

O-RAN: Disrupting the Virtualized RAN Ecosystem

Andres Garcia-Saavedra and Xavier Costa-Pérez

Abstract—The *O-RAN Alliance* is a worldwide effort to reach new levels of openness in next-generation virtualized radio access networks (vRANs). Initially launched by 5 major mobile carriers a couple of years ago, it is nowadays supported by over 160 companies (including 24 mobile operators across 4 continents) representing an outstanding example of how operators and suppliers around the world can constructively collaborate to define novel technical standards. In this paper, we provide a summary of the *O-RAN Alliance* RAN architecture along with its main building blocks. Then, a practical use case exploiting the AI/ML-based innovations enabled by O-RAN is presented, showcasing its disrupting potential. Based on this, the defined interfaces and services are described. Finally, a discussion on the pros and cons of O-RAN is provided along with the conclusions.

I. INTRODUCTION

The virtualization of radio access networks (*a.k.a.* vRAN), with the promise of considerable OPEX/CAPEX savings, high flexibility, and openness to foster innovation and competition, is the last milestone in the NFV revolution and will be a key technology for beyond-5G systems. Harnessing the strengths of NFV into the radio access arena, however, entails a number of challenges that are the object of study by multiple initiatives such as Rakuten’s greenfield deployment in Japan, Cisco’s Open vRAN Ecosystem, Facebook Telecom Infra Project’s vRAN Fronthaul Project Group and the O-RAN Alliance. Arguably, among these efforts, O-RAN is the one with most traction.

In this paper, we provide an overview of the *O-RAN Alliance* specifications to date and their capabilities. O-RAN is a major carrier-led effort to define the next generation (v)RANs for *multi-vendor* deployments. It is aimed at disrupting the vRAN ecosystem by breaking *vendors lock-in* and opening up a market that has been traditionally dominated by a small set of players. If successful, O-RAN might unleash an unprecedented level of innovation in the RAN space by lowering the market entrance barrier to new competitors.

In the following, we start off by summarizing the architecture of O-RAN and its main building blocks in §II and different deployment models in §III. In §IV we present a new use case for O-RAN concerning the joint orchestration of radio and cloud resources. We then introduce the key interfaces between the building blocks of O-RAN and available services in §V, leveraging on our use case as an illustrative example for such services. Finally, we close the paper with a discussion and the conclusions of the paper in §VI and VII.

II. O-RAN ARCHITECTURE

Fig. 1 depicts a high-level view of the O-RAN architecture [1]. Doubtlessly, the most important functional compo-

TABLE I
ABBREVIATIONS

Abbreviation	Details
AI/ML	Artificial Intelligence/Machine Learning
CU	Central Unit
CRUD	Create, Read, Update, and Delete
DU	Distributed Unit
eNB/gNB	Evolved/next-generation Node B (base station)
FCAPS	Fault, Configuration, Accounting, Performance and Security
ME	Managed Element
MF	Managed Function
NFV	Network Function Virtualization
OPEX/CAPEX	Operating Expenditures/Capital Expenditures
PNF	Physical Network Function
rApps	Non-RT RIC Applications
RIC	RAN Intelligent Controller
RT	Real-time
RU	Radio Unit
SMO	Service Management and Orchestration
VNF	Virtual Network Function
vRAN	Virtual Radio Access Network
xApps	Near-RT RIC Applications

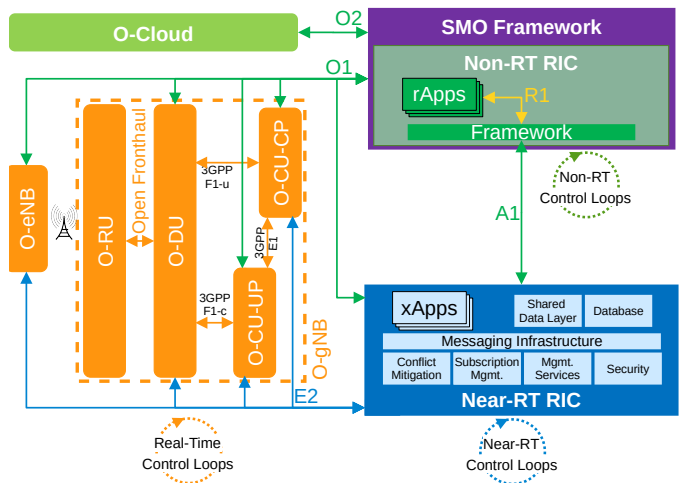


Fig. 1. O-RAN architecture [1]

nents introduced by O-RAN are the Non-Real-Time (Non-RT) Radio Intelligent Controller (RIC) and the near-RT RIC. While the former is hosted by the Service Management and Orchestration (SMO) framework of the system (e.g., integrated within ONAP), the latter may be co-located with 3GPP gNB functions, namely, O-RAN-compliant Central Unit (O-CU) and/or Distributed Unit (O-DU), or fully decoupled from them as long as latency constraints are respected. We discuss different deployment flavours in §III. Fig. 1 also depicts the O-Cloud, an O-RAN compliant cloud platform that uses

hardware acceleration add-ons when needed (e.g., to speed up fast fourier transform operations or forward error correction tasks) and a software stack that is decoupled from the hardware to deploy eNBs/gNBs as virtualized network functions (VNFs) in vRAN scenarios. In the following, we detail the jurisdiction and roles of each functional component defined above.

Service Management and Orchestration (SMO). The SMO consolidates several orchestration and management services, which may go beyond pure RAN management such as 3GPP (NG-)core management or end-to-end network slice management. In the context of O-RAN, the main responsibilities of SMO are (i) fault, configuration, accounting, performance and security (FCAPS) interface to O-RAN network functions; (ii) large timescale RAN optimization; and (iii) O-Cloud management and orchestration via O2 interface, including resource discovery, scaling, FCAPS, software management, create, read, update, and delete (CRUD) O-Cloud resources.

Non-RT RAN Intelligent Controller (Non-RT RIC). As mentioned earlier, this logical function resides within the SMO and provides the A1 interface to the Near-RT RIC. Its main goal is to support large timescale RAN optimization (seconds or minutes), including policy computation, ML model management (e.g. training), and other radio resource management functions within this timescale. Data management tasks requested by the Non-RT RIC should be converted into the O1/O2 interface; and contextual/enrichment information can be provided to the near-RT RIC via A1 interface.

Near-RT RAN intelligent Controller (Near-RT RIC). Near-RT RIC is a logical function that enables near-real-time optimization & control and data monitoring of O-CU and O-DU nodes in near-real-time timescales (between 10ms and 1s). To this end, Near-RT RIC control is steered by the policies and assisted by models computed/trained by the Non-RT RIC. One of the main operations assigned to the near-RT RIC is radio resource management (RRM) but near-RT RIC also supports 3rd party applications (so-called xApps).

This architecture inherently enables three independent—yet with sporadic interactions—control loops:

- 1) **Non-RT RIC control loop:** Large timescale operation of the orders of seconds or minutes. The goal is to perform O-RAN specific orchestration decisions such as policy configuration or ML model training.
- 2) **Near-RT RIC control loop:** Sub-second time scale operation with the goal of performing tasks such as policy enforcement or radio resource management operations.
- 3) **O-DU Scheduler control loop:** Real-time operation performing legacy radio operations such as HARQ, beamforming or scheduling—out of O-RAN’s scope.

III. SCENARIOS AND DEPLOYMENT OPTIONS

O-RAN’s disposition towards software-defined AI-assisted RAN control fosters different degrees of *openness*, namely, systems comprised of (i) O-RAN-compliant PNFs exposing and using O-RAN interfaces so different vendors can interplay (lowest degree of openness); (ii) chassis of servers and racks in a cloud shared among multiple vendors (higher degree of

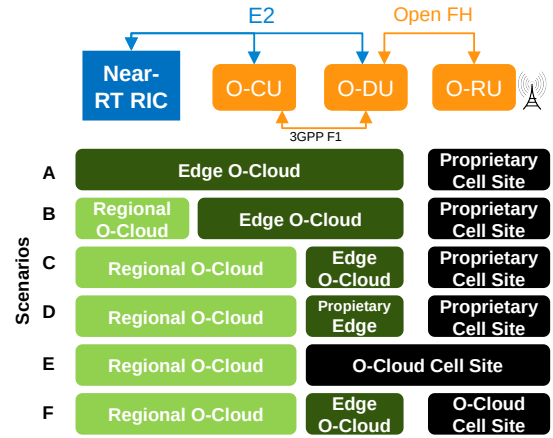


Fig. 2. Deployment scenarios [2]

openness); and (iii) one or multiple O-Clouds, a fabric of COTS servers (including FPGA or GPU accelerators) and networking infrastructure hosting O-RAN software that is decoupled from the hardware at different layers: hardware (e.g. ETSI NFVI hardware sub-layer), middle (e.g., ETSI NFVI virtualization sub-layer + VIM) and a top layer hosting virtual RAN functions (highest degree of openness).

Such openness enables substantial flexibility to deploy each of the logical functions introduced earlier, e.g., O-DU and O-RU can be co-located or not depending on the context and particular needs of the operator *and these decisions may be changed over time with minimal cost* [3]. Fig. 2 summarizes six different deployment scenarios described in the following: **Scenario A.** In this scenario, one edge cloud centralizes all near-RT RIC, virtual O-CU and O-DU functions to support very dense deployments (e.g., dense urban areas) that provide a high-capacity fronthaul network. This type of deployment expects edge clouds with substantial hardware acceleration capabilities.

Scenario B. This scenario separates the virtual O-CU and O-DU functions from the near-RT RIC, which can be placed in a regional cloud and uses E2 interface for interaction with O-CUs and O-DUs. This allows near-RT RIC to have a global view for optimization.

Scenario C. Virtual O-CU network functions are co-located with the near-RT RIC in a regional cloud. Regional cloud and edge cloud(s) must, in this case, satisfy the latency requirements of 3GPP-defined F1 interface [3]. This scenario enables deployment in locations with limited fronthaul capacity and number of O-RUs. There are two additional variations of this scenario, C.1 and C.2 to support specific network slices needs.

Scenario D. This scenario is a replica of Scenario C in which O-DU functions are *not* virtualized in an O-Cloud but rather supported by an O-RAN-capable PNF.

Scenario E. This scenario is a replica of Scenario C in which the O-RU functions are virtualized into a common O-Cloud, in addition to the O-DU functions.

Scenario F. This scenario is a replica of Scenario E in which O-DU and O-RU functions are virtualized into separate O-Clouds.

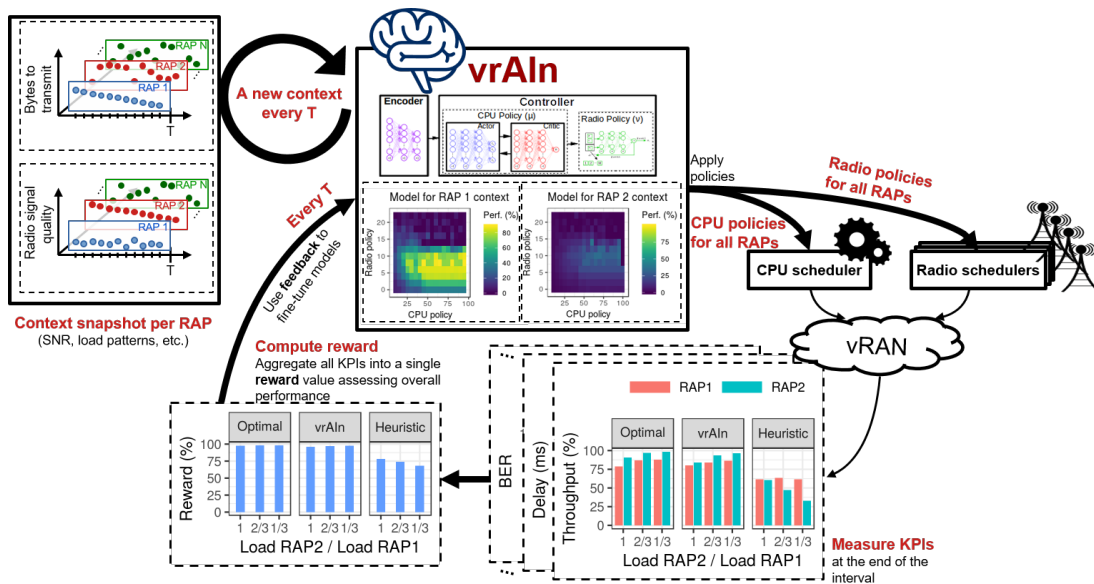


Fig. 3. AI-aided approach to joint computing/radio resource orchestration [4]

IV. JOINT ORCHESTRATION OF COMPUTING AND RADIO RESOURCES IN vRANS: AN O-RAN USE CASE

Despite the potential benefits of RAN virtualization—see a discussion in §VI—dynamic resource orchestration becomes more compounded. Specifically, the problem of optimally allocating *computing resources* and *radio resources* is now coupled and requires of joint management. This is demonstrated in several works such as [4]. Moreover, in these scenarios, the virtualized base stations (vBSs)¹ share a pool of computing resources and may or may not share radio spectrum as in [5], which further complicates the orchestration problem. In this context, the objective is twofold:

- 1) When the computing capacity is over-dimensioned, the goal shall be to minimize the allocation of computing resources in order to save operational costs.
- 2) When the computing capacity is under-dimensioned (to attain capital cost savings), the goal shall be to maximize performance, mitigating the amount of decoding errors due to deficit of computing resources.

The authors of [4] illustrate a strong coupling between computing and radio resource allocation policies. Different computing and radio resource control policies may be derived and supported by O-RAN. For instance:

• Computing control policies

- 1) Policy 1: A fraction of overall computing time is reserved for each vBS, while some computing time is left unallocated to save costs. This can be implemented using, for instance, Docker’s API for containerized O-CUs/O-DUs as in [4], and can be applied to general purpose CPUs and/or shared accelerators for specific tasks, e.g. forward error control.

¹We will use the term virtualized base station to refer to any radio stack in the edge cloud, i.e., O-DUs, O-CU+O-DUs or O-eNBs.

- 2) Policy 2: A subset of computing units (CPUs, accelerators) reserved for each vBS. This can be applied in conjunction with Policy 1 (multiplexing computing units).

• Radio control policies

- 1) Policy 1: An upper-bound eligible modulation and coding scheme (MCS) index for each vBS as in [4]. In this way, a vBS cannot select higher MCS indexes than this bound, which helps to constrain the computational demand of the vBS.
- 2) Policy 2: A fraction of the overall subcarriers or physical resource blocks per transmission time interval, as in [5]. This is required when multiple instances of a vBS share the same carrier bandwidth.

The above joint optimization may be performed with the aid of artificial intelligence and machine learning (AI/ML) models. An example of such a model is vrAIn [4]. vrAIn builds upon a contextual bandit (CB) formulation, which is a particular case of reinforcement learning (RL). In CB problems, one observes a *context* vector, chooses an action and receives a reward signal as feedback, sequentially at different time stages. The goal is to find a model that maps input contexts into compute/radio control policies or actions that maximize the expected reward:

- **State or context space.** At each stage, T context samples are collected. Each sample consists of (i) the buffer size, (ii) the mean signal-to-noise (SNR) and (iii) the variance SNR, measured for all users across all vBSs.
- **Action space.** Comprises all pairs of compute and radio control policies/decisions defined earlier.
- **Rewards.** The design of a reward function depends on the system’s goal. In [4], a two-fold objective is considered: (i) minimizing operational costs due to CPU reservations, and (ii) maximizing performance by reducing decoding error rates and latency.

Fig. 3 illustrates the decision-making closed-loop process implementing the RL formulation above. In more detail, the orchestrator consists of a construct of neural networks as follows:

- **Encoder:** A series of sparse autoencoders (SAEs) reduces the dimensionality of the *input* context samples without compromising expressiveness;
- **CPU policy:** An actor-critic neural network structure receives an encoded context as input and implements a deep deterministic policy gradient (DDPG) algorithm to compute an appropriate CPU control policy for each vBS;
- **Radio policy:** A deep classifier receives an encoded context *and* the current CPU control policy as input to derive the most appropriate radio control policy.

More details about this model can be found in [4].

AI-aided vRAN resource orchestration technologies, such as vrAI, are finally enabled in practice by O-RAN. In the following, we will use this use case as *an example* to illustrate the operation of the different interfaces and services available in the architecture of O-RAN. The interested reader may find the analysis of more use cases in [6].

V. SERVICES AND INTERFACES

In this section, we introduce the most relevant services and interfaces provided by O-RAN. We note that currently there is no public specification of an interface between the SMO and the Non-RT RIC, which is left for manufacturers to make their own design choices.

A. O1 Services: Operation, Administration and Management

O1 is in fact a logical interface to perform management operations with different deployment models:

- 1) **Flat management model:** All the entities subject of management in the architecture except O-Cloud (which is managed through O2), i.e., O-eNB, O-CU-CP/UP, O-DU and O-RU, *a.k.a* managed functions (MFs), are also managed elements (MEs) by the SMO through O1.
- 2) **Hierarchical management model:** This model allows some MFs to manage lower-level MEs, e.g., O-DU may manage O-RU through Open Fronthaul M-Plane interface.
- 3) **Hybrid management model:** In this model, the management responsibility is shared between the O-DU (through Open Fronthaul M-Plane interface) and the SMO (through O1 interface).

O-RAN's OAM architecture specification [7] illustrates different deployment examples of O1. In this way the SMO can provide a series of management services, including FCAPS, file management or software management. In the case of VNFs, the interface supports orchestration and monitoring of the infrastructure resources. In more detail, [7] specifies the following list of services:

Provisioning Management Service. This service allows a consumer to configure attributes of managed objects.

Fault Supervision Management Service. This service allows

reporting errors and event to a Fault Supervision consumer to perform fault supervision operations such as alarm handling.

Performance Assurance Management Service. This service allows to transfer bulk and/or real-time streaming performance data. Its consumer may perform performance assurance operations such as selecting the measurements to be reported and their frequency.

Trace Management Service. This service allows asynchronous streaming of trace data upon triggering event.

File Management Service. This service allows transferring files between a provider element and a consumer.

Heartbeat Management Service. This service allows a provider to send heartbeats to a consumer.

PNF Startup and Registration Management Service. This service allows acquiring network layer parameters of physical PNFs and change its operational state.

PNF Software Management Service. This service allows downloading, installing, validating and activating new software packages into physical PNFs in addition to obtaining software versions from PNFs.

B. Non-RT RIC: rApps, and AI services

The Non-RT RIC is comprised of two functions: (*i*) the Non-RT RIC framework, which terminates the A1 interface and exposes services to so-called Non-RT RIC applications (rApps) through R1 interface. rApps are modular applications in charge of providing added-value services relative to the operation of the RAN, such as driving the A1 interface, enforcing policies through O1/O2 interface, or generating enrichment information for other rApps. In turn, R1 is an interface internal to the Non-RT RIC connecting rApps and Non-RT RIC framework. It is a collection of services, such as service registration and discovery services, AI/ML workflow services, and A1-related services. In the context of our illustrative use case, presented in §IV, the actor-critic neural network structure giving light the CPU policy and the deep classifier implementing vrAI's radio policy is implemented as two rApps, as shown in Fig. 5. The radio policy uses information from the CPU policy, which is communicated via R1 interface and the Non-RT RIC framework.

In turn, A1 is a logical interface that connects the non-RT RIC with the near-RT RIC [8]. The main goal of this interface is to enable non-RT RIC to provide policy-based guidance, AI/ML model management and enrichment information to the near-RT RIC for the optimization of certain RAN functions. Moreover, A1 can provide basic feedback from near-RT RIC to allow the non-RT RIC monitor the use of policies. To this end, A1 provides essentially three services:

- 1) **Policy Management Service:** Declarative policies based on A1 policy feedback and network status provided over O1 interface. O-RAN uses a consumer/producer model where Non-RT RIC hosts the A1 policy (A1-P) consumer and the A1-P producer resides within the near-RT RIC. The A1-P producer cannot modify or delete a policy. Examples of policy statements specified by O-RAN are policy objectives: QoS, QoE, KPI, KQI




Fig. 4. O-Cloud at the edge serving shared computing resources to multiple O-DUs.

targets; and policy resources: traffic steering preferences and system efficiency. The specification of policy management functions (create, query, update, delete and feedback subscription) can be found in [8]. In our use case, this service is employed to communicate the aforementioned radio policy to the Near-RT RIC.

- 2) **ML Model Management Service:** AI/ML is an integral part to O-RAN. O-RAN specifies different AI/ML scenarios where A1 may be involved. Given the important role of AI/ML in O-RAN, we provide extended details later in §V-F.
- 3) **Enrichment Information Service:** External information that may be exposed to the Near-RT RIC internal functions or applications (e.g. context information for ML models) that is not directly reachable to the near-RT RIC from network function data.

C. O2 Services: Cloudification and orchestration

The O-Cloud pools computing resources including general-purpose CPUs and *shared* task accelerators (based on GPUs, FPGAs or ASIC) for fast fourier transform tasks or forward error coding. This computing resources are brokered by an abstraction layer² (see Fig. 4). O-RAN provides a cloud reference design in [9].

The O2 interface corresponds to a collection of services and associated interfaces between the O-Cloud and the SMO. Specifically, O-RAN organizes these services into two logical groups:

- 1) **Infrastructure Management Services:** A subset of O2 functions that are responsible for deploying and managing cloud infrastructure.
- 2) **Deployment Management Services:** A subset of O2 functions that are responsible for managing the life-cycle of virtualized/containerized deployments on cloud infrastructure.

In the context of our our case, presented in §IV, the O2 interface is used by the CPU policy in the Non-RT RIC to enforce CPU policies, as shown in Fig. 5.

D. Near-RT RIC and E2

E2 nodes are logical functions that support all the protocol layers and interfaces defined by 3GPP RAN (eNB for E-UTRAN and gNB/ng-eNB for NG-RAN). One Near-RT RIC may be connected through transport functions to one or

multiple E2 nodes, though each E2 node may be connected to a single Near-RT RIC. The near-RT RIC uses the A1 interface to receive policies, enrichment data or ML models from the Non-RT RIC, and E2 interface to collect near-real-time information from E2 nodes and carry out fine-grained radio resource management (RRM) actions over E2 nodes. The architecture of the near-RT RIC is shown in Fig. 1 and its key functions described as follows:

xApps. These are 3rd party applications, which can be implemented by multiple microservices. The Near-RT RIC hosts one or more xApps that also use A1 and E2 interface to provide value-added services and enhance the RRM capabilities of the near-RT RIC.

A database. Stores data from xApp applications and (near real-time) data from E2 nodes and provides data to xApp applications.

Interface termination. For O1, A1 and E2 interfaces.

xApps subscription management. Consolidates all subscriptions and data distribution operations into an unified functional block.

Conflict mitigation. Resolves conflicting interactions (e.g. requests) from different xApps.

Security. Prevents hazard to the near-RT RIC from 3rd party xApps such as exporting unauthorized data or abusing radio resource allocations. The description of concrete security functions are not defined yet.

Management Services. Including fault management, configuration, accounting and performance management (FCAPS), including such as collection of logs, traces and metrics; life-cycle management for xApps, including onboarding, deployment, resource management and termination.

Messaging infrastructure. A common message distribution system for different elements within the near-RT RIC.

Following ETSI NFV directions, an xApp consists of an xApp descriptor and its image. The xApp descriptor provides xApp management services including the necessary information for life-cycle management, health management and FCAPS. Note, importantly, that, although xApps may belong to 3rd parties, they shall expose an open API for A1, O1 and E2 termination, control and shared data management.

The protocols over E2 interface are based upon control-plane protocols, defined in [10]. O-RAN specifies two types of procedures over E2: Functional and Global. Information elements (IEs) may be used to incorporate information in control messages. O-RAN specifies different IEs including cause IE, global RIC ID IE, global E2 node ID IE or RIC control IE, among others—see [10] for details.

To integrate our use case, an xApp implements the context encoder, which encodes contextual data collected from the O-DU via the E2 interface and stored in the Near-RT RIC's database. An additional xApp forwards radio policies to the O-DU according to the Non-RT RIC's radio policy received via the A1-P service. This is illustrated in Fig. 5.

²See, e.g., bbdev (https://doc.dpdk.org/guides/prog_guide/bbdev.html)

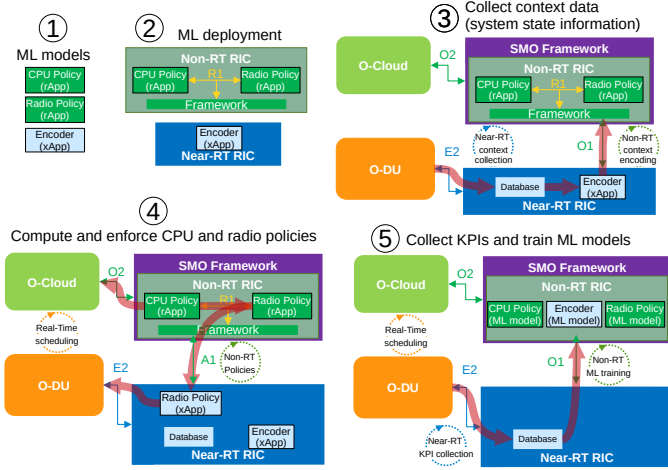


Fig. 5. Integration of an AI-aided vRAN resource orchestrator into O-RAN architecture.

E. Open Fronthaul

The choice of a functional split for next-generation RANs have attracted substantial research activity in the last few years [11], [3] as there exists an inherent trade-off between keeping the O-RU as simple as possible to reduce costs—centralizing functions in CU—and distributing functions towards the RU to alleviate congestion on the fronthaul network. O-RAN has selected “7-2x” functional split, following 3GPP nomenclature, though O-RAN is flexible to allow the precoding function be located on either side.

O-RAN’s open fronthaul is a logical interface consisting of Lower-Layer Split (LLS) control plane (LLS-CP), LLS user plane (LLS-UP) and synchronization plane [12], management plane (M-plane) [13], in addition to specifying a new cooperative transport interface (CTI). CTI is intended to support real-time and non-real-time cooperation between the eNB/gNB and the resource allocation based transport network. In case the transport network (fronthaul) consists of a point-to-point link (e.g. optical fiber) between each O-RU and the corresponding O-DU, CTI is not required because transport resources are not shared. However, when the transport network consists of a packet-based system interconnecting multiple O-DUs to multiple O-RUs, CTI is used to identify each fronthaul flow and trigger appropriate scheduling decisions by the transport nodes so latency, bandwidth and jitter requirements are met across all flows.

F. Artificial Intelligence and Machine Learning Services

AI/ML is a cornerstone in the design of the O-RAN architecture. The goal is to exploit AI/ML models to carry out tasks that have traditionally been done quasi-statically by human operators in the past or are overly complex tasks that never made the transfer from academia into production systems. These include tasks such as, for instance, zero-touch and automated resource control tasks, anomaly detection or traffic classification.

The use of AI/ML models for next-generation RANs is paramount in the design of O-RAN’s architecture. Regard-

ing AI/ML, O-RAN follows the general principles described in [14]. O-RAN defines an *ML training host* as the entity (network function) that builds the ML model and performs its training offline. Similarly, an *ML inference host* corresponds to the network function that executes the ML model and/or performs online training. A ML model will usually be part of a larger decision-making solution, i.e., an ML-assisted solution, which is in turn hosted by *the actor*, the ultimate responsible to perform decisions or *actions*. These actions may be of different nature, including configuration management (CM) changes over O1 interface, policy management over A1 interface, or O-eNB (O-CU/O-DU/O-RU) control/policy parameters over E2 interface, depending on the deployment flavour of the ML hosts.

Three deployment scenarios are considered:

- 1) Non-RT RIC takes up both roles of ML training and inference host. In this case, the process of building the ML model, its life-cycle management and data provisioning is handled internally within the SMO. Two types of actions are considered in this case:
 - a) A policy for the near-RT RIC, which is transferred through the policy service of A1 interface.
 - b) An O-CU/O-DU/O-RU configuration parameter, which is enforced using O1 interface.
- 2) Non-RT RIC takes the role of ML training host, while the near-RT RIC acts as ML inference host. In this case, both O1 and O2 interfaces are used for creating and maintaining the model. The nature of the action may be twofold:
 - The near-RT RIC itself, e.g., forecasting information for internal mechanisms, where A1’s enrichment data service is used for data provisioning between non and near-RT RICs.
 - An O-CU/O-DU/O-RU configuration parameter, where E2 is used for both data collection and enforcement of control or policy parameters.
- 3) Non-RT RIC acts as the ML training host and the O-CU or O-DU takes the role of ML inference host.

Regardless the deployment option and the type of AI/ML algorithm (supervised, unsupervised or reinforcement learning), there are a series of key steps that are relevant:

ML model capability query/discovery. Whenever an ML-assisted solution requires to build an AI/ML model, the SMO shall discover some capabilities in the ML inference host, namely: HW processing capabilities (e.g., CPU/GPU resources, memory, etc.), supported ML models and engines (e.g., JSON, protobuf, etc.), NFVI-based architecture support, and data sources available.

ML model selection and training. The ML model designer needs to make a series of choices, e.g., exploration-vs-exploitation intervals in reinforcement learning, format of the input and out data, etc.

ML model deployment and inference. Models may be deployed via containerized images into the inference host.

ML model performance monitoring. Explicit feedback of

the performance of the ML model (e.g., for training in reinforcement learning mechanisms).

ML model update. Online model updates (e.g., online training) or major model updates done by the ML designer.

In the case of our AI-assisted resource orchestrator case, we have 3 machine learning models, CPU policy, radio policy and context encoder (see Fig. 5 (1) and (2)). On the one hand, both the CPU actor-critic implementing the CPU policy and the deep classifier implementing the radio policy are hosted by two respective rApps, which communicate through R1 interface, i.e., the Non-RT RIC acts as their inference host and the resulting policies are enforced via O2 (CPU) and A1 (radio) interfaces (see Fig. 5 (3)). On the other hand, the autoencoders implementing our context encoder are deployed into an xApp hosted by the Near-RT RIC, which acts as ML inference host (see Fig. 5 (4)). During training, which could be done in pre-production (offline), all ML models are trained by the non-RT RIC as shown by Fig. 5 (5), with data stored in the near-RT RIC's database.

VI. DISCUSSION

State-of-the-art vRAN solutions applied today in the market, which rely on *dedicated* hardware acceleration, jeopardize the very reasons that make virtualization appealing for the RAN in the first place: flexibility and cost-efficiency. First, research has shown that *cloud RANs require many more resources than legacy RAN platforms* to attain similar performance guarantees in real mobile networks. Second, *dedicated accelerators make vDUs more expensive and power-hungry than their legacy counterparts*—let alone the fact that the much-longed hardware/software decoupling is not achieved.

O-RAN's O-Cloud approach strives to address the above issues: while hardware acceleration is still required for specialized, compute-intensive and repetitive tasks, such as fast Fourier transform and forward error coding, O-RAN's approach is to provide pools of shared accelerators, brokered by an abstraction layer as shown in Fig. 4. The goal is to preserve the carrier-grade performance that only hardware accelerators can provide without sacrificing the flexibility and cost-efficiency of RAN virtualization.

O-RAN does not only target virtualized RAN scenarios, but open RAN deployments overall to enable competition in the RAN, traditionally monopolized by a small set of manufacturers. This should accelerate innovation and help reduce costs. However, according to recent market forecasts [15], Open RAN is expected to cover only about 10% of the overall market by 2025. Thus, despite the new business opportunities opened to small and medium-scale vendors (traditionally alien to large-scale RAN deployments) significant hurdles will need to be overcome to reach the economies of scale of major vendors in the RAN ecosystem in order to be competitive.

VII. CONCLUSIONS

The *O-RAN Alliance* is a major carrier-led effort aimed at disrupting the next generation virtualized RAN (vRAN)

ecosystem and unleash an unprecedented level of innovation. Its large carrier and vendor support by more than 160 companies has given it an exceptional momentum, producing over 40 technical specification documents within 2 years and 1.3 million lines of open-source code. In this paper we summarized the main content of the O-RAN specifications available focusing on the proposed architecture and building blocks. To illustrate the innovations enabled by O-RAN, we used an state-of-the-art AI-aided orchestrator that jointly manages radio and computing control policies in vRANs, named vrAIn. Finally, a discussion on O-RAN pros and cons was provided, summarizing its disrupting potential together with major technical and market challenges ahead.

REFERENCES

- [1] O-RAN Alliance, "O-RAN-WG1-O-RAN Architecture Description - v04.00.00." Technical Specification, Mar. 2021.
- [2] —, "Cloud Architecture and Deployment Scenarios v02.01 (O-RAN.WG6.CAD-v02.01)." Technical Report, Jul. 2020.
- [3] A. Garcia-Saavedra, J. X. Salvat, X. Li, and X. Costa-Perez, "Wizhaul: On the centralization degree of cloud ran next generation fronthaul," *IEEE Transactions on Mobile Computing*, vol. 17, no. 10, pp. 2452–2466, 2018.
- [4] J. A. Ayala-Romero *et al.*, "vrAIn: A Deep Learning Approach Tailoring Computing and Radio Resources in Virtualized RANs," in *The 25th Annual International Conference on Mobile Computing and Networking*, 2019, pp. 1–16.
- [5] J. Mendes *et al.*, "Cellular access multi-tenancy through small-cell virtualization and common rf front-end sharing," *Computer Communications*, vol. 133, pp. 59–66, 2019.
- [6] O-RAN Alliance, "O-RAN Use Cases and Deployment Scenarios," White Paper, Feb. 2020.
- [7] —, "O-RAN Working Group 1. O-RAN Operations and Maintenance Interface Specification (O-RAN.WG1.OI-Interface.0-v04.00)," Technical Specification, Nov. 2020.
- [8] —, "O-RAN Working Group 2. Non-RT RIC: Functional Architecture (O-RAN.WG2.Non-RT-RIC-ARCH-TR-v01.01)," Technical Report, Nov. 2021.
- [9] —, "Cloud Platform Reference Designs (O-RAN.WG6.CLOUD-v02.00)," Technical Specification, Nov. 2020.
- [10] —, "O-RAN Working Group 3, Near-Real-time RAN Intelligent Controller, E2 Application Protocol (E2AP) (ORAN-WG3.E2AP-v02.00.00)," Technical Specification, Mar. 2021.
- [11] A. Garcia-Saavedra, G. Iosifidis, X. Costa-Perez, and D. J. Leith, "Joint optimization of edge computing architectures and radio access networks," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 11, pp. 2433–2443, 2018.
- [12] O-RAN Alliance, "O-RAN Fronthaul Working Group. Control, User and Synchronization Plane Specification (O-RAN-WG4-CUS.0 - v03.00)," Technical Specification, Apr. 2020.
- [13] —, "O-RAN Fronthaul Working Group. Management Plane Specification (O-RAN-WG4-MP.0 - v03.00)," Technical Specification, Apr. 2020.
- [14] —, "O-RAN Working Group 2. AI/ML workflow description and requirements (O-RAN.WG2.AI/ML-v01.02)," Technical Report, 2021.
- [15] Dell'Oro Group, "Open RAN Market Expected to Approach \$10 B, According to Dell'Oro Group," Online: <https://www.delloro.com/news/open-ran-market-expected-to-approach-10-b/>, Feb. 2021.