1

Introduction to O-RAN

Cellular technologies are becoming increasingly complex. For example, fifth generation (5G) is significantly more complex than fourth-generation (4G) cellular technology Long-Term Evolution (LTE). Open Radio Access Network (O-RAN), also referred to as Open RAN, is a revolutionary technology that is transforming the cellular landscape. Unlike traditional cellular networks, O-RAN is an open, programmable, interoperable, and virtualized Radio Access Network (RAN) architecture that allows network operators to mix and match hardware and software components from different vendors. O-RAN provides increased flexibility, scalability, and cost savings for network operators. Furthermore, O-RAN leads to enhanced competition among vendors, which ultimately benefits consumers and enterprises.

This chapter first creates a foundation for O-RAN by illustrating evolution of cellular technologies, reviewing components of a cellular system, and describing evolution of RAN architectures. Then, a brief introduction to the O-RAN ALLIANCE is given, followed by illustration of a simplified O-RAN architecture defined by the O-RAN ALLIANCE. Deployments of O-RAN require the use of standards and software frameworks developed by organizations other than the O-RAN ALLIANCE in addition to the specifications developed by the O-RAN ALLIANCE. Roles of such non-O-RAN ALLIANCE organizations in the O-RAN ecosystem are discussed. Driving forces behind O-RAN's rising importance are explained. Numerous use cases of O-RAN that help optimize various RAN operations are described. Finally, the overall accomplishments of O-RAN are summarized along with O-RAN's technical priorities and potential future work.

1.1 Evolution of Cellular Technologies

O-RAN is a comprehensive framework for implementing the radio access network portion of a cellular network. The O-RAN specifications are developed by the O-RAN ALLIANCE and support 4G and 5G cellular communications. As the cellular communication technologies evolve, O-RAN specifications can be augmented to support evolved cellular technologies.

Figure 1.1 shows the evolutionary path of cellular technologies from the first generation (1G) to the sixth generation (6G).

The concept of cellular communications was developed at Bell Labs in the 1970s. The 1G cellular systems were analog. Diverse types of 1G cellular systems were deployed around the world. Examples of such 1G systems included Advanced Mobile Phone System (AMPS) in the United States and NMT 500 and European Total Access Communication System (ETACS) in Europe. The 1G systems provided voice services and typically used frequency-division multiple access on the radio interface between the user device and the RAN. The analog 1G systems were replaced by

Fundamentals of O-RAN, First Edition. Nishith D. Tripathi and Vijay K. Shah.

Companion website: www.wiley.com/go/tripathioran

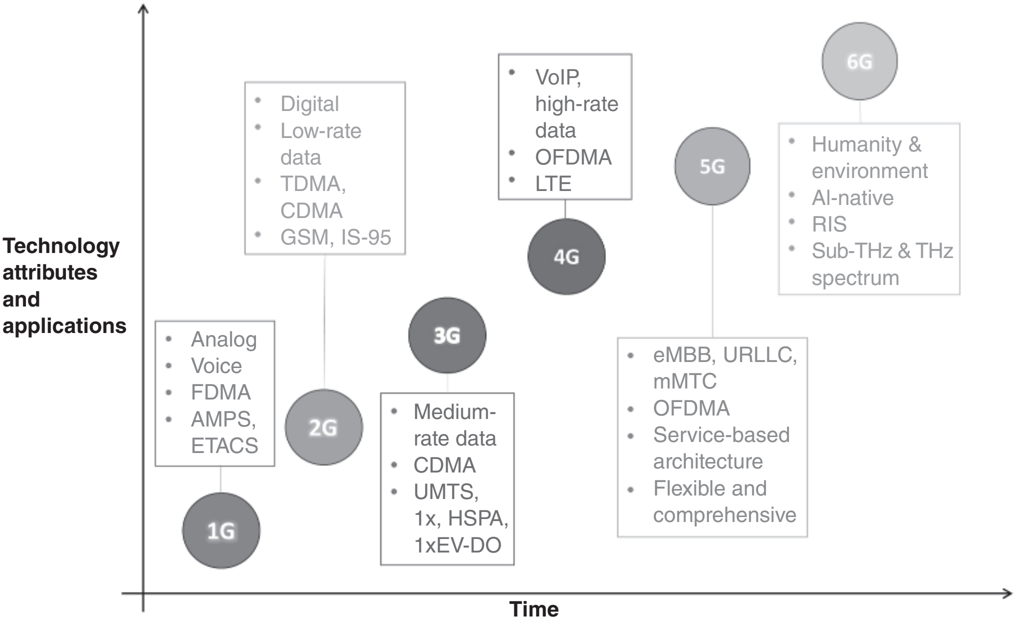

Figure 1.1 Evolutionary Path of Cellular Technologies.

digital second-generation (2G) systems. Examples of 2G systems include Global System for Mobile Communication (GSM) and Interim Standard-95 (IS-95). GSM utilized time-division multiple access (TDMA) and IS-95 used code-division multiple access (CDMA) on the radio interface. The 2G systems provided voice and low-rate data services. The 2G GSM systems typically transitioned to 2.5G General Packet Radio Service (GPRS) and Enhanced Data rates for GSM Evolution (EDGE) and then to third-generation (3G) Universal Mobile Telecommunication System (UMTS) and UMTS enhancements such as High-Speed Downlink Packet Access and High-Speed Uplink Packet Access. 2G IS-95 systems evolved to 3G CDMA200 and 1xEV-DO systems. While 2.5G systems provided higher data rates such as peak data rates of few hundred kilobits per second, 3G systems provided peak data rates of few megabits per seconds. The 3G cellular systems use two types of core networks, a circuit-switched core network to support circuit-switched voice services and a packet-switched core network to support Internet Protocol (IP)-based services such as web browsing, video streaming, and e-mail. The competing 3G technologies of CDMA2000 and UMTS evolved to the common 4G LTE.

4G LTE utilizes orthogonal frequency-division multiple access (OFDMA) on the radio interface. Furthermore, LTE utilizes a single type of core network, a packet-switched core network called Evolved Packet Core (EPC). Since there is no circuit-switched core network in LTE, voice calls are supported using Voice over IP (VoIP), widely known as Voice over LTE (VoLTE). The use of OFDMA and wide radio channel bandwidths enable LTE to support much higher data rates such as tens of megabits per seconds.

5G is a technology with the potential to transform many industries and segments of the economy. While 5G continues to use OFDMA on the radio interface, it significantly increases the flexibility such as large channel bandwidths, variable subcarrier spacing, support for millimeter wave spectrum, and cartier bandwidth parts. A high-performance and flexible radio interface enables 5G to support a much broader set of use cases. These use cases may belong to one or more ITU usage categories of enhanced mobile broadband (eMBB), ultra-reliable and low-latency communication (URLLC), and massive Internet of Things (mIoT). 5G introduces network slicing, where different logical networks are created to meet diverse service and customer requirements using the same physical infrastructure. 5G also introduces a Service-Based Architecture (SBA) that makes use of virtualization technologies for efficiency, scalability, and flexibility.

6G is the first cellular technology that emphasizes humanity and environment right from the requirements phase instead of focusing primarily on user experience and network performance like previous generations of cellular technologies (NGA 2022). 6G is expected to be even more transformational than 5G. 6G is likely to introduce new technologies such as native artificial intelligence (AI), Reflective Intelligent Surface (RIS), and higher-frequency spectrum, such as sub-THz spectrum and THz spectrum.

1.2 Components of Cellular Systems

While wireless communication is one of the most prominent features of a cellular system, the overall cellular system includes both wireless and wired links. Figure 1.2 illustrates a high-level architecture of a 5G system (Tripathi and Reed 2019).

A 5G system consists of the user equipment (UE), the 5G access network (AN), and the 5G core network (5GC or 5GCN) (3GPP TS23.501 2024). Different types of 5G UEs are supported including smartphones, Internet of Things (IoT) devices, smart watches, and augmented reality/virtual reality headsets. The 5G access network may be a Third Generation Partnership Project (3GPP)-defined RAN such as next-generation RAN (NG-RAN) or a non-3GPP access network. The NG-RAN may have next-generation-Node Bs (gNBs) or next-generation-evolved Node Bs (ng-eNBs). The gNB communicates with 5G UEs using a New Radio (NR) air interface, while the ng-eNB communicates with UEs using the 4G LTE air interface. Both the gNB and the ng-eNB connect to the 5GC, which is also referred to as the next-generation core (NGC). *A 5G RAN can be defined as the RAN that utilizes the 5G NR air interface for communications with UEs*. Hence, 5G RAN is a subset of NG-RAN, because NG-RAN can encompass both gNBs and ng-eNBs. An example of a non-2GPP access network is a Wi-Fi access network. This book primarily focuses on the 5G NR air interface between the 5G UE and the gNB while discussing the O-RAN.

A 5G UE exchanges Access Stratum (AS) signaling such as Radio Resource Control (RRC) signaling with the gNB in support of various RAN operations such as RRC connection setup, bearer setup for user traffic exchange, and handover. A 5G UE exchanges Non-Access Stratum (NAS) signaling with the NGC in support of operations such as registration, mutual authentication,

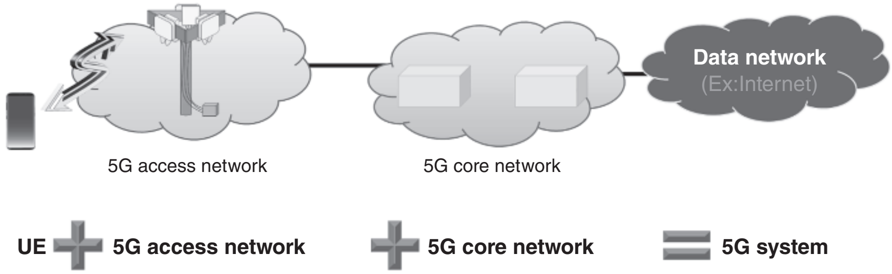

Figure 1.2 5G System: A High-Level Overview.

security activation, and protocol data unit (PDU) session setup. The gNB transparently transports the NAS signaling messages between the UE and the NGC. A UE can connect to and exchange user traffic with data networks (DNs) such as the internet through the 5G AN and the NGC.

The 5G O-RAN represents one possible way of designing and deploying a 5G AN. An operator or a service provider may deploy a 5G RAN with or without O-RAN, although O-RAN is expected to be a dominant way of deploying 5G RAN. The 5G O-RAN can be viewed as the realization of the 5G RAN, where gNBs or gNB entities such as gNB-Central Units (gNB-CUs) and gNB-Distributed Units (gNB-DUs) comply with O-RAN specifications. The O-RAN specifications utilize concepts such as disaggregation, openness, AI, and virtualization and cloud technologies. The 5G O-RAN does not alter the 3GPP-defined NR air interface between the UE and the 5G RAN in any way. A 5G UE does not need to be aware of the O-RAN and does not need to do anything differently to be able to work with the 5G O-RAN. A 5G O-RAN appears to a 5G UE as a 3GPP-defined 5G RAN. Similarly, the NGC does not need to be aware of the O-RAN and does not need to do anything differently to be able to work with the 5G O-RAN.

1.3 Evolution of the RAN

The 3GPP has defined the architectures for 3G, 4G, and 5G. Such architectures are referred to as logical architectures, where logical links between the network nodes, network elements, or network functions (NFs) are defined. For example, two gNBs can communicate with each other using the Xn interface. Specifically, two gNBs exchange Xn Application Protocol (XnAP) signaling with each other. The user traffic can also traverse on the Xn interface between two gNBs via a GTP[1] tunnel in support of packet forwarding during the inter-gNB handover. The 3GPP does not dictate or mandate any specific physical implementation or realization of the architecture. Hence, the operator has flexibility in deciding how to design, provision, and deploy the network architecture to meet target objectives.

Figure 1.3 illustrates how RAN evolution has occurred through generations of cellular technologies.

The older generations of cellular technologies from 1G to 3G utilize the RAN architecture with centralized RAN controllers called Base Station Controllers (BSCs) or Radio Network Controllers (RNCs). A BSC or RNC typically controls hundreds of base stations (BSs) or Node Bs. Certain radio resource management (RRM) tasks, such as the handover algorithm, are implemented at the BSC/RNC. The RAN thus consists of one or more BSCs and numerous BSs or Node Bs. The BSs are distributed across the cellular service area.

In commercial deployments, it is common to divide a BS into two parts, a baseband unit (BBU) and a radio unit (RU). An RU includes components such as radio frequency (RF) amplifiers and filters and is often situated next to the antenna that transmits and receives RF signals in support of wireless communications. A BBU carries out technology-specific processing at the baseband. The BBU and the RU can be connected via an RF cable to carry an analog signal. It is also possible to use an optical fiber between the BBUs and the RU, which reduces the overall loss for the link between the mobile device and the BS. Common Public Radio Interface (CPRI) is a widely used protocol between the BBU and the RU.

In a distributed RAN, BSs are distributed across the service area and a centralized RAN controller is absent. In such flat or distributed architecture, the BS manages its own radio resources and can

---

[1] GTP stands for GPRS Tunneling Protocol. GPRS is General Packet Radio Service.

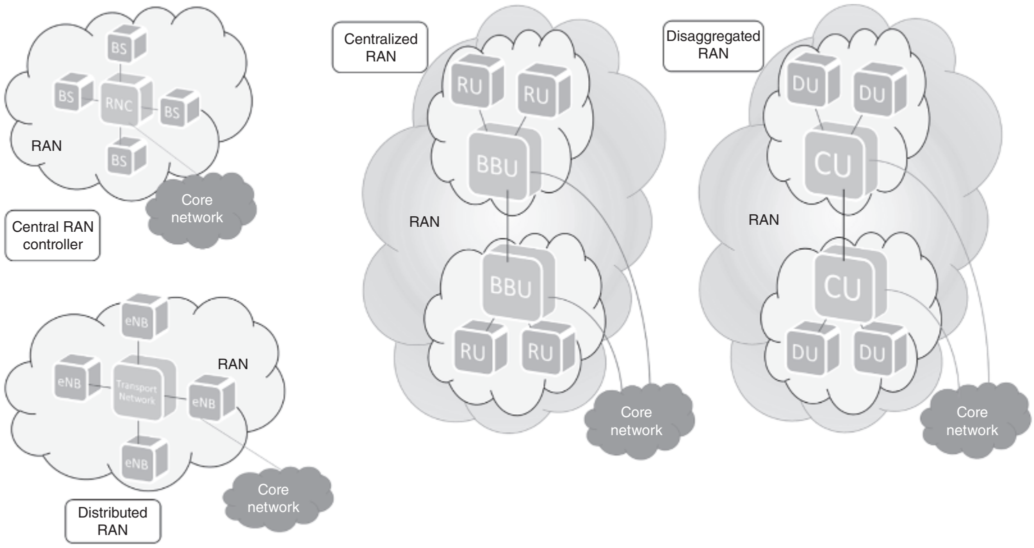

Figure 1.3 Evolution of the RAN Architecture.

communicate with other BSs through an interface. For example, in 4G LTE, the BSs called evolved Node Bs (eNBs or eNodeBs) communicate with each other using the Xn interface. An internal IP network of the service provider is used to implement the Xn interface between two eNBs. Such distributed or flat RAN is quite scalable, and new eNBs can be added to meet capacity and coverage needs in a given service area. Each eNB needs to connect to the core network. In commercial deployments, the eNB can be deployed using the BBU and the RU with relatively short BBU-RU distance such as few or tens of meters. Several 4G LTE and 5G deployments have used such distributed RAN.

In a centralized RAN (C-RAN) approach, a set of BBUs is placed in a relatively central data center and remote radio units (RRUs) or remote radio heads (RRHs) are placed at cell sites. Since the distance between the BBU and the RU is quite long (e.g., hundreds of meters or few kilometers), such RUs are called remote radio units. The C-RAN architecture is commercially deployed in LTE networks. The C-RAN provides benefits such as energy cost savings due to colocation of multiple BBUs, real estate cost savings due to a smaller footprint at the cell site, and performance enhancement due to enhanced collaboration of radio resources across multiple cells.

1.3.1 Horizontal Disaggregation of RAN: A Protocol Perspective

A disaggregated RAN divides the BS into a central unit (CU) and a distributed unit (DU). For example, while 5G gNBs have been deployed using the distributed RAN architecture, the 3GPP allows supports disaggregated gNBs, where a gNB is divided into a gNB-CU and a gNB-DU. The NR radio interface protocol stack is implemented partially at the gNB-CU and partially at the gNB-DU. Time-sensitive processing (e.g., acknowledgments for received data and any needed retransmissions) is carried out at the gNB-DU. Relatively less time-sensitive processing such as AS signaling is carried out at the gNB-CU. The gNB-DUs are placed at cell sites. While gNB-CUs can be colocated with gNB-DUs, gNB-CUs are often located in data centers. A gNB-CU connects to other gNB-CUs and the NGC. The 3GPP also allows decomposition of the gNB-CU into the gNB-CU-control plane (gNB-CU-CP) and the gNB-CU-user plane (gNB-CU-UP).

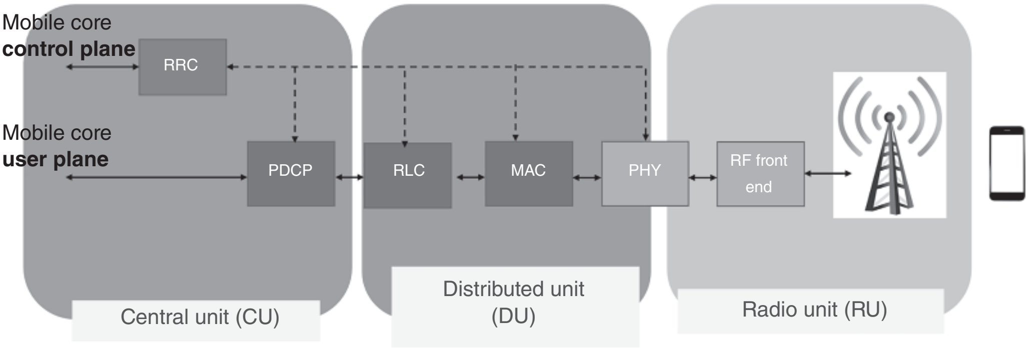

Figure 1.4 Disaggregation of the Base Station.

Figure 1.4 illustrates disaggregation of the BS from the perspectives of the radio interface protocol stack and control plane (CP) and user plane (UP) (Peterson et al. 2023). The UP, responsible for carrying user traffic, utilizes the protocols such as Packet Data Convergence Protocol (PDCP), Radio Link Control (RLC), Medium Access Control (MAC), and Physical (PHY) layer with the RF front end. In Figure 1.4, the CU implements RRC and PDCP protocols; the DU implements RLC, MAC, and selected PHY layer functions; and the RU implements selected PHY layer functions and the RF front end.

Software-Defined RAN: Figure 1.5 illustrates the implementation of the RAN using software-defined networking (SDN) principles, which leads to the development of a software-defined RAN. This approach facilitates the incorporation of a programmatic API that supports software-based control over the RAN UP pipeline. In the O-RAN architecture, this programmatic API is commonly known as the RAN Intelligent Controller (RIC). Refer to Chapter 2 for more details on RICs.

The 5G O-RAN makes use of the 3GPP-defined disaggregated 5G RAN architecture, carries out further disaggregation of the gNB-DU, and facilitates management of the RAN using intelligent controllers as described in Section 1.4 and Chapters 2 and 3.

When the RAN components are implemented using suitable RAN component software (e.g., the gNB-CU software) on generic commercial-off-the-shelf (COTS) hardware, such RAN implementation is termed as virtualized RAN (vRAN). The vRAN implementation no longer requires the RAN software to be coupled with specific or proprietary hardware. Hence, a service provider can purchase COTS hardware such as servers and instantiate suitable RAN software of its network vendors on such COTS servers to create a vRAN. While a vRAN can be deployed on a small scale using in-house network resources, cloud resources can also be exploited to build a vRAN. When cloud

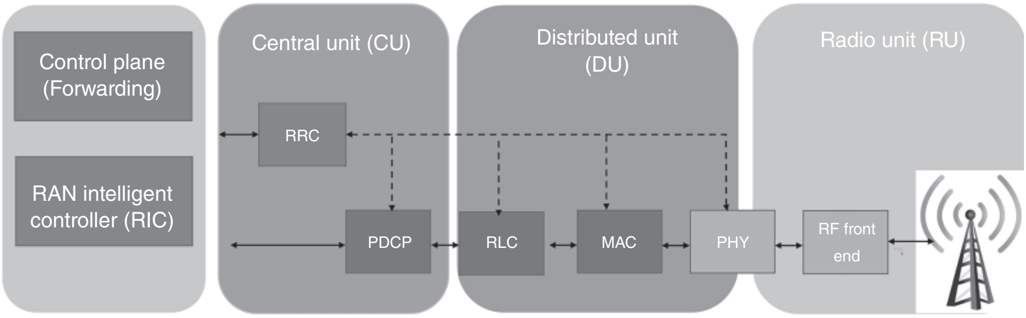

Figure 1.5 Disaggregated Base Station with SDN Principles.

resources are used to build a vRAN, such vRAN is also known as Cloud-RAN. For example, cloud resources of traditional cloud service providers or cloud vendors, such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP), can be used to create Cloud-RAN.

1.4 Introduction to the O-RAN ALLIANCE and the O-RAN Architecture

The O-RAN ALLIANCE is an operator-driven industry alliance that has defined the O-RAN architecture. In this book as well as general public literature, O-RAN usually refers to the O-RAN architecture defined the O-RAN ALLIANCE. Section 1.4.1 provides a brief overview of the O-RAN ALLIANCE. Section 1.4.2 describes a high-level O-RAN architecture. Section 1.4.3 summarizes key focus areas of various groups operating in the O-RAN ALLIANCE. Further details on the O-RAN ALLIANCE and its groups can be found in Chapter 5. A comprehensive view of the O-RAN architecture is given in Chapter 2.

1.4.1 O-RAN ALLIANCE: A Brief Overview

The O-RAN ALLIANCE was founded in February 2018 by AT&T, China Mobile, Deutsche Telekom, NTT DOCOMO, and Orange. It was established as a German entity in August 2018. Today's O-RAN ALLIANCE is a global community of mobile network operators (MNOs), vendors, and research and academic institutions. The mission of the O-RAN ALLIANCE is to "*re-shape the RAN industry towards more intelligent, open, virtualized and fully interoperable mobile networks*" (O-RAN ALLIANCE 2024a).

The O-RAN ALLIANCE emphasizes *openness* and *intelligence* for the next-generation wireless networks and beyond (O-RAN ALLIANCE 2018). Open interfaces enable operators to build cost-effective and agile RAN through multivendor deployments in a more competitive RAN supplier ecosystem. The use of open-source software and reference hardware designs can accelerate the pace of RAN innovations. Intelligence is important to manage increasing complex RAN by minimizing human involvement, harnessing the power of AI/machine learning (ML), and automating network operations. Embedded intelligence at the component and the network level facilitates dynamic and efficient RRM, optimizing networkwide efficiency. The development of O-RAN specifications by the O-RAN ALLIANCE enables a more competitive and diverse RAN supplier ecosystem and can accelerate the pace of RAN innovations to increase efficiency, reduce costs, and enhance user experience.

Figure 1.6 shows the three focus streams of the O-RAN ALLIANCE. The specification efforts involve extending RAN specifications such as the 3GPP 5G specifications to achieve openness and intelligence. Several work groups (WGs)[2] have been formed to address different aspects of O-RAN and define the overall O-RAN specifications. The O-RAN Software Community (OSC) of the O-RAN ALLIANCE, in cooperation with the Linux Foundation, is in charge of the release and the development of open-source software for O-RAN. The testing and integration efforts facilitate O-RAN members in testing and integration of O-RAN implementations to ensure interoperability of products and compliance with O-RAN specifications.

There are several governing bodies in the O-RAN ALLIANCE. The O-RAN ALLIANCE board consists of up to 15 members with 5 founding members and up to 10 elected members. The elections for the (elected) board members take place every two years. The executive committee

---

[2] Section 1.4.3 provides a summary of O-RAN work groups. Chapter 5 provide more details about the work groups.

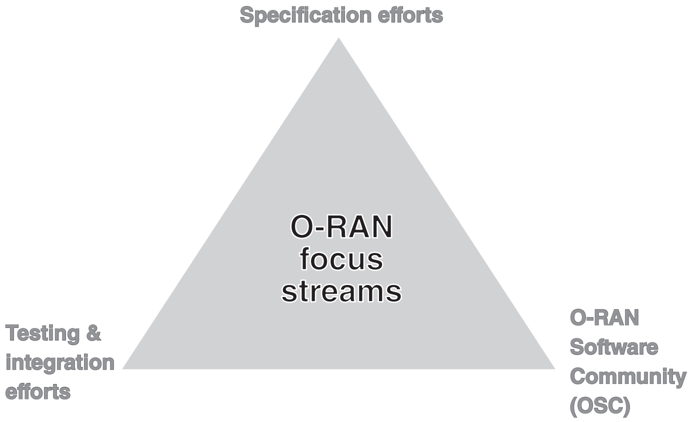

Figure 1.6 The Three Focus Streams of the O-RAN ALLIANCE.

(EC) supports the board by proposing agendas, priorities, projects, and releases for the board to consider and approve. The EC consists of representatives of the five O-RAN founding members and two elected representatives from the board members. The Technical Steering Committee (TSC) decides or gives guidance on O-RAN technical topics and approves O-RAN specifications prior to the Board approval and publication. The TSC consists of member representatives and the WG and focus group cochairs.

1.4.2 The O-RAN Architecture in a Nutshell

Figure 1.7 illustrates a simplified O-RAN architecture (O-RAN ALLIANCE 2024c).

The 3GPP-defined gNB-CU-CP and gNB-CU-UP are transformed into O-RAN-CU-CP (O-CU-CP) and O-RAN-CU-UP (O-CU-UP) by the O-RAN ALLIANCE, respectively. Furthermore,

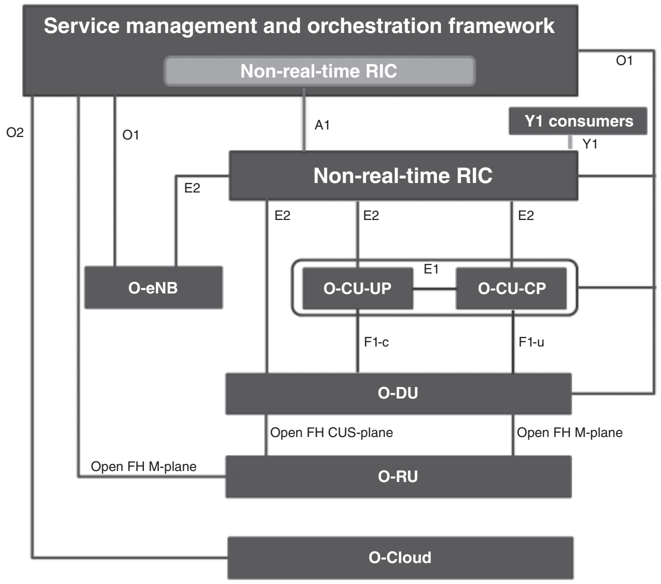

Figure 1.7 Simplified Logical O-RAN Architecture.

the 3GPP-defined gNB-DU is disaggregated and transformed into O-RAN-Distributed Unit (O-DU) and O-RAN-Radio Unit (O-RU) by the O-RAN ALLIANCE. A new interface called the open fronthaul between the O-DU and the O-RU is defined by the O-RAN ALLIANCE. Such disaggregation of the gNB-DU is often referred to as functional split 7.2x. The 3GPP had studied numerous functional splits of the NR radio protocol stack and finally standardized functional split Option 2 corresponding to the gNB-CU and the gNB-DU. The O-RAN ALLIANCE, in addition to supporting and augmenting such functional split, utilizes an additional split 7.2x. The O-DU and the O-RU interact with each other using two interfaces, open fronthaul control–user plane–synchronization (CUS)-plane and open fronthaul management (M)-plane. Details on this functional split 7.2x are given in Chapter 5.

The O-RAN ALLIANCE introduces two intelligent controllers in the spirit of the intelligence aspect of the O-RAN architecture, Non-Real-Time RAN Intelligent Controller (Non-RT RIC) and Near-Real-Time RAN Intelligent Controller (Near-RT RIC). The Non-RT RIC is part of the Service Management and Orchestration (SMO) framework and operates with a time scale exceeding 1 second. The Near-RT RIC operates with a time scale between 10 ms and 1 second. The O-RAN ALLIANCE-defined E2 interface enables the O-CU-CP, the O-CU-UP, and the O-DU to provide measurements to the Near-RT RIC. The Near-RT RIC executes suitable software applications called extended applications (in short, xApps) to manage the O-CU-CP, the O-CU-UP, and the O-DU. For example, an xApp in the Near-RT RIC can influence the scheduling algorithm implemented in the O-DU. Similarly, the Non-RT RIC executes suitable software applications called remote applications (in short, rApps) to manage the O-CU-CP, the O-CU-UP, and the O-DU. For example, an rApp may adapt handover parameters to optimize handover performance. In addition to supporting 5G RAN, the O-RAN also supports 4G LTE as seen from the existence of the O-RAN-eNB (O-eNB) in the O-RAN architecture.

The SMO framework interacts with the Near-RT RIC, the O-CU-CP, the O-CU-UP, O-DU, and the O-eNB through the O1 interface. The O-RAN-Cloud (O-Cloud) can be viewed as the cloud infrastructure on which suitable O-RAN functions, such as the O-CU-CP and the O-CU-UP, are implemented.

Cellular RAN architectural evolution and SDN evolution: an analogy

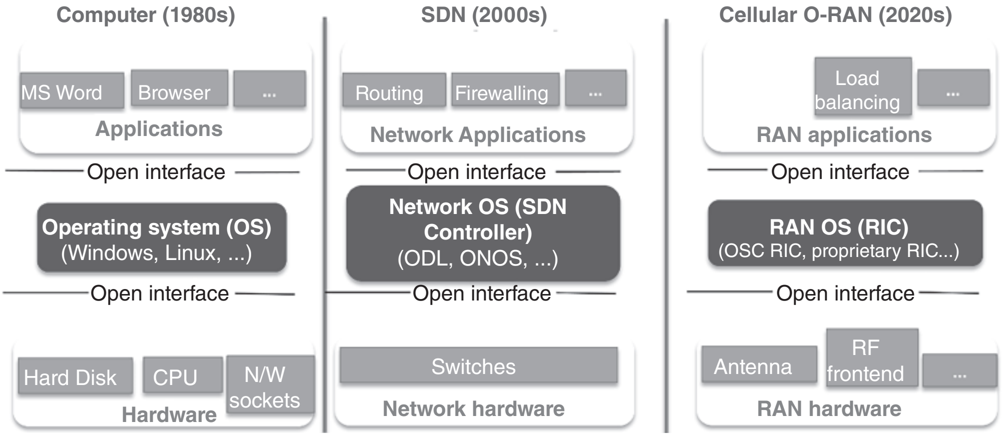

(Continued)

(Continued)

Let us draw parallels between the cellular RAN architectural evolution and computer and software-defined networking (SDN) evolution.

Consider what happened with computers back in the 1980s, when the proprietary, black box "Mainframes" were disaggregated into (i) computer hardware (e.g., hard disk and CPU) and (ii) computer applications (e.g., MS Word and a browser) with operating system (OS). The OS acts as a glue between the applications and computer hardware. This changed the computer world for better forever and set the foundation of today's computer world where every device, whether it is smartphone, smart watch, computer, or tablet, adopts the disaggregated computer architecture with OS as the middleware, which ensures that the applications (developed by independent application developers) run on the system without them worrying about how to work the computer hardware.

The same thing happened to the internet world in the early 2000s when SDN was introduced. Internet switches/routers were disaggregated into hardware switches (for data forwarding only), and all the intelligence in form of network applications such as routing and firewalling were moved to servers. And the middleware called SDN controller (informally referred to as network operating system) serves as the glue between network applications and hardware switches.

Similar evolution is happening with wireless networks with the introduction of the O-RAN paradigm. With O-RAN, the cellular RAN is disaggregated into RAN hardware (e.g., antennas and RF frontend) and RAN applications called xApps and rApps. The RAN Intelligent Controller (RIC) serves as the RAN Operation System (RAN OS) and functions as the glue between RAN applications and RAN hardware.

1.4.3 Groups in the O-RAN ALLIANCE

The technical O-RAN work is performed by O-RAN WGs and O-RAN focus groups and research groups. Table 1.1 lists these groups and summarizes their key focus areas.

As seen from Table 1.1, each of the WGs focuses on a part of the O-RAN architecture. In contrast, the focus and research groups deal with the topics that are overarching the technical WGs or relevant for the entire organization. See Chapter 5 for more details about the groups listed in Table 1.1.

1.5 Driving Forces Behind O-RAN

Several forces are driving the emergence of O-RAN as an approach to design and deploy 5G (and 4G) RAN as listed in Figure 1.8. O-RAN driving forces include RAN vendor diversity, complexity of RAN, accelerated innovations, and geopolitical environment.

RAN Vendor Diversity. In early generations of cellular technologies, a cellular service provides numerous RAN vendors to choose from. Examples of these vendors included Alcatel, Lucent Technologies, Motorola, Siemens, Nokia, Nortel Networks. Huawei, Samsung, and ZTE. Due to factors such as rising competition and economic environments, consolidation among RAN vendors occurred over time and the number of 4G and 5G RAN vendors has reduced significantly.

Table 1.1 Groups in the O-RAN ALLIANCE.

Group	Key Area of Focus
WG1: Use Cases and Overall Architecture Work Group	Define overall O-RAN architecture and use cases
WG2: The Non-Real-Time RAN Intelligent Controller and A1 Interface Work Group	Support Non-RT intelligent radio resource management, higher layer procedure optimization, policy optimization in RAN, and transfer of AI/ML models to Near-RT RIC
WG3: The Near-Real-Time RIC and E2 Interface Work Group	Define the Near-RT RIC architecture and support for data collection and actions over E2 interface
WG4: The Open Fronthaul Interfaces Work Group	Define open fronthaul interfaces to enable multivendor DU-RU interoperability
WG5: The Open F1/W1/E1/X2/Xn Interface Work Group	Define 3GPP-compliant multi-vendor profile specifications for F1/W1/E1/X2/Xn interfaces and propose 3GPP specification enhancements (if any)
WG6: The Cloudification and Orchestration Work Group	Enable decoupling of RAN software from the underlying hardware platforms to produce technology and reference designs that leverage commodity hardware platforms
WG7: The White-box Hardware Work Group	Specify and release a complete reference design to foster a decoupled software and hardware platform
WG8: Stack Reference Design Work Group	Develop the software architecture, design, and release plan for the O-CU and O-DU based on O-RAN and 3GPP specifications for the NR protocol stack
WG9: Open X-haul Transport Work Group	Specify transport equipment, physical media, and control/management protocols associated with the transport network
WG10: OAM Work Group	Specify the OAM requirements, OAM architecture, and the O1 interface
WG11: Security Work Group	Address security aspects of the O-RAN ecosystem
SDFG: Standard Development Focus Group	Determine the standardization strategies and interface with other Standard Development Organizations (SDOs)
OSFG: Open-Source Focus Group	Successfully launched the O-RAN Software Community (OSC) (The group is now dormant group because the open-source software development activities are being carried out by the OSC)
TIFG: Testing and Integration Focus Group	Define the overall approach for testing and integration including coordination of test specifications across various WGs
nGRG: Next-Generation Research Group	Carry out research of open and intelligent RAN principles in 6G and future network standards
IEFG: Industry Engagement Focus Group	Promote and accelerate O-RAN based technology adoption, proliferation and continuous innovation by engaging with industry/ecosystem
SuFG: Sustainability Focus Group	Optimize energy consumption, reduce environmental impact, and create more energy-efficient and environmentally friendly mobile networks

O-RAN provides opportunities for new players to be an integral part of the O-RAN ecosystem due to open interfaces, a high degree of disaggregation, and support for independent development of AI-based applications RAN controllers.

Complexity of the RAN. 5G is significantly more complex than 4G LTE. Remember that flexibility is a euphemism for complexity! There are several newer RAN operations and degrees of

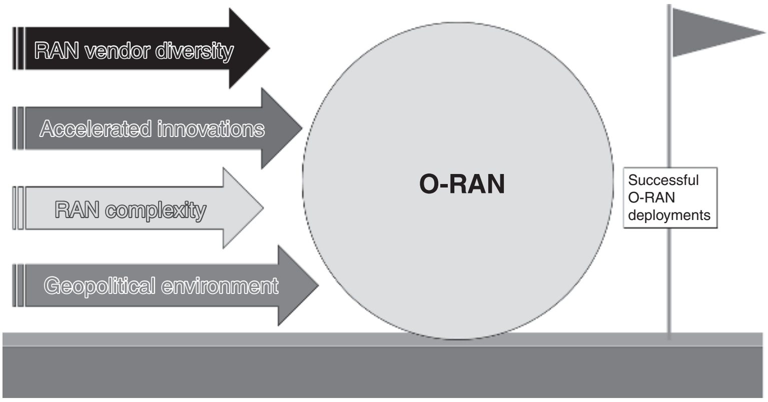

Figure 1.8 Major Driving Forces Behind O-RAN.

freedom in 5G relative to LTE such as flexible OFDM numerologies, carrier bandwidth parts, highly granular time-frequency radio resources for resource allocation, and support for network slicing for diverse applications and customer requirements. The definitions of open interfaces and support for intelligent RAN controllers enable the service providers to automate RAN operations and optimization to a large degree, reducing the probability of human errors and enhancing the overall RAN performance.

Accelerated Innovations. According to an anonymous Facebook philosopher,[3] when two people with one idea each work separately, each has one idea. But when they work together, they both have two ideas! That is the power of collaboration. When more companies in the ecosystem, the pace of innovations can accelerate significantly compared to the case of limited RAN vendors.

Geopolitical Environment. Security concerns of the US government about Chinese RAN vendors such as Huawei led to the banning of the Huawei equipment in the United States. Indeed, the US Congress passed a bill to help operators replace the Huawei equipment. Additionally, supply chain management has become an important policy consideration for the US government. Governments of several other countries such as Germany and the United Kingdom also share the security concerns of the US government. There is a significant push by the US government to support O-RAN to increase supplier diversity, ensure security and trustworthiness of RANs, and alleviate supply chain management issues, with the US congress passing a bill to allocate funding to support adoption of O-RAN. In summary, a geopolitical environment is a major driving force behind O-RAN.

1.6 O-RAN Ecosystem

While the O-RAN ALLIANCE defines a complete architecture for the RAN, design and implementation of such architecture in practice require support from other organizations. The overall O-RAN ecosystem includes (i) entities directly working on the three focus streams of O-RAN, (ii) organizations that develop specifications, and (iii) organizations that develop, maintain, and

---

[3] People share humorous and often philosophical quotes on social media such as Facebook. Some of these quotes are extremely valuable but quite difficult to fully implement in practice. All that matters is we move in the right direction!

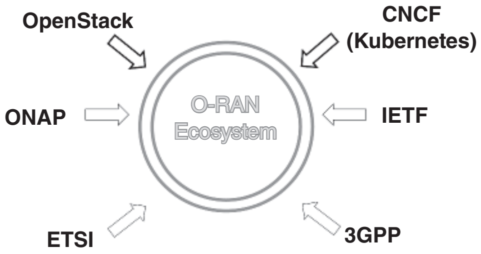

Figure 1.9 Building Blocks of the O-RAN Ecosystem.

publish software packages that are utilized by O-RAN. Figure 1.9 mentions examples of key entities that are building blocks for designing, deploying, and managing the O-RAN architecture. These entities include 3GPP, European Telecommunication Standards Institute (ETSI), Internet Engineering Task Force (IETF), Cloud Native Computing Foundation (CNCF), OpenStack, and Open Network Automation Protocol (ONAP).

3GPP. The O-RAN ALLIANCE utilizes the logical RAN architecture defined by the 3GPP as the baseline. For example, the O-CU performs all the functions of the gNB-CU defined by the 3GPP. Similarly, the O-DU and O-RU together carry out all the tasks of the gNB-DU defined by the 3GPP. The O-CU, the O-DU, and the O-RU together implement the 5G NR radio interface protocol.

ETSI. O-RAN functions or nodes such as the O-CU, the O-DU, the Non-RT RIC, and the Near-RT RIC can be implemented using the cloud infrastructure O-Cloud. O-Cloud makes use of technologies such as network functions virtualization (NFV) can be used in O-Cloud to implement O-RAN nodes. The NFV architecture is defined by ETSI. See O-RAN ALLIANCE (2024g) for a list of the relevant ETSI NFV standards documents that are useful for realizing O-RAN. These documents are related to virtualized network function (VNF) packaging, Life Cycle Management (LCM) operations, and security for managing virtual machines (VMs) and containers. Examples of ETSI NFV standards relevant to the O-RAN include (i) SOL003, SOL005, and SOL006 standards for VNFs and (ii) ETSI GS NFV-SOL 013 for O2 interface security. Hence, ETSI plays an important role in the realization of O-RAN.

IETF. O-RAN nodes can interact with each other using IP. To ensure security of communications between the O-RAN nodes, IP-level security mechanisms can be applied. Examples of such mechanisms applicable to O-RAN include IPsec and Transport Layer Security (TLS) (O-RAN ALLIANCE 2024f). The 3GPP also supports the use of these mechanisms in a 5G network to provide security at the IP layer. The IETF is responsible for specifying these IP-based mechanisms.

CNCF. Kubernetes, abbreviated as K8s, is a portable, extensible, and open-source platform for managing containerized workloads and services (Kubernetes 2024). K8s is a container orchestration engine for automating deployment, scaling, and management of containerized applications. Google originally developed K8s and open sourced it in 2014. The K8s project is hosted by the CNCF. K8s can be used in the design and operation of an O-RAN platform such as the Near-RT RIC platform.

OpenStack. OpenStack is a widely deployed open-source cloud software. OpenStack controls large pools of compute, storage, and networking resources through APIs or a dashboard. OpenStack provides functionalities such as infrastructure-as-a-service, orchestration, fault management, and service management (OpenStack 2024). Since O-RAN is highly virtualized, OpenStack can be used to manage O-Cloud compute, storage, and networking resources.

ONAP. ONAP is an open-source platform that can be used to manage functions and/or components of a virtualized network. Since virtually all (!) O-RAN nodes can be virtualized, ONAP can be viewed as a candidate platform for automation and orchestration in an O-RAN implementation. ONAP can be used for a specific purpose such as facilitating SON and network slicing use cases. For example, the sixth release of ONAP, called Frankfurt, focuses on support for 5G network slicing (Rethink 2024). It also features increased support for the O-RAN O1 interface for fault, performance, and configuration management. New functions to support A1 policies have been added to ONAP (Bonneau and Keeney 2023). A1 policies enable the Non-RT RIC to provide policy guidance to the Near-RT RIC. Specifically, ONAP supports O-RAN's A1 interface, between the orchestration layer (which could be ONAP or another MANO framework) and the RIC.

1.7 O-RAN Use Cases

The O-RAN architecture is intended to facilitate RAN operations and optimization. The O-RAN ALLIANCE has published white papers and specifications related to O-RAN use cases. Initial O-RAN use cases are discussed in a 2020 White Paper (O-RAN ALLIANCE 2020), and more use cases are described in a 2021 White Paper (O-RAN ALLIANCE 2021). Furthermore, multiple use cases are briefly described at a high level by O-RAN WG1 as part of Release R003 (O-RAN ALLIANCE 2024d). Details of the O-RAN use cases can be found in O-RAN ALLIANCE (2024e).

Table 1.2 lists the O-RAN use cases mentioned in O-RAN ALLIANCE (2024d). *These use cases provide a glimpse of the types of RAN management that the O-RAN architecture can facilitate. As new 5G features and capabilities are defined, O-RAN use cases are expected to grow.*

These use cases are briefly summarized below. Note that openness and intelligence are two central themes of O-RAN as mentioned in Section 1.4.1. The use cases briefly described below should be viewed as examples of potential ways in which intelligence aspects of O-RAN could be utilized. An operator may use RICs or other means to exploit intelligence in the RAN. Furthermore, operators may or may not utilize these use cases in their networks. Additionally, RAN functions other than the ones mentioned below could be exploited in O-RAN. Since RAN functions are typically implementation specific and proprietary and not standardized, there is a significant degree of flexibility and innovations in realizing RAN functions. However, motivation and typical processing are briefly explained for each use case below.

1.7.1 Context-Based Dynamic HO Management for V2X

The use of vehicle-to-everything (V2X) communication helps realize benefits such as increased safety, reduced emissions, and shorter commute times. However, the V2X environment is quite dynamic due to the dynamic nature of vehicular traffic and heterogenous environment, resulting in a challenging situation for handover. This use case aims to prevent or minimize performance degradation of V2X applications due to ineffective handover using two main functionalities. The *first functionality* makes use of long-term analytics, where the Non-RT RIC utilizes the UE-based handover events and mobility data maintained by the V2X application server. The Non-RT RIC identifies causes of handover anomalies and determines optimal handover sequences. A database records resolutions to anomalies. The *second functionality* involves a trained ML model in the Near-RT RIC that performs real-time (RT) optimization. The Near-RT RIC monitors UE-specific mobility context, detects or predicts unexpected handover events, and generates a target handover sequence to address the handover anomaly. The Near-RT RIC creates and optimizes the UE-specific Neighbor Relation Table (NRT) that contains the list of neighbor cells.

Table 1.2 Examples of O-RAN Use Cases.

Use Case ID	Brief Use Case Description	Use Case ID	Brief Use Case Description
1	Context-Based Dynamic HO Management for V2X	14	Massive MIMO SU/MU-MIMO Grouping Optimization
2	Flight Path-Based Dynamic UAV Radio Resource Allocation	15	O-RAN Signaling Storm Protection
3	Radio Resource Allocation for UAV Application Scenario	16	Congestion Prediction and Management
4	QoE Optimization	17	Industrial IoT Optimization
5	Traffic Steering	18	BBU Pooling to Achieve RAN Elasticity
6	Massive MIMO Beamforming Optimization	19	Integrated SON Function
7	RAN Sharing	20	Shared O-RU
8	QoS-Based Resource Optimization	21	Energy Saving
9	RAN Slice SLA Assurance	22	MU-MIMO Optimization
10	Multivendor Slices	23	Sharing Non-RT RIC Data with the Core
11	3GPP 4G/5G Dynamic Spectrum Sharing (DSS)	24	Industrial Vision SLA Assurance
12	NSSI Resource Allocation Optimization	25	Non-Public Network (NPN) RAN-Sharing via Midhaul for Multi-Operator Coverage
13	Local Indoor Positioning in RAN	26	Interference Detection and Optimization

1.7.2 Flight Path-Based Dynamic UAV Radio Resource Allocation

An uncrewed/unmanned aerial vehicle (UAV) experiences a unique radio environment compared to a traditional device on the ground. A UAV's radio propagation environment is characterized by line-of-sight propagation, increased amount of interference, signal variations due to the sidelobes of antenna radiation patterns, and more sudden changes in signal strengths. This use case aims to address the radio environment challenges for UAVs by exploiting flight path-based radio resource allocation. The Non-RT RIC obtains UAV-related metrics from suitable measurement reports from UAV UEs. The Non-RT RIC also obtains supporting information from a suitable application (e.g., unmanned traffic management or UTM), including the flight path information of the UAV, climate information, forbidden or restricted area information, and space load information, and builds (and later updates) an AI/ML model. The AI/ML model may consider uplink (UL)/downlink (DL) interference from/to UAVs, detection of UAVs, and prediction of available radio resources. The Near-RT RIC deploys and executes the AI/ML model developed by the Non-RT RIC and allocates radio resources allocation to UAV UEs.

1.7.3 Radio Resource Allocation for UAV Application Scenario

This use case envisions a UAV control vehicle that houses 5G O-RAN to control operations of various UAV UEs to provide reliable communications services. A large amount of data needs to be processed locally in a UAV control vehicle so that the overall UAV system can meet its operational objectives. Examples of data being communicated between the UAV UE and the UAV control vehicle include both the application data and the control data. The application data such as 4k high-definition videos require asymmetric DL/UL data rates and high UL data rates. In contrast, the control data such as navigation commands, configuration changes, and flight status data reporting require low latency, but low data rates are adequate. The Non-RT RIC obtains UE-level radio resource requirements from a suitable AS to determine UE-specific radio resource allocation policies. The Near-RT RIC executes these policies and creates configuration parameters for the E2 nodes. The E2 nodes adjusts radio resources and provide information about the UE status to the Near-RT RIC.

1.7.4 QoE Optimization

Diverse quality of experience (QoE) requirements of applications such as Cloud VR, industrial automation, online multiplayer gaming, and connected vehicles cannot be met by a semi-static quality of service (QoS) framework. QoE requirements can often vary during an application's lifetime. QoE estimation/prediction from the application layer and exposure of RAN performance to the application layer can be used to address such uncertainty, increase the efficiency of radio resource utilization, and enhance user experience. This use case aims to develop an automated closed-loop approach to optimize QoE. Specifically, the use case envisions a "User RAN Policy" that resides at the Non-RT RIC or Near-RT RIC and applies the operator's desired QoE configuration for a specific user, slice or 5QI flow in response to (i) requests from external systems or (ii) UE mobility. This policy supports reporting of the service-level agreement (SLA) information to external systems hosting the user applications. The use case exposes per-user or per-cell radio performance analytics information to external applications so that these applications can optimize user QoE.

1.7.5 Traffic Steering

Multiple access technologies such as 5G NR, LTE, and Wi-Fi with different amounts of spectrum may be available for use in a given geographic area. 5G NR may use licensed spectrum, unlicensed spectrum, or shared spectrum. Variations in the traffic load across cells of different access technologies or variances in available bandwidths and QoS attributes may lead to suboptimal spectrum utilization, limiting the achievable user experience. This use case aims to configure the desired optimization policies and utilize suitable performance criteria to proactively manage user traffic across different access technologies. The use case also enables the use of the enrichment information such as radio fingerprint information that is based on the data analytics of the historical RAN data. The Non-RT RIC observes the user experience by the UE-level performance measurements and the cell-level resource utilization and compares the observed performance against the expected service-level requirements. If the expected requirements are not met, the Non-RT RIC determines the related cell and evaluates the radio resource conditions in suitable cells. The Non-RT RIC may decide to relocate one or more users to other cells or even split user traffic across technologies. The Non-RT RIC creates traffic management policies to relocate UEs along with the prioritized list of cells. These policies from the Non-RT RIC are used by the Near-RT RIC to manage radio resources of different technologies.

1.7.6 Massive MIMO Beamforming Optimization

This use case aims to adjust coverage of massive MIMO beams proactively to enhance capacity/throughput and mobility performance by adapting beam configurations in a non-RT or near-RT situations. Example parameters that may be adjusted include the number of beams, beam boresights, vertical beamwidths, horizontal beamwidths, beam black lists or white lists, and beam mobility thresholds. Non-RT situations involve factors such as 3D construction, 3D terrain topology, network, weather seasons, intraday cell splitting/merging/shaping, traffic distribution, beam conflicts. Near-RT situations include moving users or hotspots, changing traffic distribution, and crowd source data. The O-RAN framework is highly attractive in this use case due to the large number of configuration parameters per antenna array, the amount of available measurements, and the overall complexity of MIMO. This use case envisions three optimization loops. An outer Non-RT loop implements massive MIMO grid of beams (GoB) beamforming optimization. An inner Near-RT loop implements massive MIMO beam-based mobility robustness optimization. Another inner Near-RT loop handles massive MIMO beam selection optimization.

1.7.7 RAN Sharing

Network sharing is often considered to be an efficient and sustainable way to accelerate the deployment of 5G due to a common pool of physical infrastructures shared by multiple operators. O-RAN can enable efficient RAN sharing including whole or part of O-CU and O-DU due to open interfaces and RICs. In one implementation approach, two operators share the same RAN but have their own independent core networks. Operators can share the state of their radio resources to facilitate optimization of radio resource utilization.

In an example RAN sharing approach, multiple service providers with business relationship with their customers can share the same RAN. Such a RAN can advertise public land mobile network (PLMN) identities of the supported service providers. Such a RAN sharing arrangement can be quite attractive to service providers economically, because they do not need to deploy their dedicated RANs to serve their customers. We observe that a Multi-Operator Radio Access Network (MORAN) is a widely used term to refer to RAN sharing that allows multiple operators to share the same RAN infrastructure. In an example implementation of MORAN, the infrastructure resources such as the antennas, tower, cell site, and power are shared among multiple operators, but operators have their dedicated radio channels or carrier frequencies. Operators can still independently implement desired features on a given MORAN.

1.7.8 QoS-Based Resource Optimization

This use case aims to optimize resource allocation for a special category of users such as public safety personnel (e.g., firefighters). In a statis implementation, a certain amount of radio resources may be typically used or reserved as part of a networks slice for such a special category of users. However, statis provisioning may be inadequate to fulfill the service-level specifications (SLS). In this use case, the Non-RT RIC monitors the performance of a RAN slice by obtaining QoS metrics from the E2 nodes (e.g., O-CU and O-DU) and the SMO functions. The Non-RT RIC may also receive information about user priorities from an external server. When needed, the Non-RT RIC modifies policies to fulfill the SLS. The Near-RT RIC receives the policies and provide suitable guidance to the E2 nodes so that the E2 nodes can execute the QoS enforcement decisions issued by the Near-RT RIC.

1.7.9 RAN Slice SLA Assurance

5G supports end-to-end network slicing, with the RAN being an important component of a network slice. This use case clarifies mechanisms and parameters for the SLA assurance for the RAN slice. In an example implementation, the SMO obtains RAN slice SLA. The Non-RT RIC observes long-term trends of the performance of RAN slices, evaluates this performance against the SLAs, and applies AI/ML techniques to adjust policies. The SMO can influence suitable policies, or the Non-RT RIC can provide policies to the Near-RT RIC. It is possible for the Non-RT RIC to develop AI/ML models that are executed by the Near-RT RIC. The policies may include information such as the network slice ID and key performance indicator (KPI) targets. The Near-RT RIC can optimize actions of E2 nodes by considering policies and slice-specific measurements and executing the deployed AI/ML models.

1.7.10 Multivendor Slices

This use case enables deployments of multivendor slices, where one slice is provided by a combination of O-CU(s) and O-DU(s) from one vendor and another slice is provided by a combination of O-CU(s) and O-DU(s) from another vendor. Different vendors may have unique features or advantages for different network slices. Hence, such multivendor slices provide the operator with significant flexibility in choosing a vendor for a given network slice. Supply chain risks are also mitigated because of the availability of products from multiple vendors.

1.7.11 3GPP 4G/5G Dynamic Spectrum Sharing (DSS)

The 3GPP has defined 4G/5G dynamic spectrum sharing (DSS) to enable spectrum sharing between 5G NR and LTE. This use case aims to exploit O-RAN to realize 4G/5G DSS. Specifically, an RIC can coordinate between an LTE scheduler and a 5G NR scheduler to share the same spectrum efficiently. The 4G/5G DSS can be realized using multiple applications. A resource management application, the 4G/5G DSS-App in the Non-RT RIC, manages the shared spectrum resources by adapting to dynamic 4G and 5G workload requirements. The 4G/5G DSS-App in the Non-RT RIC creates a long-term policy or intent as a scheduling guidance to 4G and 5G schedulers to reflect business, user, spatial, and temporal workload factors. The 4G/5G DSS-App in the Non-RT RIC converts the global 4G/5G DSS objectives such as workload requirements for a region and time-of-day into 4G/5G DSS policies (e.g., the maximum or minimum bandwidth threshold). Another application, RAT-App in the Non-RT RIC, maps the global 4G/5G DSS policies from the Non-RT 4G/5G DSS-App to RAT-specific policies. The 4G/5G DSS-App in the Near-RT RIC is involved in a closed-loop decision-making process by using the KPIs from the RAN and adapting to the needs of the 4G and 5G cells. The RAT-App in the Near-RT RIC performs RAT-specific configuration, control, and data subscription with the E2 nodes.

1.7.12 NSSI Resource Allocation Optimization

5G supports diverse use case scenarios, such as eMBB, URLLC, and mMTC, and utilizes Network Slice Subnet Instances (NSSIs) to offer these services. This use case trains an AI/ML model based on the huge volumes of performance data collected at O-RAN nodes over a long period (e.g., days, weeks, or even months) to predict the traffic demand patterns of 5G networks for each network slice subnet at target times and geographic locations. The AI/ML model then automatically reallocates the network resources to optimize network performance and prevent occurrence of any network performance issues. In an example implementation, the Non-RT RIC monitors the network performance by collecting measurements from the radio nodes such as O-CU and O-DU and utilizes an AI/ML model to process the measurements and predict the future traffic demand. The AI/ML model at the Non-RT RIC also determines the actions needed to adjust radio resources for the RAN NFs. The Non-RT RIC reconfigures the NSSI attributes (e.g., those related to the RRM policy) via the Operation, Administration, and Maintenance (OAM) functions in the SMO to configure the E2 nodes.

1.7.13 Local Indoor Positioning in RAN

The 3GPP supports a comprehensive framework to support positioning. For example, the 3GPP has defined the location management function (LMF) in the core network to support 5G NR-based positioning. However, in certain scenarios such as industrial manufacturing, it is more effective to carry out positioning locally in the RAN itself to reduce latency and expose such location to local applications. This use case utilizes the O-RAN architecture to deploy a positioning application in the Near-RT RIC. This positioning application determines the UE location and velocity using the measurements obtained through the E2 interface. This application can potentially make use of an AI/ML model that is trained based on historical signal measurements and known user locations. The positioning application can then expose the UE location to an edge application securely and/or to the SMO.

1.7.14 Massive MIMO SU/MU-MIMO Grouping Optimization

5G supports both single-user MIMO (SU-MIMO) and multi-user MIMO (MU-MIMO). In a given cell, different users are allocated different Physical Resource Blocks (PRBs) when SU-MIMO is used for such users. In contrast, MU-MIMO allocates the same PRB(s) to multiple users in the cell. The SU-MIMO and MU-MIMO are suitable for different scenarios. The performance of MU-MIMO is highly susceptible to the UE speed compared to the SU-MIMO performance. Furthermore, the MU-MIMO performance may be relatively lower in a low-traffic volume scenario compared to the SU-MIMO performance. Additionally, SU-MIMO and MU-MIMO have different resource requirements (e.g., reference signals). This use case adapts the transmission method (i.e., SU-MIMO versus MU-MIMO) for a given user based on factors such as the radio environment and traffic.

In an example implementation, the Non-RT RIC obtains necessary configurations, performance indicators, measurement reports, and auxiliary information such as user GPS information and traffic information to create suitable AI/ML models. The Non-RT RIC uses such trained AI/ML model to determine a list of UEs for an SU-MIMO group and an MU-MIMO group. The Non-RT RIC can also determine RRC configurations (e.g., sounding reference signal [SRS] and connected mode discontinuous reception [DRX] configurations) for the SU-MIMO and the MU-MIMO groups. The Near-RT RIC can send the configurations to E2 nodes by policy. The RAN nodes convey RRC configurations to UEs in both SU-MIMO and MU-MIMO groups and carry out scheduling for SU-MIMO and MU-MIMO groups. The RAN nodes also collect and report performance measurements to SMO related to SU-MIMO and MU-MIMO spectral efficiency (e.g., average number of spatial multiplexing layers, rank, and throughput for SU-MIMO and MU-MIMO groups).

1.7.15 O-RAN Signaling Storm Protection

Some devices may aggressively attempt to register or attach to the network at a rate of a few thousand times per hour. An attacker that can manipulate a large number of such vulnerable devices remotely can cause an attach or registration-related signaling storm, leading to potentially long outages. Intelligent and fine-grained security control at the RAN can stop such attacks, insulating the core network from in the core.

In an example implementation, an xApp can be built to address the signaling storm with two main functionalities, a distributed denial-of-service (DDoS) detection capability and a DDoS mitigation capability. A near-RT detection portion of the DDoS detection capability occurs in an xApp, and a non-RT detection portion (which relies on information available in the 5GC) occurs at the SMO.

The detection algorithms detect an abnormally high volume of CP messages originating from a group of devices. The detection algorithm may identify (i) a list of devices responsible for the signaling storm and/or (ii) geographic areas or gNBs expiring the signaling storm. Normal behavior may be learned by an AI/ML algorithm over time.

1.7.16 Congestion Prediction and Management

This use case proactively addresses congestion in a gNB by analyzing the radio resource utilization and taking a timely corrective action to mitigate any potential congestion. Congestion in a commercial cellular network leads to problems such as Radio Link Failures (RLFs), handover failures, and low data rates. Operators currently do not have a well-defined mechanism to predict congestion. This use case aims to utilize the embedded intelligence of O-RAN to predict congestion to initiate a cell congestion mitigation solution.

This use case proposes congestion prediction and management (CPM) architecture to proactively detect and mitigate congestion. In the CPM architecture, E2 node statistics such as counters are collected and preprocessed (e.g., addition of cell IDs and conversion of counters into KPIs) by the SMO. The Non-RT RIC invokes a suitable training model in an AI server inside SMO for data cleaning and training. Predicted KPIs are obtained by the CPM rApp. Suitable AI/ML models can be used to learn and predict the future traffic for the next hour/day/month. The CPM rApp in the Non-RT RIC forms the inference about cell congestion.

Two options to mitigate cell congestion are envisioned by this use case. In Option A, the CPM rApp in the Non-RT RIC transfers the inference to the CPM xApp in the Near-RT RIC through A1 interface, and the Near-RT RIC decides the inference-based mitigation solutions such as switching to dual connectivity mode, debarring of user access, and load sharing. These solutions can be controlled using the E2 interface. In Option B, the Non-RT RIC directly helps to mitigate congestion using the O1 interface using mitigation solutions such as cell splitting, addition of more carrier frequencies, switching to higher-order MIMO, and transferring some users to Wi-Fi.

1.7.17 Industrial IoT Optimization

In support of high reliability requirements for the industrial Internet of Thing (IIoT) scenarios such as factory automation and transportation industry, 5G makes use of features including data duplication (e.g., PDCP duplication), Time-Sensitive Networking (e.g., Ethernet header compression [EHC]), and different prioritized transmissions (e.g., preemption of lower priority transmissions by higher priority transmissions).

This IIoT optimization use case aims to guarantee reliability and resource utilization efficiency by optimizing mechanisms of PDCP duplication, EHC, and different prioritized transmissions. The enhancement mechanisms for IIoT are semi-static or dynamically configured for certain users and certain services based on the AI/ML-based service prediction and performance prediction. The solutions envisioned by this use case include Non-RT RIC optimization solution and Near-RT RIC optimization solution.

In the Non-RT RIC-based solution, the Non-RT RIC retrieves network data (e.g., network load, performance metrics, and 5QIs) and external service information from the SMO. The SMO also obtains external enrichment information such as service type, traffic period, duration, packet size, and KPI requirements from sources such as a multi-access edge computing (MEC) server, APP server, or industrial control platform. The Non-RT RIC retrieves configured policy information such as cell edge user reliability KPI target, Ethernet header compression efficiency from the SMO. The Non-RT RIC deploys and updates an AI/ML model received from the SMO. The Non-RT RIC uses the AI/ML model to infer and decide suitable IIoT configurations and actions (e.g., mapping between the QoS flow and the DRB that supports duplication and EHC and preemption decision). The Near-RT RIC retrieves IIoT-related configurations from Non-RT RIC and sends the configurations to E2 nodes by policy. In support of this use case, E2 nodes such as the O-CU and the O-DU provide UE measurement reports and performance metrics to the SMO. E2 nodes enforce the E2 control policy or O1 configuration.

In the Near-RT RIC solution of IIoT optimization, the SMO obtains external enrichment information such as service type, traffic period, duration, packet size, and KPI requirements from an MEC server, APP server, or industrial control platform. The SMO then provides to the Non-RT RIC the information including external service information and network data such as network load, performance metrics, and 5QIs The Non-RT RIC trains an AI/ML model for the Near-RT RIC. Furthermore, the Non-RT RIC also evaluates the collected data and A1 policy feedback and generates or updates suitable optimizations (e.g., reliability targets) and sends it to the Near-RT RIC.

The Near-RT RIC deploys or updates the AI/ML model. The Near-RT RIC also receives an A1 policy from the Non-RT RIC and initiates relevant optimization procedure(s). The Near-RT RIC evaluates the performance data from the E2 nodes to determine compliance to the A1 policy-dictated KPI targets. For example, cell edge user performance cannot meet the A1 policy. The Near-RT RIC considers a variety of factors such as service characteristics, channel quality estimation, and A1 policy to determine actions such as PDCP duplication on/off, RLC entity selection, and high priority traffic preemption. The Near-RT RIC may generate or update E2 policies. The Near-RT RIC sends reports to the Non-RT RIC for evaluation and optimization as needed. The E2 nodes such as the O-CU and the O-DU provide network metrics to the SMO. The E2 nodes also provide UE measurement reports, performance metrics, and cell resource utility to the Near-RT RIC through E2 interface. The E2 nodes enforce the received E2 control policy.

1.7.18 BBU Pooling to Achieve RAN Elasticity

This use case focuses on cloudified baseband units (BBUs) that are deployed on a common hardware pool, centralized at a given cloud location. Such BBU pooling enables flexible mapping between O-RUs and BBUs, achieving RAN elasticity and providing benefits such as reduced costs, increased resiliency, increased agility of the RAN to adapt to traffic demands, and increased RAN utilization efficiency. Both the O-CUs and the O-DUs can be pooled.

In an example implementation of O-DU pooling, O-DU resources for multiple cell sites are pooled at a given centralized location to create an O-DU pool. The Cloud Access Switch (CAS) at the O-DU pool aggregates traffic from the multiple fronthaul gateways and routes the traffic to a suitable O-DU blade to realize load balancing. For the O-DU pooling use case, it is assumed that an open fronthaul exists with O-RAN-defined 7.2x split. Based on the association between O-RUs and O-DUs, and the granularity at which the traffic is assigned from the O-RUs to the O-DUs, we define the following.

Three classes of O-DU pooling are defined based on the association between O-RUs and O-DUs and the granularity at which the traffic is assigned from the O-RUs to the O-DUs. In Class 0 pooling, an O-RU is assigned to a single-specific O-DU statically. Reassigning an O-RU to a different O-DU requires significant manual operations and is performed infrequently and during specific maintenance windows. In Class 1 pooling, a set of O-RUs is initially assigned to a single O-DU for a specific period, but these O-RUs can be reassigned to a different O-DU at any time using SMO-assisted orchestration. Finally, in case of Class 2 pooling, a set of O-RUs is assigned to a set of O-DUs, and a cluster aware scheduler dynamically distributes subsets of traffic from one O-RU to the O-DU resources within the O-DU pool.

1.7.19 Integrated SON Function

A self-organizing network (SON) includes self-configuration (e.g., configuration of parameters and automated interface setup), self-optimization (e.g., optimization of parameters), and self-healing (i.e., recovery from a failure). When a SON algorithm is executed in the 3GPP management system, it is called centralized SON. When a SON algorithm is executed in the NF layer, it is called distributed SON. Finally, when part of the SON algorithm is executed at the NF layer and part of the SON algorithm is executed at the management layer, it is called hybrid SON.

This SON use case focuses on following SON functions: (i) initial PCI allocation, conflict resolution, and Automatic Neighbor Relations (ANR); (ii) mobility load balancing (MLB) and mobility robustness optimization (MRO); (iii) RACH optimization; (iv) energy savings (ES); and (v) Coverage and Capacity Optimization (CCO).

In centralized SON deployments, one challenge is that the operator may need to choose the same vendor for both the SON solution and the RAN nodes in case of centralized SON. In the case of distributed SON, the interface between the SON and RAN nodes could be proprietary and internal to the vendor's product. The performance of such vendor product-specific SON solutions affects overall KPIs, because typical 5G deployments utilize products from multiple vendors.

This use case proposes the realization of SON algorithms are as rApps or as xApps or as a combination of rApps and xApps. For example, the A1 and E2 and optionally O1 interfaces are utilized. SON functions operating with long time scales such as initial PCI, initial ANR, and ES can be implemented at the Non-RT RIC using an rApp or in the management entities like EMS, NMS, or at the SMO. In contrast, SON functions that need to operate with short time scales are realized at the Near-RT RIC using one or more xApps of RT-RIC.

1.7.20 Shared O-RU

The current O-RAN architecture associates a given O-RU node with a single O-DU node. Such constraint limits the ability of fronthaul systems to be used effectively. For example, use cases such as RAN sharing and multivendor network slices involve a common O-RU that works with multiple O-DUs. A shared O-RU provides enhanced scalability and/or enhanced node availability in fronthaul systems, where multiple O-DU nodes deployed by a single operator are connected to a shared O-RU. When multiple O-DU nodes are deployed by different operators, the benefit of a shared O-RU is enhanced sharing capabilities across operators.

The initial phase of the proposed O-RU sharing is to enhance the current open fronthaul architecture to enable a common O-RU to work with multiple O-DUs with static allocation of resources to distinct O-DUs. This initial phase supports sharing of the O-RU by O-DUs of (i) a single-operator case or (ii) multiple operators.

1.7.21 Energy Saving

Energy consumption of the RAN is important to network operators to manage costs. RAN ES depends on suitable planning and configuration. Due to the varying nature of traffic load and user mobility, optimization of energy consumption of the RAN is quite complex, and it is entirely possible that the RAN equipment would consume a significant amount of energy even when traffic is light.

The motivation behind this use case is to leverage O-RAN's AI/ML capabilities and open interfaces to introduce optimized ES solutions that switch off/on different network components at different time scales when appropriate. The ES use case is divided into *three sub-use cases* based on the time scale of the control and the controlled system.

The first sub-use case of "carrier and cell switch off/on ES" operates in non-RT. Its goal is to reduce O-CU/DU/RU power consumption by switching off/on one or more carriers or a cell of a given technology. Implementation-specific AI/ML-assisted solutions in the Non-RT RIC can be used to control the traffic load of the carriers and the cell and determine when to switch off/on one or more carriers or a cell using O1 and/or open fronthaul management plane (M-plane) parameter configurations. Suitable traffic steering is also used in conjunction with off/on switching for service continuity and QoS. The carrier switch off/on can be carried out using (i) a hibernate mode where the radio's power amplifiers remain on but with minimum electric current draw or (ii) a complete switch off mode.

The second sub-use case of "RF channel switch off/on ES" operates in non-RT or near-RT. This sub-use case aims to reduce power consumption of an O-RU with massive MIMO by switching off/on certain RF channels. An rApp or xApp, with assistance from AI/ML algorithms triggers switching off/on certain RF channels, based on information such as traffic load, user location, and mobility. The number of RF channels, the maximum number of spatial multiplexing layers, and number of MU-MIMO users can be reduced. The O-RU can be reconfigured using the open fronthaul M-plane from an E2 node or SMO.

The last sub-use case of "Advanced Sleep Mode ES" operates in near-RT and aims to reduce power consumption by partially switching off O-RU components. The Near-RT RIC configures suitable cell parameters (e.g., the SSB periodicity) based on traffic load, user service type, and energy efficiency measurements.

1.7.22 MU-MIMO Optimization

MU-MIMO beamforming can be optimized using the Near-RT RIC for both DL and UL and TDD and FDD, enhancing user performance and cell performance for a wide range of vehicular speeds. The Near-RT RIC collects information from the E2 nodes; selects users; determines suitable time and frequency resources, precoding coefficients and modulation and coding schemes for selected users; and sends control messages to the E2 nodes using the E2 interface. Specifically, the Near-RT RIC collects UE states from the O-CU and configures the O-CU/O-DU to report RRC parameters and UL/DL traffic and radio channel information and provides users and transmission parameters for selected users to the O-DU over E2 interface. The O-DU uses the information received from the Near-RT RIC to schedule MU-MIMO transmissions.

1.7.23 Sharing Non-RT RIC Data with the Core

Cellular networks are becoming more complex with 5G being significantly more complex than 4G LTE. Increased network complexity necessitates complex analytics to be available in the RAN as well as the core network. For example, the 3GPP-defined network data analytics function (NWDAF) provides analytics to other 5GC NFs so that they can optimize their processing.

The 3GPP has defined an analytics logical function (AnLF) in support of the NWDAF, which can use AI/ML algorithms to generate analytics information useful to 5GC NFs.

rApps can produce analytics for multiple data consumers, such as other rApps, Near-RT RIC/xApps, general external entities (e.g., an application function [AF]). This use case seeks to add to this list the 5G Core NFs. In such a multiconsumer environment, it is a good architectural practice to avoid coupling of the data producer with the data consumer, in fact even hiding from the producer the identity of its data consumers or even whether the data consumers are single or multiple. In keeping with this principle, rApps can structure the information content of the data published over R1 such that it is “generally useful” and not customized for a specific data consumer.

This use case envisions a scenario where the Non-RT RIC-produced analytics are useful to the 5GC. Hence, this use case proposes a mechanism whereby such analytics can be shared with the 5GC. Two options (called Option 1 and Option 2) for the SMO/Non-RT RIC to perform the role of a 5GC NWDAF without a separate intervening NWDAF are suggested by this use case. Furthermore, two options (called Option 3 and Option 4) for the SMO/Non-RT RIC to interact with a separate intervening NWDAF are also suggested.

Option 1. SMO/Non-RT RIC Direct to 5GC Consumer with façade of an NWDAF. In this approach, the SMO/Non-RT RIC exposes the analytic directly to the consuming 5GC NF, acting as an NWDAF relative to such 5GC NF. The AnLF functionality exists only in the Non-RT RIC. Option 1 requires an rApp that makes its data available to 5G core NFs and registers its R1 data type as an analytic ID. The 5GC NF discovers the appropriate “NWDAF” (associated with SMO/Non-RT RIC) through the UDM for UE-related analytics and the NRF for non-UE-related analytics.

Option 2. SMO/Non-RT RIC to DCCF/MFAF with façade of an NWDAF. In this approach, the SMO/Non-RT RIC shares its analytics with a DCCF/MFAF that makes those analytics available to 5GC NFs. Like Option 1, the AnLF functionality is not duplicated but exists only in the Non-RT RIC. The SMO/Non-RT RIC acts as an NWDAF toward the DCCF/MFAF to leverage the NWDAF services for data subscription and delivery. The DCCF/MFAF receives analytics requests from 5GC NFs, forwards them to the SMO/Non-RT RIC, receives analytics from the SMO/Non-RT RIC, and provides these analytics to the 5GC NFs. In this approach, the 5GC NFs only need to discover via the NRF the DCCF/MFAF instances. The DCCF/MFAF contacts suitable NWDAF façade instances.

Option 3. SMO/Non-RT RIC to Separate NWDAF with façade of OA&M. In this approach, the SMO/Non-RT RIC shares its analytics with a separate NWDAF that provides these analytics to 5GC NFs. The AnLF functionality is duplicated in the Non-RT RIC and the NWDAF. In this approach, the SMO/Non-RT RIC acts as the OAM system toward the NWDAF. The separate NWDAF receives analytics requests from 5GC NFs, determines the source data to be collected, constructs analytics, and forwards the analytics to 5GC NFs.

Option 4. SMO/Non-RT RIC to Separate NWDAF with façade of RAN NF. In this approach, as in Option 3, the SMO/Non-RT RIC shares its analytics with a separate NWDAF that provides these analytics to 5GC NFs. Also, like Option 3, the AnLF functionality is duplicated in the Non-RT RIC and the NWDAF. However, the SMO/Non-RT RIC acts as a RAN NF instead of the OAM system in Option 3. The SMO/Non-RT RIC provides performance management OAM relative to the NWDAF. As in Option 3, the separate NWDAF receives analytics requests from 5GC NFs, determines the source data to be collected, constructs analytics, and forwards the analytics to 5GC NFs.

1.7.24 Industrial Vision SLA Assurance

Industrial vision is an image recognition technology that uses machines instead of human eyes for target measurements and evaluations to facilitate decision-making. It is suitable for automatic inspection, work piece processing and assembly automation, and monitoring and control of production process. The industrial vision system obtains images of the products being tested or processed and transmits these images to the vision server for detection and providing the feedback with the results of a suitable analysis.

O-RAN enables adaptive prescheduling parameters configuration to efficiently meet the QoS requirements of this use case. Key components for the industrial vision use case include the Non-RT RIC, Near-RT RIC, E2 nodes, industrial cameras, and the industrial vision server. Non-RAN data such as production line and camera information and data transmission information are transmitted to Non-RT RIC as enrichment information. With enrichment information from the vision server, the Near-RT RIC dynamically calculates prescheduling parameters (e.g., prescheduling data size, prescheduling period, and prescheduling start time) to meet the requirements of the industrial vision use case so as to configure the E2 node via E2 control/policy service with these prescheduling parameters as scheduling recommendations. The E2 node can synchronize the camera data UL resource allocations with work piece arrival times through dynamic prescheduling.

1.7.25 Non-Public Network (NPN) RAN-Sharing via Midhaul for Multioperator Coverage

This use case enables multiple operators to securely and cost-effectively share a RAN in the case of Non-Public Network (NPN) small cell (e.g., microcell or picocell) deployments in private spaces such as buildings, airports, hospitals, and hotels. Both the NPN owner and the MNO derive the benefits. The NPN owner obtains better coverage and better return on investment (ROI), while the MNO can extend its coverage without spending CapEx. O-RAN can facilitate the implementation of this use case by creating multivendor interoperability profiles for the 3GPP-defined F1 interface.

In an example architecture, a hierarchical cell structure can be used, where the 3GPP-defined F1 interface (i.e., the midhaul) and the 3GPP feature of the Integrated Access and Backhaul (IAB) are utilized to integrate the NPN-owned small cells with the macrocells of MNOs. A new NPN-RAN sharing manager function is defined in the O-RAN SMO and RAN-orchestrator to divide the NPN-RAN into multiple RAN partitions through the distribution of NPN-RAN PRBs among partner MNOs. Specifically, the shared NPN-RAN O-DU performs PRB distribution for NPN O-RUs and connects to the partner MNO O-CUs. The NPN-RAN spectrum is partitioned in a suitable manner to support partner MNOs.

1.7.26 Interference Detection and Optimization

In scenarios such as spectrum sharing between 4G and 5G, heterogeneous networks, and ultra-dense networks, effective interference management can increase the spectrum utilization efficiency and enhance performance. This use case exploits UE and network level data collection and implements interference detection, interference relationships construction, and interference optimization to optimize overall O-RAN performance.

The Near-RT RIC enables UE-level interference detection through proper configuration of the reference signal for the serving cell and relevant neighboring cells for the UE. The Near-RT RIC determines the E2 control or policy containing resource allocation for the reference signal for the configuration of the UE based on the interference detection-related A1 policy specified by the Non-RT RIC.

The Near-RT RIC utilizes QoS metrics from the SMO and network-level measurement data to construct interference relationships (e.g., interference graphs) that describe the interference relationships among UEs, groups of UEs, or RAN slices.

In one interference optimization approach, the Near-RT RIC optimally allocates radio resources to UEs, groups of UEs, RAN slices, or cells based on the interference relationships. The Non-RT RIC provides the interference optimization-related A1 policy to the Near-RT RIC. The Near-RT RIC obtains suitable measurement data through E2 interface and creates an interference optimization policy for the serving cell and relevant intra-frequency neighboring cells by allocating suitable radio resources. Note that interference optimization can be realized via slow loop optimization by the Non-RT RIC or fast loop optimization by the Near-RT RIC depending on traffic trends or use cases.

1.8 O-RAN: Accomplishments, Technical Priorities, and Potential Future Work

The current state of O-RAN is summarized in the form of accomplishments and technical priorities in Sections 1.8.1 and 1.8.2, respectively. Section 1.8.3 comments on the potential areas on which the O-RAN ALLIANCE may need to work on as cellular technologies advance from 5G to 5G-Advanced and 6G.

1.8.1 O-RAN Accomplishments

O-RAN has carried out significant work in its three workstreams—specifications, testing and integration, and OSC. As mentioned earlier, the O-RAN specifications work is carried out by O-RAN WGs with each WG focusing on a specific aspect of O-RAN. Release 3 versions of specifications are available from July 2023. As mentioned earlier, the testing and integration WG focuses its efforts on the overall approach for testing and integration and coordinates test specifications across various WGs. Several PlugFests have been held by the O-RAN ALLIANCE to validate compliance of products to the O-RAN specifications and interoperability between RAN vendors. The OSC typically releases new software every six months. The OSC software is aligned with the architecture and specifications released by the O-RAN ALLIANCE WGs. The OSC software is available publicly on GitHub. The first OSC software release, Amber, occurred in December 2019. As of early 2024, the first release occurred in December 2023.

1.8.2 O-RAN Technical Priorities

The O-RAN ALLIANCE has published three releases for *technical priorities* as of early 2024 (O-RAN ALLIANCE 2024b). The overall goal of these priorities is to accelerate the development of competitive O-RAN solutions in Europe and other regions to accelerate the global adoption of O-RAN. The technical priorities are the priorities for O-RAN solutions determined by the signatories (e.g., Deutsche Telekom, Orange, Telefónica, TIM, and Vodafone). The technical priorities provide guidance to the RAN supplier industry to accelerate market deployments.

These priorities serve as an input into TIP's O-RAN Release Framework to develop requirements as the basis for certification to promote an efficient supply chain.

Release 1 technical priorities were published in June 2021, and Release 2 technical priorities were published in March 2022. Release 3 updated Release 1 and Release 2 documents of technical priorities in April 2023. Release 1 focused on the main scenarios and technical requirements for the building blocks of a multivendor RAN. Release 2 focused on intelligence, orchestration, transport and cloud infrastructure, and energy efficiency goals to support a sustainable O-RAN. The focus of Release 3 is on developing additional requirements on the SMO and the RIC. Furthermore, significant enhancements are targeted for the cloud infrastructure, the O-CU, the O-DU, and the O-RU. Release 3 also focuses on details of security topics and the challenges of disaggregation in O-RAN.

1.8.3 Role of O-RAN Beyond 5G

The 3GPP is enhancing 5G specifications to introduce new features and capabilities and enhance existing features and capabilities compared to Release 15 Phase 1 5G and Release 16 Phase 2 5G. For example, the 3GPP introduced the nonterrestrial network (NTN) with a transparent payload in Release 17 and may consider the NTN with a regenerative payload in Release 19. The 3GPP has envisioned mobile IAB. Such new architectures can impact O-RAN design and deployments. Implementation of network slicing for diverse deployment scenarios can also affect the O-RAN design, configurations, and operational guidelines for optimal performance.

Organizations around the globe such as ATIS in North America are defining a vision and requirements for 6G (NGA 2022). 6G is anticipated to be deployed in 2030s. 6G is expected to be AI native and sharing native. In other words, 6G is expected to have built-in AI right from the design stage. Similarly, 6G would support sharing natively by supporting features such as spectrum sharing. 6G is also expected to have an even more advanced radio interface than 5G NR. For example, 6G may use RISs, which would alter the overall RAN design. See Section 8.2 in Chapter 8 for a brief overview of 6G.

In summary, O-RAN would need to keep up with evolving cellular technologies including 5G-Advanced and 6G. As mentioned earlier in the chapter, the O-RAN ALLIANCE has created a focus group called nGRG to monitor developments in 6G and future standards and carry out research of open and intelligent RAN principles so that suitable actions can be taken in the O-RAN ALLIANCE to keep O-RAN relevant to future technologies beyond 5G.

1.9 Major Takeaways of the Chapter

O-RAN is an emerging trend to deploy 5G RANs. O-RAN can also be used to deploy 4G LTE. Here are the major takeaways of this chapter.

- Cellular technologies have evolved from analog 1G technology to high-performance and flexible 5G technology. 5G is a transformational cellular technology that is expected to transform multiple industries. O-RAN is expected to play a significant role in 5G RAN deployments.
- Key components of a cellular system are the UE, the RAN, and the core network.
- RAN architectures have evolved from highly centralized architectures in 1G, 2G, and 3G to distributed architectures in 4G and 5G. Furthermore, 5G introduced disaggregation of the gNB such as the gNB-CU and the gNB-DU to distribute the radio protocol stack protocols. The gNB-CU

can be further disaggregated into the gNB-CU-CP and the gNB-CU-UP to separate CP and UP functions and provide additional deployment flexibility.

- The O-RAN ALLIANCE was founded in February 2018 by AT&T, China Mobile, Deutsche Telekom, NTT DOCOMO, and Orange. It was established as a German entity in August 2018. Today's O-RAN ALLIANCE is a global community of MNOs, vendors, and research and academic institutions.
- The O-RAN ALLIANCE has defined the O-RAN architecture that emphasizes openness and intelligence to cost-effectively deploy RANs and optimize RANs using AI/ML mechanisms inherently supported in the O-RAN architecture. For example, Non-RT RIC and Near-RT RIC provide intelligence to O-RAN by making use of AI/ML model to optimize various RAN operations.
- Deployments of O-RAN require the use of standards and software frameworks developed by organizations such as 3GPP, ETSI, IETF, CNCF/Kubernetes, OpenStack, and ONAP.
- Major forces driving O-RAN include RAN vendor diversity, complexity of RAN, accelerated innovations, and geopolitical environment.
- O-RAN helps optimize various RAN operations. Examples of O-RAN use cases are context-based dynamic handover management for V2X, flight path-based dynamic UAV radio resource allocation, traffic steering, massive MIMO optimization, RAN sharing, RAN slice management, DDS, local indoor positioning in RAN, signaling storm protection, CPM, industrial IoT optimization, BBU pooling, integrated SON functions, shared O-RU, ES, sharing Non-RT RIC data with the core network, and industrial vision SLA assurance.
- O-RAN has conducted significant work in its three workstreams—specifications, testing and integration, and OSC. The O-RAN specifications work is conducted by O-RAN WGs with each WG focusing on a specific aspect of O-RAN.
- The testing and integration WG focuses its efforts on the overall approach for testing and integration and coordinates test specifications across various WGs. Several PlugFests have been held by the O-RAN ALLIANCE to validate compliance of products to the O-RAN specifications and interoperability between RAN vendors.
- The OSC typically releases new software every six months. The OSC software is aligned with the architecture and specifications released by the O-RAN ALLIANCE WGs. The OSC software is available publicly on GitHub.
- The overall goal of O-RAN's technical priorities is to accelerate the development of competitive O-RAN solutions in Europe and other regions to accelerate the global adoption of O-RAN. The technical priorities provide guidance to the RAN supplier industry to accelerate market deployments.
- nGRG is an O-RAN focus group that monitors developments in 6G and future standards and conducts research of open and intelligent RAN principles to keep O-RAN relevant to future technologies beyond 5G.

References

3GPP TS23.501 (2024). *System Architecture for the 5G System (5GS)*. The Third Generation Partnership Project. https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144 (accessed 18 July 2024).

Bonneau, M. and Keeney, J. (2023). *O-RAN A1 Policies in ONAP*. April 2023, https://wiki.onap.org/display/DW/O-RAN+A1+Policies+in+ONAP (accessed 18 July 2024).

Kubernetes (2024). *Overview*. https://kubernetes.io/docs/concepts/overview/ (accessed 18 July 2024).

(NGA (2022). *Roadmap to 6G. Next G ALLIANCE Report*. February 2022, https://www.nextgalliance.org/wp-content/uploads/2022/02/NextGA-Roadmap.pdf (accessed 18 July 2024).

O-RAN ALLIANCE (2018). O-RAN: Towards an Open and Smart RAN. White Paper, October 2018, https://mediastorage.o-ran.org/white-papers/O-RAN.White-Paper-2018-10.pdf (accessed 18 July 2024).

O-RAN ALLIANCE (2020). O-RAN Use Cases and Deployment Scenarios: Towards Open and Smart RAN. White Paper, February 2020, https://mediastorage.o-ran.org/white-papers/O-RAN.WG1.Use-Cases-and-Deployment-Scenarios-White-Paper-2020-02.pdf (accessed 18 July 2024).

O-RAN ALLIANCE (2021). O-RAN Minimum Viable Plan and Acceleration towards Commercialization. White Paper, June 2021, https://mediastorage.o-ran.org/white-papers/O-RAN.Minimum-Viable-Plan-and-Acceleration-towards-Commercialization-white-paper-2021-06.pdf (accessed 18 July 2024).

O-RAN ALLIANCE (2024a). *About Us*. https://www.o-ran.org/about (accessed 18 July 2024).

O-RAN ALLIANCE (2024b). *O-RAN Ecosystem Resources*. https://www.o-ran.org/ecosystem-resources (accessed 18 July 2024).

O-RAN ALLIANCE (2024c). *O-RAN Architecture Description*. O-RAN.WG1.OAD-R003-v09.00. The O-RAN ALLIANCE. https://specifications.o-ran.org (accessed 18 July 2024).

O-RAN ALLIANCE (2024d). *Use Cases Analysis Report*. O-RAN.WG1. Use-Cases-Analysis-Report-v14.00. The O-RAN ALLIANCE, March 2024, https://specifications.o-ran.org (accessed 18 July 2024).

O-RAN ALLIANCE (2024e). *Use Cases Detailed Specification*. O-RAN.WG1.Use-Cases-Detailed-Specification-R003-v14.00. The O-RAN ALLIANCE, March 2024, https://specifications.o-ran.org (accessed 18 July 2024).

O-RAN ALLIANCE (2024f). *Security Protocols Specifications*. O-RAN.WG11.Security-Protocols-Specification.O-R003-v05.00. The O-RAN ALLIANCE, https://specifications.o-ran.org (accessed 18 July 2024).

O-RAN ALLIANCE (2024g). *O2 Interface General Aspects and Principles*. O-RAN.WG6.O2-GA&P-R003-v03.00. The O-RAN ALLIANCE. https://specifications.o-ran.org (accessed 18 July 2024).

OpenStack (2024). *Software*. https://www.openstack.org/software/ (accessed 18 July 2024).

Peterson, L., Oguz Sunay, O. and Bruce Davie, B. (2023). *Private 5G: A Systems Approach*. Systems Approach LLC*, 2023*, https://5g.systemsapproach.org/ran.html (accessed 18 July 2024).

Rethink (2024). *ONAP's New Release Focuses on 5G Slicing and O-RAN Integration*. June 2020, https://rethinkresearch.biz/articles/onaps-new-release-focuses-on-5g-slicing-and-o-ran-integration/ (accessed 18 July 2024).

Tripathi, N. and Reed, J. (2019). *5G Cellular Communications—Journey and Destination*. 1st edn. The Wireless University. https://thewirelessuniversity.com/ (accessed 18 July 2024).