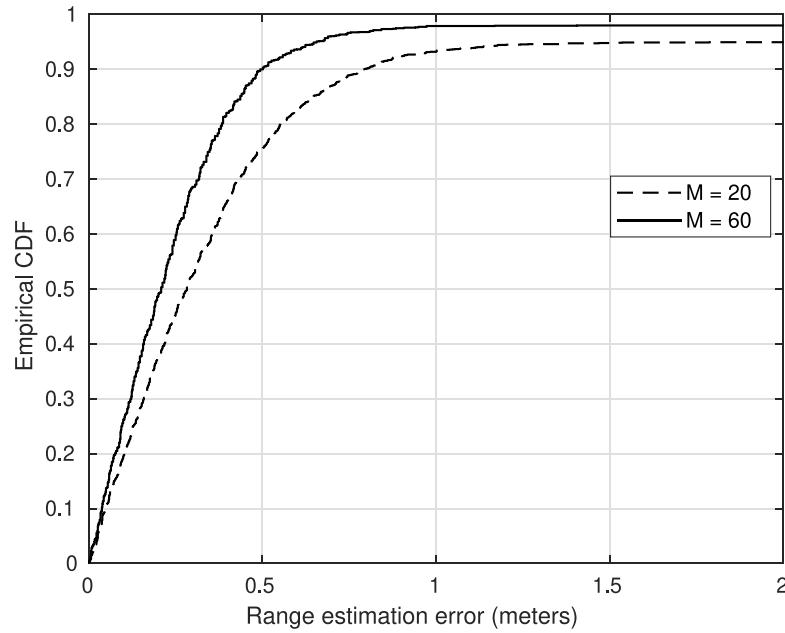


Fig. 13. Cumulative Distribution Function (CDF) of the range estimation error.

new use cases for RAN optimization, including inference based on I/Q samples, reference signals available at the physical layer, or as complex ISAC systems. Based on the data and control requirements defined by these use cases, we discussed the architecture for the integration of dApps and RAN nodes. We also provided insights on the design and implementation of a set of APIs that RAN nodes need to expose to dApps, coordinated by an API broker called E3 agent. We then discussed the LCM for dApps, including steps for development, onboarding, and deployment of dApps, as well as procedures for interaction with the RICs and SMO.

Then, we proposed a reference implementation for dApps based on OAI, which is publicly available to enable research and development

of applications and use cases relying on real-time control loops in Open RAN. We described the dApp framework, based on Python, and the integration with OAI. We then presented two use cases based on our reference implementation. For the spectrum sharing one, a dApp performs spectrum sensing at the DU, understands which portions of the spectrum are occupied by incumbents, and coordinates with the DU scheduler to avoid scheduling over it. In the analysis of this use case, results show that while the presence of a continuous incumbent affects the performance of the 5G network, the use of a dApp enables the sharing of the spectrum bringing a slight improvement of the 5G network performance.

For the positioning use case, the dApp processes uplink reference signals to extract the uplink channel response and perform ranging for the UE. We profiled the performance of our dApp framework, demonstrating that real-time control loops under 1 ms are achievable in O-RAN. In our implementation, the average control loop duration is slightly less than 400 μ s, on average.

As part of our future work, we are extending the dApp framework to other open stacks, including NVIDIA ARC-OTA [55], and integrate the dApps in a continuous integration, deployment, and testing framework to continuously test it on and up-to-date gNB protocol stack.

CRediT authorship contribution statement

Andrea Lacava: Writing – review & editing, Writing – original draft, Software, Resources, Methodology, Data curation. **Leonardo Bonati:** Writing – review & editing, Writing – original draft, Software, Methodology. **Niloofar Mohamadi:** Writing – original draft, Formal analysis. **Rajeev Gangula:** Software, Conceptualization. **Florian Kaltenberger:** Writing – review & editing, Supervision. **Pedram Johari:** Writing – review & editing, Project administration, Formal analysis. **Salvatore D’Oro:** Writing – review & editing, Validation, Investigation. **Francesca Cuomo:** Supervision. **Michele Polese:** Writing – original draft, Validation, Supervision, Project administration, Funding acquisition. **Tommaso Melodia:** Writing – review & editing, Validation, Supervision, Project administration, Investigation, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This article is based upon work partially supported by OUSD(R&E) through Army Research Laboratory Cooperative Agreement Number W911NF-24-2-0065. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. The work was also partially supported by SERICS (PE00000014) 5GSec project, CUP B53C22003990006, under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU, and by the U.S. National Science Foundation under grants CNS-1925601, CNS-2117814, and CNS-2112471.

Appendix. dApp deployment process

The following table provides a detailed breakdown of the LCM process for the dApp, outlining the steps, roles, and conditions involved in its deployment and operation.

Table A.4
dApp Deployment.

Use case stages and Evolution/Specification	
Goal	To instantiate and configure a dApp on the CU/DU, facilitating its deployment, management, and operationalization directly via the SMO and the CU/DU.
Actors and Roles	<ul style="list-style-type: none"> SMO (Service Management and Orchestration): Orchestrates the deployment, configuration, and lifecycle management of the dApp. DU (Distributed Unit): Possible host of the dApp and provides necessary network and compute resources for its operation. CU (Central Unit): Possible host the dApp and provides necessary network and compute resources for its operation. Network Function Orchestration (NFO): Interfaces with the DU for deploying and configuring the dApp. Deployment Management Services (DMS): Manages the deployment resources on the DU.
Assumptions	<ul style="list-style-type: none"> The SMO and CU/DU are available and operational. The dApp package has been validated, verified, and cataloged in the SMO’s runtime library. The DU is pre-configured and ready to host dApp instances.
Preconditions	<ul style="list-style-type: none"> The dApp is onboarded and certified within the SMO. The CU/DU has the necessary resources and access to the container images required for the dApp deployment.
Begins When	A Network Function Install Project Manager initiates a request to the SMO for deploying a new dApp instance on the CU/DU.
Step 1 (M)	The SMO receives a service request to deploy the dApp instance on the CU/DU.
Step 2 (M)	The SMO decomposes the service request and identifies all dApps to be deployed and their deployment order.
Step 3 (M)	The SMO determines which DU, or CU, and deployment parameters to use. This is based on policies or explicit input from the Network Function Install Project Manager.
Step 4 (M)	The SMO retrieves the CloudNativeDescriptor for the dApp from the runtime library.

(continued on next page)

Table A.4 (continued).

Use case stages and Evolution/Specification	
Step 5 (M)	The SMO directs the O-Cloud DMS using O2 to create the dApp deployment.
Step 6 (M)	DMS allocates the necessary compute, storage, and network resources on the CU/DU as per the dApp deployment request.
Step 7 (M)	The SMO sets up the initial configuration for the dApp, such as environment variables, network policies, and access parameters.
Step 8 (M)	DMS deploys the dApp container(s) on the CU/DU, setting up the necessary resources and connectivity.
Step 9 (M)	DMS notifies the SMO that the dApp deployment has been successfully instantiated and provides a Deployment ID.
Step 10 (M)	The SMO updates its dApp inventory with the new Deployment ID and deployment status.
Step 11 (M)	The deployed dApp instance reads its initial configuration from the provided parameters and begins its operation.
Step 12 (O)	The SMO continuously monitors the dApp's health, performance, and connectivity. It can also perform scaling, updates, or redeployment as required.
Step 13 (M)	The SMO informs the Network Function Install Project Manager of the overall success or failure of the request.
Ends When	The dApp instance is successfully deployed, configured, and operational on the CU/DU, with the SMO actively managing its lifecycle.
Post Conditions	<ul style="list-style-type: none"> The dApp is actively running and functional on the CU/DU. The SMO has an updated inventory reflecting the deployed dApp, and lifecycle management is in place.
Exceptions	If the deployment fails, the DMS notifies the SMO, and the SMO informs the Network Function Install Project Manager to take corrective actions.

Data availability

Data will be made available on request.

References

- [1] D. Johnson, D. Maas, J. Van Der Merwe, *NexRAN: Closed-loop RAN slicing in POWDER-A top-to-bottom open-source open-RAN use case*, in: Proceedings of the 15th ACM Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization, 2022, pp. 17–23.
- [2] M. Polese, L. Bonati, S. D’Oro, S. Basagni, T. Melodia, *Colo-RAN: Developing machine learning-based xApps for open RAN closed-loop control on programmable experimental platforms*, *IEEE Trans. Mob. Comput.* 22 (10) (2022) 5787–5800.
- [3] A. Lacava, M. Polese, R. Sivaraj, R. Soundararajan, B.S. Bhati, T. Singh, T. Zugno, F. Cuomo, T. Melodia, *Programmable and customized intelligence for traffic steering in 5G networks using open RAN architectures*, *IEEE Trans. Mob. Comput.* 23 (4) (2024) 2882–2897.
- [4] M. Dryjański, Ł. Kufacz, A. Kliks, *Toward modular and flexible open ran implementations in 6G networks: Traffic steering use case and O-RAN xApps*, *Sensors* 21 (24) (2021) 8173.
- [5] K. Suzuki, J. Nakazato, Y. Sasaki, K. Maruta, M. Tsukada, H. Esaki, *Toward B5G/6G connected autonomous vehicles: O-RAN-Driven millimeter-wave beam management and handover management*, in: *IEEE Conference on Computer Communications Workshops, INFOCOM WKSHPS*, 2024, pp. 1–6.
- [6] L. Kundu, X. Lin, R. Gadiyar, *Towards energy efficient RAN: From industry standards to trending practice*, 2024, arXiv preprint [arXiv:2402.11993](https://arxiv.org/abs/2402.11993).
- [7] O.T. Başaran, M. Başaran, D. Turan, H.G. Bayrak, Y.S. Sandal, *Deep autoencoder design for RF anomaly detection in 5G O-RAN near-RT RIC via xApps*, in: *2023 IEEE International Conference on Communications Workshops, ICC Workshops*, IEEE, 2023, pp. 549–555.
- [8] A. Tripathi, J.S.R. Mallu, M.H. Rahman, A. Sultana, A. Sathish, A. Huff, M. Roy Chowdhury, A.P. Da Silva, *End-to-end O-RAN control-loop for radio resource allocation in SDR-based 5G network*, in: *IEEE Military Communications Conference, MILCOM*, 2023, pp. 253–254.
- [9] A. Al-Shawabka, F. Restuccia, S. D’Oro, T. Jian, B.C. Rendon, N. Soltani, J. Dy, S. Ioannidis, K. Chowdhury, T. Melodia, *Exposing the fingerprint: Dissecting the impact of the wireless channel on radio fingerprinting*, in: *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*, IEEE, 2020, pp. 646–655.
- [10] S.-D. Wang, H.-M. Wang, C. Feng, V.C.M. Leung, *Sequential anomaly detection against demodulation reference signal spoofing in 5G NR*, *IEEE Trans. Veh. Technol.* 72 (1) (2023) 1291–1295.
- [11] D. Uvaydov, S. D’Oro, F. Restuccia, T. Melodia, *Deepsense: Fast wideband spectrum sensing through real-time in-the-loop deep learning*, in: *IEEE INFOCOM 2021-IEEE Conference on Computer Communications*, IEEE, 2021, pp. 1–10.
- [12] R. Chen, J.-M. Park, Y.T. Hou, J.H. Reed, *Toward secure distributed spectrum sensing in cognitive radio networks*, *IEEE Commun. Mag.* 46 (4) (2008) 50–55.
- [13] F. Liu, Y. Cui, C. Masouros, J. Xu, T.X. Han, Y.C. Eldar, S. Buzzi, *Integrated sensing and communications: Toward dual-functional wireless networks for 6G and beyond*, *IEEE J. Sel. Areas Commun.* 40 (6) (2022) 1728–1767.
- [14] Z. Qadir, K.N. Le, N. Saeed, H.S. Munawar, *Towards 6G Internet of Things: Recent advances, use cases, and open challenges*, *ICT Express* 9 (3) (2023) 296–312.
- [15] S. D’Oro, M. Polese, L. Bonati, H. Cheng, T. Melodia, *dApps: Distributed applications for real-time inference and control in O-RAN*, *IEEE Commun. Mag.* 60 (11) (2022) 52–58.
- [16] W.-H. Ko, U. Ghosh, U. Dinesha, R. Wu, S. Shakkottai, D. Bharadia, *EdgeRIC: Empowering real-time intelligent optimization and control in nextg cellular networks*, in: *21st USENIX Symposium on Networked Systems Design and Implementation, NSDI 24*, 2024, pp. 1315–1330.
- [17] P.S. Upadhyaya, N. Tripathi, J. Gaedert, 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.* 37 (5) (2023) 7–15.
- [18] C. Liu, G. Aravindhan, A. Kak, N. Choi, *TinyRIC: Supercharging O-RAN base stations with real-time control*, in: *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, in: *ACM MobiCom ’23*, Association for Computing Machinery, New York, NY, USA, 2023.
- [19] Northeastern University, NVIDIA, Mavenir, MITRE, Qualcomm, *dApps for Real-Time RAN Control: Use Cases and Requirement*, Research Report, O-RAN next Generation Research Group (nGRG), 2024, report ID: RR-2024-10, [Online]. Available: <https://mediastorage.o-ran.org/nrg-rr/nGRG-RR-2024-10-dApp%20use%20cases%20and%20requirements.pdf>.
- [20] M. Polese, L. Bonati, S. D’Oro, P. Johari, D. Villa, S. Velumani, R. Gangula, M. Tsampazi, C.P. Robinson, G. Gemmi, et al., *Colosseum: The open RAN digital twin*, *IEEE Open J. Commun. Soc.* (2024).
- [21] L. Bertizzolo, L. Bonati, E. Demirors, A. Al-shawabka, S. D’Oro, F. Restuccia, T. Melodia, *Arena: A 64-antenna SDR-based ceiling grid testing platform for Sub-6GHz 5G-and-beyond radio spectrum research*, *Comput. Netw.* 181 (2020) 107436, [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1389128620311257>.
- [22] M. Polese, L. Bonati, S. D’oro, S. Basagni, T. Melodia, *Understanding O-RAN: Architecture, interfaces, algorithms, security, and research challenges*, *IEEE Commun. Surv. Tutor.* (2023).

[23] A.S. Abdalla, P.S. Upadhyaya, V.K. Shah, V. Marojevic, Toward next generation open radio access networks: What O-RAN can and cannot do! *IEEE Netw.* 36 (6) (2022) 206–213.

[24] X. Foukas, B. Radunovic, M. Balkwill, Z. Lai, Taking 5G RAN analytics and control to a new level, in: Proceedings of the 29th Annual International Conference on Mobile Computing and Networking, in: ACM MobiCom '23, Association for Computing Machinery, New York, NY, USA, 2023.

[25] S. Maxenti, S. D'Oro, L. Bonati, M. Polese, A. Capone, T. Melodia, ScalO-RAN: Energy-aware network intelligence scaling in open RAN, in: Proc. of IEEE Intl. Conf. on Computer Communications, INFOCOM, 2024.

[26] L. Baldesi, F. Restuccia, T. Melodia, ChARM: NextG spectrum sharing through data-driven real-time O-RAN dynamic control, in: IEEE INFOCOM 2022 - IEEE Conference on Computer Communications, 2022, pp. 240–249.

[27] D. Villa, D. Uvaydov, L. Bonati, P. Johari, J.M. Jornet, T. Melodia, Twinning Commercial Radio Waveforms in the Colosseum Wireless Network Emulator, in: Proceedings of the 17th ACM Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization, 2023, pp. 33–40.

[28] M. Polese, F. Restuccia, T. Melodia, DeepBeam: Deep waveform learning for coordination-free beam management in mmWave networks, in: Proceedings of the Twenty-Second International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing, 2021, pp. 61–70.

[29] J. Groen, M. Belgiovine, U. Demir, B. Kim, K. Chowdhury, TRACTOR: Traffic analysis and classification tool for open RAN, in: ICC 2024 - IEEE International Conference on Communications, 2024, pp. 4894–4899.

[30] O-RAN Working Group 3, O-RAN E2 service model (E2SM), lower layers control 1.0, 2025, ORAN-WG3.TS.E2SM-LLC-R004-v01.00 Technical Specification.

[31] Y. Chen, Y.T. Hou, W. Lou, J.H. Reed, S. Kompella, M3: A sub-millisecond scheduler for multi-cell MIMO networks under C-RAN architecture, in: IEEE INFOCOM 2022 - IEEE Conference on Computer Communications, 2022, pp. 130–139.

[32] Qualcomm, Dell Technologies, Spectrum Sharing based on Shared O-RUs, Research Report, O-RAN next Generation Research Group (nGRG), 2023, Report ID: RR-2023-05, [Online]. Available: https://mediastorage.o-ran.org/ngrg-rr/nGRG-RR-2023-05-Spectrum_Sharing_with_Shared_O-RU-v1_0.pdf.

[33] M. Giordani, M. Polese, A. Roy, D. Castor, M. Zorzi, A tutorial on beam management for 3GPP NR at mmWave frequencies, *IEEE Commun. Surv. & Tutorials* 21 (1) (2019) 173–196.

[34] H. Cheng, P. Johari, M.A. Arfaoui, F. Periard, P. Pietraski, G. Zhang, T. Melodia, Real-time AI-enabled CSI feedback experimentation with open RAN, in: 2024 19th Wireless on-Demand Network Systems and Services Conference, WONS, 2024, pp. 121–124.

[35] I. Palamà, S. Bartoletti, G. Bianchi, N.B. Melazzi, 5G positioning with SDR-based open-source platforms: Where do we stand? in: 2022 IEEE 11th IFIP International Conference on Performance Evaluation and Modeling in Wireless and Wired Networks, PEMWN, 2022, pp. 1–6.

[36] 3GPP, Management and Orchestration; 5G Performance Measurements, Technical Specification (TS) 28.552, 3rd Generation Partnership Project (3GPP), 2022, Version 17.6.0, [Online]. Available: <http://www.3gpp.org/DynaReport/28552.htm>.

[37] 3GPP, Telecommunication Management; Performance Management (PM); Performance Measurements Evolved Universal Terrestrial Radio Access Network (E-UTRAN), Technical Specification (TS) 32.425, 3rd Generation Partnership Project (3GPP), 2021, Version 17.1.0, [Online]. Available: <http://www.3gpp.org/DynaReport/32425.htm>.

[38] O-RAN Working Group 3, O-RAN near-real-time RAN intelligent controller E2 service model 2.00, 2021, ORAN-WG3.E2SM-v02.00 Technical Specification.

[39] O-RAN Working Group 10, O-RAN Operations and maintenance architecture, 2024, ORAN-WG10.OAM-Architecture-R004 Technical Specification.

[40] O-RAN Working Group 6, Cloudification and orchestration use cases and requirements for O-RAN virtualized RAN, 2024, O-RAN-WG6.ORCH-USE-CASES-R003-v11.00 Technical Specification.

[41] P. Hintjens, ZeroMQ: Messaging for Many Applications, O'Reilly Media, 2013.

[42] zhchai, libzmq not support SCTP transport, 2017, Github issue on the ZeroMQ project repository, [Online]. Available: <https://github.com/zeromq/libzmq/issues/2620>.

[43] F. Kaltenberger, T. Melodia, I. Ghauri, M. Polese, R. Knopp, T.T. Nguyen, S. Velumani, D. Villa, L. Bonati, R. Schmidt, et al., Driving innovation in 6G wireless technologies: The OpenAirInterface approach, 2024, arXiv preprint [arXiv:2412.13295](https://arxiv.org/abs/2412.13295).

[44] F. Kaltenberger, A.P. Silva, A. Gosain, L. Wang, T.-T. Nguyen, OpenAirInterface: Democratizing innovation in the 5G era, *Comput. Netw.* 176 (2020) 107284, [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1389128619314410>.

[45] F.A. Bimo, F. Feliana, S.-H. Liao, C.-W. Lin, D.F. Kinsey, J. Li, R. Jana, R. Wright, R.-G. Cheng, OSC community lab: The integration test bed for O-RAN software community, in: 2022 IEEE Future Networks World Forum, FNWF, 2022, pp. 513–518, ISSN: 2770-7679, [Online]. Available: <https://ieeexplore.ieee.org/document/10056724>.

[46] R. Schmidt, M. Irazabal, N. Nikaein, FlexRIC: an SDK for next-generation SD-RANs, in: Proceedings of the 17th International Conference on Emerging Networking EXperiments and Technologies, CoNEXT '21, Association for Computing Machinery, New York, NY, USA, 2021, pp. 411–425, <http://dx.doi.org/10.1145/3485983.3494870>.

[47] R. Mundlamuri, R. Gangula, F. Kaltenberger, R. Knopp, 5G NR positioning with OpenAirInterface: Tools and methodologies, in: 20th Wireless on-Demand Network Systems and Services Conference, WONS, 2025.

[48] R. Gangula, A. Lacava, M. Polese, S. D'Oro, L. Bonati, F. Kaltenberger, P. Johari, T. Melodia, Listen-While-Talking: Toward dApp-based Real-Time Spectrum Sharing in O-RAN, in: MILCOM 2024-2024 IEEE Military Communications Conference, MILCOM, IEEE, 2024, pp. 651–652.

[49] 3GPP, NR; Physical Layer Procedures for Data, Technical Specification (TS) 38.214, 3rd Generation Partnership Project (3GPP), 2024, Version 18.4.0, [Online]. Available: <http://www.3gpp.org/DynaReport/38214.htm>.

[50] A. Bourdoux, A.N. Barreto, B. van Liempd, C. de Lima, D. Dardari, D. Belot, E.-S. Lohan, G. Seco-Granados, H. Sarieddeen, H. Wymeersch, et al., 6G White Paper on Localization and Sensing, 2020, arXiv preprint [arXiv:2006.01779](https://arxiv.org/abs/2006.01779).

[51] A. Nessa, B. Adhikari, F. Hussain, X.N. Fernando, A survey of machine learning for indoor positioning, *IEEE Access* 8 (2020) 214945–214965.

[52] X. Li, K. Pahlavan, Super-resolution TOA estimation with diversity for indoor geolocation, *IEEE Trans. Wirel. Commun.* 3 (1) (2004) 224–234.

[53] R. Gangula, T. Melodia, R. Mundlamuri, F. Kaltenberger, Round trip time estimation utilizing cyclic shift of uplink reference signal, 2024, arXiv preprint [arXiv:2410.04528](https://arxiv.org/abs/2410.04528).

[54] SignalCraft, <https://www.signalcraft.com/products/test-measurement/microwave-systems/sc2430/>.

[55] D. Villa, I. Khan, F. Kaltenberger, N. Hedberg, R.S. da Silva, S. Maxenti, L. Bonati, A. Kelkar, C. Dick, E. Baena, et al., X5G: An open, programmable, multi-vendor, end-to-end, private 5G O-RAN testbed with NVIDIA ARC and OpenAirInterface, 2024, arXiv preprint [arXiv:2406.15935](https://arxiv.org/abs/2406.15935).