

ARGOS: Anomaly Recognition and Guarding through O-RAN Sensing

Stavros Dimou
Northeastern University
Boston, MA, USA
dimou.s@northeastern.edu

Guevara Noubir
Northeastern University
Boston, MA, USA
g.noubir@northeastern.edu

Abstract—Rogue Base Station (RBS) attacks, particularly those exploiting downgrade vulnerabilities, remain a persistent threat as 5G Standalone (SA) deployments are still limited and User Equipment (UE) manufacturers continue to support legacy network connectivity. This work introduces *ARGOS*, a comprehensive O-RAN compliant Intrusion Detection System (IDS) deployed within the Near Real-Time RIC, designed to detect RBS downgrade attacks in real time, an area previously unexplored within the O-RAN context. The system enhances the 3GPP KPM Service Model to enable richer, UE-level telemetry and features a custom xApp that applies unsupervised Machine Learning models for anomaly detection. Distinctively, the updated KPM Service Model operates on cross-layer features extracted from *Modem Layer 1 (ML1)* logs and *Measurement Reports* collected directly from Commercial Off-The-Shelf (COTS) UEs. To evaluate system performance under realistic conditions, a dedicated testbed is implemented using Open5GS, srsRAN, and FlexRIC, and validated against an extensive real-world measurement dataset. Among the evaluated models, the Variational Autoencoder (VAE) achieves the best balance of detection performance and efficiency, reaching 99.5% Accuracy with only 0.6% False Positives and minimal system overhead.

Index Terms—5G, O-RAN, xApps, ML, RBS, IDS

I. INTRODUCTION

The rise of 5G mobile networks represents a major shift in telecommunications, opening the door to innovative applications through faster connectivity and minimal delay. As the telecom industry nears the midpoint of 5G adoption, a growing number of Mobile Network Operators (MNOs), particularly in Asia and North America, are focusing on network densification and transitioning to 5G Core Standalone (SA) architectures [1]. This transition is intended to unlock the full potential of 5G SA, with global subscriptions expected to reach 3.6 billion by 2030 [2]. Nonetheless, studies indicate that LTE and earlier mobile networks will continue to be used globally, suggesting that 5G SA will coexist with legacy networks well into the next decade, especially as user devices remain compatible with older cellular technologies. This raises the problem of backward compatibility with older generations, which continue to expose legacy vulnerabilities. Downgrade attacks take advantage of this, forcing devices to connect through less secure legacy networks, compromising the connection’s integrity. Such attacks are commonly carried out by adversaries using Rogue Base Stations (RBS), which operate as International Mobile Subscriber Identity (IMSI) catchers

[3]–[8]. Thus, as we move beyond 5G and towards 6G, networks must be equipped to detect and classify malicious entities and traffic within their vicinity.

From both operational and security perspectives, a key advancement in 5G is the adoption of the software-defined Open Radio Access Network (O-RAN) architecture, which introduces a new level of programmability to traditional cellular infrastructures. O-RAN redefines the traditionally monolithic and vendor-proprietary RANs by introducing a disaggregated, modular architecture that promotes openness, interoperability, and programmability, as illustrated in Figure 2. Embracing the principles of Software-Defined Networking (SDN), O-RAN enables centralized control and dynamic network optimization across the RAN. Central elements of this architecture are the RAN Intelligent Controllers (RICs), located in the Control Plane, both supporting modular, “plug-and-play” applications, also known as xApps and rApps. Recent studies have increasingly leveraged the RIC for both network optimization and the implementation of Intrusion Detection Systems (IDS) and Intrusion Prevention Systems (IPS), addressing a range of attack vectors.

However, no prior work has explicitly addressed downgrade attacks by proposing a practical solution aligned with current and future cellular network deployments and smartphone manufacturing constraints. Moreover, much of the existing research relies on synthetic data or constrained measurements, failing to reflect the characteristics of existing cellular network deployments. To address these gaps, we introduce *ARGOS*¹, the first comprehensive system integrated with O-RAN for detecting RBS attempting to launch downgrade attacks, combining an IDS xApp with an enhanced telemetry collection mechanism within the Near Real-Time RIC (NearRT-RIC). It collects diverse physical layer (PHY-layer) indicators, such as Reference Signal Received Power (RSRP), Reference Signal Received Quality (RSRQ), and Signal-to-Interference-plus-Noise Ratio (SINR), extracted from User Equipment (UE) Measurement Reports and Modem Layer 1 data, and leverages Machine Learning (ML) models to detect the presence of malicious cells within the surrounding area. Furthermore, to address the limitations of restricted real-world measurements, we evaluate

¹*Argos Panoptes*, the “all-seeing” giant in Greek mythology, had a hundred eyes and symbolized perpetual vigilance—reflecting *ARGOS*’s continuous monitoring against RBS threats.

our system using an extensive, real-world dataset collected from over 10 areas across two major U.S. cities, covering two MNOs and four Commercial Off-The-Shelf (COTS) UEs. Finally, to assess the performance of our framework, we built a custom testbed using Open5GS [9], srsRAN [10], and FlexRIC [11] as reference platforms for the core network and O-RAN infrastructure. Our system achieves up to 99.5% Accuracy and 96.7% Precision, demonstrating both its reliability in detecting rogue cells over time and its robustness against false alarms. The remainder of the paper is organized as follows: Section II reviews related work, while Section III provides the necessary background. Section IV outlines the threat model and architectural design of the proposed IDS. Section V presents the benchmarking and experimental results, whereas Section VI details the ethical considerations adhered to throughout this work. Finally, Section VII concludes the paper with implications for future work.

II. RELATED WORK

5G introduces substantial security improvements over previous generations; however, it still inherits vulnerabilities that persist from legacy systems such as LTE [12], enabling various attacks, such as IMSI Catching. IMSI catchers, implemented through RBS, have been widely studied across all cellular generations [5], [13]. These attacks allow adversaries to actively impersonate legitimate base stations, prompting UEs to reveal their IMSI in plaintext, leading to subscriber identity exposure, tracking, and localization [14]–[16]. Beyond IMSI disclosure, RBS facilitates a range of threats, drawing attention from both academic and standardization communities [13]. The 3rd Generation Partnership Project (3GPP) introduced an optional RBS detection framework within its technical specifications [17], and further dedicated an entire technical report to this issue [18]. The report identifies critical RBS threat scenarios and introduces mitigation strategies, including enhanced UE-Measurement Reporting. However, these proposals lack concrete implementation strategies or timelines for integration into the specification.

The ongoing risk posed by RBS is further amplified by the continued reliance on legacy mobile networks like LTE [19]. This enables bidding-down attacks, downgrading UEs from 5G to less secure generations, and exploiting weaker authentication and encryption protocols [3], [7], [20]. In [3], researchers demonstrate a downgrade attack from 5G-SA to 2G on commercial networks. Similarly, [7] reveals a vulnerability in LTE where UEs reveal their capabilities before establishing RRC security, allowing adversaries to intercept and manipulate these messages to initiate a downgrade.

Given the limitations of current defenses, recent work has explored O-RAN as a promising path forward. Its architecture has motivated extensive research into leveraging O-RAN as a foundation for IDS [21]–[29], enabling intelligent threat monitoring within the RAN infrastructure. A prominent example is 5G-SPECTOR [21], a framework targeting Layer 3 protocol exploit detection, utilizing a security audit and xApp. Similarly, [22] presents an AI/ML-driven IDS that also functions

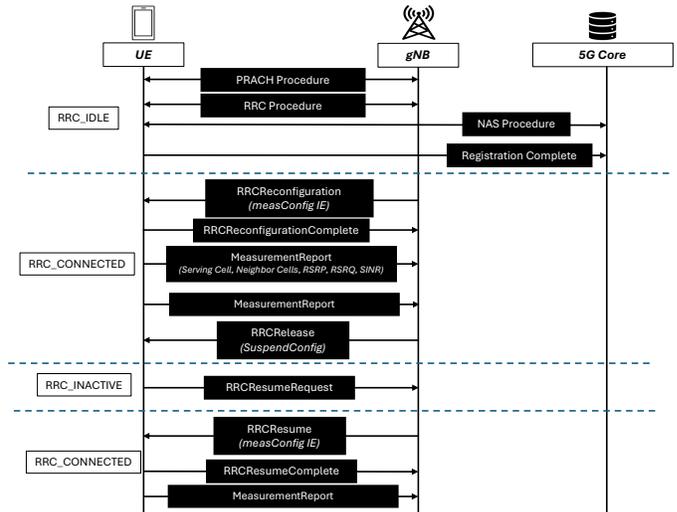


Fig. 1. RRC state transitions and Measurement Reports.

as a real-time resource allocator. In [23], the authors propose UE-level detection of RBS by training ML models on signal stability metrics within the NearRT-RIC and distributing them back to the UEs. A related effort, [27], focuses on jamming detection using UE-reported Channel Quality Indicator (CQI) and RSRP values, employing the Kolmogorov-Smirnov test to flag anomalies. [24] deploys an IDS within the NearRT-RIC security module, targeting model poisoning attacks in ensemble learning setups. [28] similarly uses cross-domain AI models embedded in xApps, combining data from both the RAN and transport networks. Pushing detection to lower layers, Det-RAN [26] proposes a real-time IDS at the gNB-DU, leveraging PHY-layer features such as IQ samples and CSI. Meanwhile, [25] focuses on the Open Fronthaul (O-FH) interface, applying deep learning to detect and mitigate DDoS attacks. Finally, 6G-XSec [29] introduces a two-stage IDS combining unsupervised anomaly detection via xApp with a Large Language Model (LLM) for threat interpretation.

Although these studies present insights for IDS, none has explicitly focused on detecting RBS, particularly in the context of downgrade attacks. Moreover, they lack actual implementation, relying on simulations and artificial datasets that fail to capture the constraints of real-world mobile networks and UE behavior. In this work, we address these gaps by designing, implementing, and evaluating a real-time IDS system within O-RAN, utilizing a real-world setup. The system is evaluated using real-world measurements collected directly from COTS UEs operating on public commercial networks across multiple MNOs. The system accurately detects malicious cells within a given area, offering a practical implementation that advances ongoing research in RBS detection.

III. 3GPP/O-RAN TELEMETRY MECHANISMS

In this section, we present the necessary background, introducing the architecture of LTE/5G networks, the telemetry mechanisms underpinning the Measurement Reporting pro-

cess, and the key components and interfaces that define O-RAN.

A. Cellular Network Operations

1) *5G & LTE Cellular Networks*: The architecture of 5G systems is composed of three principal entities, as depicted in Figure 1: (1) the UE, typically a smartphone equipped with a Universal Subscriber Identity Module (USIM) subscribing to commercial networks and identified by a unique user identifier known as the Subscription Permanent Identifier (SUPI), referred to as IMSI in LTE and earlier generations; (2) the gNodeB (gNB), the 5G base station operating within the RAN, which connects the UE to the MNO’s core network; and (3) the 5G Core Network (5G-CN), a service-based architecture (SBA) enabling authentication, security, and session management through different Network Functions (NFs). The gNB may interface either with a 5G Core Network (5G-CN) in the 5G SA architecture or with an LTE Evolved Packet Core (EPC) in the 5G Non-Standalone (NSA) architecture.

The initial procedure for UE connectivity begins with cell attachment and network registration. The UE performs initial cell selection by detecting and decoding System Information Block (SIB) messages broadcast by nearby gNBs. Subsequently, it initiates random access via the Physical Random Access Channel (PRACH) to achieve uplink synchronization. Upon successful random access, the establishment, maintenance, and release of radio connections are managed by the Radio Resource Control (RRC) protocol [30]. To establish an RRC connection, the UE sends an RRCSetupRequest message. If the gNB accepts, it responds with an RRCSetup message providing configuration information. The handshake is then completed with an RRCSetupComplete message. Following the RRC establishment, the Non-Access Stratum (NAS) procedures commence to facilitate UE registration with the core network. The UE exchanges NAS messages with either the Access and Mobility Management Function (AMF) in the 5G-CN or the Mobility Management Entity (MME) in the EPC. The NAS procedure is initiated with a Registration Request (in 5G) or Attach Request (in LTE), containing the UE’s temporary (TMSI) or permanent identifiers, such as the Subscription Concealed Identifier (SUCI) in 5G or the IMSI in earlier generations. Authentication and security procedures follow, involving the Authentication and Key Agreement (AKA) protocol. Upon successful authentication, the NAS procedure concludes with a Registration Complete or Attach Complete message.

Once a secure radio connection between the UE and the network has been established, the UE enters the RRC_Connected state. In this state, the network can transmit system information to the UE through dedicated signaling, primarily using the *RRCReconfiguration* message. This procedure is used to modify an already established RRC connection by configuring various parameters [30]. Once the UE successfully acknowledges the reconfiguration process, it responds with an *RRCReconfigurationComplete* message to confirm the changes. If the network is in an idle state, the UE can optionally stay

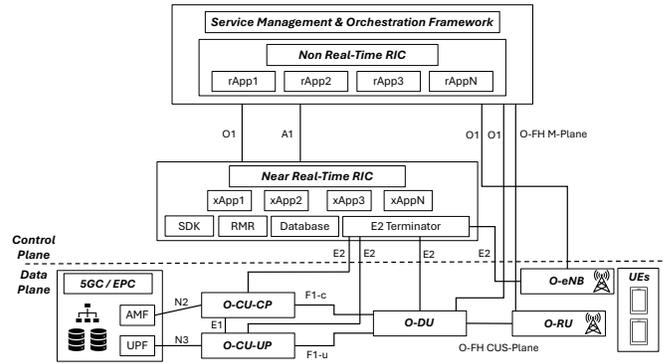


Fig. 2. Open Radio Access Network (O-RAN) architecture.

in RRC_Inactive state, as depicted in Figure 1, instead of completely releasing the RRC connection and recover it via the *RRC_Resume* procedure. During the *RRC_Resume* procedure, the exchange of messages such as *RRCResumeRequest*, *RRCResume*, and *RRCResumeComplete* occurs.

2) *UE Measurement Reports*: Following the successful establishment of the RRC connection, the network can instruct the UE to perform specific measurements through the *meas-Config* Information Element (IE), typically conveyed within the *RRCReconfiguration* [30], as shown in Figure 1. This IE defines the measurement configuration that the UE should follow, specifying which frequencies, cells, or signals to monitor. *measConfig* IE may also be included within the *RRCResume* message. The measurement configuration can direct the UE to perform intra- and inter-frequency NR measurements, defined through the *MeasObjectNR* IE, as well as inter-Radio Access Technology (RAT) measurements, including E-UTRA (LTE) measurements via the *MeasObjectEUTRA* IE and UTRA-FDD (UMTS) measurements via the *MeasObjectUTRA-FDD* IE. Depending on the configuration, measurements can be based either on Synchronization Signal/Physical Broadcast Channel blocks (SSB/PBCH) or on Channel State Information Reference Signals (CSI-RS). Measurement Reporting can be configured to occur periodically, via the *reportInterval* IE or based on event-triggered conditions. Additionally, the network may provide lists of specific cells that the UE should prioritize or cells that should be ignored. Furthermore, the network specifies the radio quantities to be included in the reports, such as RSRP, RSRQ, or SINR. Through these configurations, Measurement Reporting enables the UE to provide the network with critical information regarding the radio environment, supporting functions such as mobility management, beam selection, handover, and connection optimization.

B. O-RAN Architecture

Figure 2 illustrates the O-RAN architecture, highlighting both the data and control planes. Below, we outline the key architectural principles governing each plane.

1) *O-RAN Data Plane*: The O-RAN architecture, illustrated in Figure 2, follows the 3GPP disaggregation model [31], splitting the gNB into three main units: the Radio Unit (O-

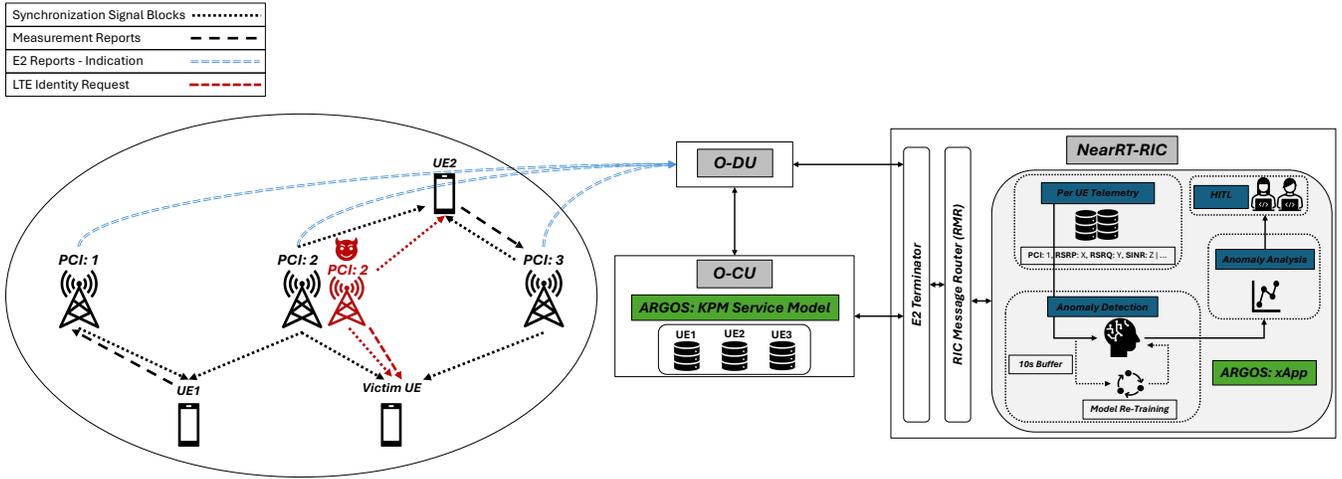


Fig. 3. Even if an RBS mimics the PCI of a legitimate cell, it cannot replicate the RF propagation profile observed by all UEs. These measurements are transmitted to the Near-RT RIC via ARGOS’s enhanced KPM Service Model, where a dedicated xApp leverages trained ML models to detect anomalies.

RU), Distributed Unit (O-DU), and Central Unit (O-CU). O-RUs, located near the antennas in the fronthaul, handle PHY-layer operations. O-DUs and O-CUs, deployed at the network edge, manage Layers 2 and 3. The O-DU oversees Medium Access Control (MAC) and Radio Link Control (RLC) functions, while the O-CU, divided into O-CU-CP (control plane) and O-CU-UP (user plane), handles RRC and forwards control/user traffic to the core network (AMF/UPF). Standardized interfaces connect these components, with F1 linking O-DU and O-CU and E1 connecting O-CU-CP to O-CU-UP.

2) *O-RAN Control Plane*: The O-RAN control plane is distinct from the data plane and centers on the RICs, which are split into NearRT-RIC and Non Real-Time RIC (NonRT-RIC). These programmable components provide centralized network visibility, enabling closed-loop control and orchestrating RAN operations. The NearRT-RIC operates on timescales from 10 milliseconds to 1 second and hosts modular control applications called xApps for different use cases. Communication between xApps and other internal components is facilitated by the RIC Message Router (RMR), which ensures efficient message routing within the NearRT-RIC. NearRT-RIC interfaces directly with the O-DU and O-CU, also referred to as E2 Nodes, via the E2 interface, allowing real-time telemetry and control. Interactions are governed by four core E2 procedures: *Report, Insert, Control, and Policy*. Each xApp implements specific E2 Service Models (E2SMs) layered over the E2 Application Protocol (E2AP). The NonRT-RIC, part of the Service Management and Orchestration (SMO) framework, operates on longer timescales (above 1 second), supporting similar control applications to xApps, known as rApps, which are higher-level applications that generate long-term policies impacting overall network behavior.

IV. ARGOS OVERVIEW

In this section we outline the threat model targeted by our system and present an architectural overview of ARGOS,

detailing UE telemetry acquisition and employed ML models.

A. Exploiting LTE Compatibility in COTS UEs

The threat model addressed in this work concerns the exploitation of plaintext IMSI transmission over the air. As discussed in Section III, during the initial NAS message exchange between the UE and the network, the UE transmits either a permanent or temporary identifier. In the absence of prior interaction or under malicious intent, the network may issue an Identity Request, prompting the UE to disclose its permanent identifier. In LTE networks specifically, the IMSI is transmitted without encryption or integrity protection, enabling adversaries to intercept it. Such behavior enables IMSI catching attacks, in which RBS masquerade as legitimate cells, broadcasting identical network identifiers with stronger signal or with a different Tracking Area Code (TAC), deceiving the UE into believing that it has entered a new tracking area. RBS can either broadcast a different Physical Cell ID (PCI) than nearby legitimate cells or carry out a more sophisticated attack by reusing the same PCI to impersonate a valid cell, as illustrated in Figure 3. Once camped at RBS, the UE is coerced to disclose its IMSI, or in some cases even its IMEI [3], enabling subsequent user tracking and localization [13]. This vulnerability stems from specifications in the 3GPP standard rather than implementation flaws, rendering all LTE-compatible devices susceptible regardless of vendor. The issue is expected to persist until UEs fully transition to support only the latest standards, disabling previous generations. As highlighted in [2], current deployment limitations of 5G-SA make it infeasible to release UEs that are exclusively 5G-compatible. Therefore, coordinated efforts between UE manufacturers and telecom vendors are essential to promote end-to-end adoption of secure, modern standards. In the interim, we propose ARGOS as a practical and deployable solution to detect and mitigate RBS effectively within existing network environments.

Algorithm 1 ARGOS Telemetry-Based Anomaly Detection

```
1: Input:  $T_u = \{([f_u], [c_u], [r_u], [q_u], [s_u], [t_u])\}$  for each  
   UE  $u = 1, \dots, N$ ; ML model  $M$ ; Anomaly threshold  $\tau$   
2: Output: Anomaly score per second  $\alpha_u(t)$ ; MSE per  $u$   
3: Connect to Near-RT RIC.  
4: Subscribe to E2 nodes using KPM Service Model.  
5: for  $u$  do  
6:   Circular Buffer  $B_u \leftarrow \emptyset$   
7: end for  
8: while true do  
9:   for  $u$  do  
10:     $B_u \leftarrow T_u$   
11:    if  $|B_u| \geq 1$  second of new telemetry then  
12:       $X_u(t) \leftarrow \text{encoded}(B_u)$   
13:       $\hat{X}_u(t) \leftarrow M(X_u(t))$   
14:       $\alpha_u(t) \leftarrow \text{MSE}(X_u(t), \hat{X}_u(t))$   
15:      if  $\alpha_u(t) > \tau$  then  
16:        Anomaly  
17:      else  
18:        Legitimate  
19:      end if  
20:    end if  
21:  end for  
22:  if  $\sum_{u=1}^N |B_u| \geq 10$  seconds then  
23:     $D \leftarrow \bigcup_{u=1}^N B_u$   
24:     $M \leftarrow \text{Train}(D)$   
25:     $\tau \leftarrow \text{GetThreshold}(M, D)$   
26:  end if  
27: end while
```

B. ARGOS Architecture

The E2 Setup procedure serves as the starting point of the architecture and is independent of any specific xApp. This procedure establishes application-level communication between E2 Nodes and the NearRT-RIC, replacing any prior configurations with the latest agreed-upon parameters [32]. Importantly, during this phase, E2 Nodes expose the types of telemetry and control capabilities they support. Following successful setup, xApps can query information about connected E2 Nodes and initiate telemetry collection via the E2 Subscription procedure [33], as shown in Figure 3. ARGOS uses only the *Report* procedure, which involves *E2 RIC Indication* messages containing telemetry from E2 nodes, sent either periodically or in response to specific trigger events. However, the system is extensible to support the remaining services introduced in Section III, enabling a shift from IDS to IPS.

To collect telemetry at the xApp level from the E2 nodes, we adopt the latest Key Performance Measurement (KPM) Service Model defined by 3GPP [34]. Several KPMs defined in the latest Service Model encapsulate Measurement Report messages from UEs, which, based on the measConfig IE, include metrics such as RSRP, RSRQ, and SINR. This Model is extended by ARGOS, incorporating both intra- and inter-

frequency measurements (5G and LTE), enabling a comprehensive view of all neighboring cells as reported directly by UEs.

Each E2 Node, particularly the O-CU handling RRC signaling, aggregates Measurement Reports per UE, identified by SUPI or SUCI, with a dedicated memory buffer assigned to each UE. Similarly, the implemented xApp maintains its own per UE memory buffers to enable continuous telemetry processing. Upon receipt at the O-CU, reports are parsed to extract per cell RSRP, RSRQ and SINR measurements. After one second of telemetry is accumulated, the data are encapsulated in E2 RIC Indication messages, structured according to the extended KPM Service Model, and sent to the xApp via the E2 interface, where anomaly detection is performed using deep learning techniques, as depicted in Figure 3. Before being passed to the models for evaluation, the received telemetry undergoes additional processing to generate per-second vectors. As shown in Algorithm 1, each per-second vector T_u encodes the presence or absence of known legitimate cells, identified by their Absolute Radio Frequency Channel Number (ARFCN) (f_u) and PCI (c_u). It includes corresponding measurements of RSRP (r_u), RSRQ (q_u), and SINR (s_u), each independently normalized. The vector also records the measurement timestamp associated with each observation. During inference, the xApp evaluates each per-second vector, generating a binary anomaly verdict along with the associated Mean Squared Error (MSE) value. To adapt to evolving network behavior, the model is retrained every 10 seconds using newly accumulated legitimate telemetry. The telemetry collection and preprocessing workflow is summarized in Algorithm 1, along with the model's inference and re-training processes.

C. Deep Learning Based RBS Detection

To ensure safe integration with commercial networks and adhere to ethical standards, outlined in Sections V and VI, ARGOS is trained and tested exclusively on legitimate, non-malicious data collected from commercial MNOs using passive observation setups. Given the absence of labeled attack data, our xApp is evaluated using four unsupervised learning models: Autoencoders, Denoising Autoencoders, Variational Autoencoders, and Isolation Forests. These models are inherently suited for anomaly detection tasks where only benign patterns are available during training.

1) *Autoencoders*: Autoencoders (AEs) are artificial neural networks designed to learn compact representations of unlabeled data. Their architecture includes an encoding function that compresses the input vector into a lower-dimensional space, $A : \mathbb{R}^n \rightarrow \mathbb{R}^p$, and a decoding function that reconstructs the original vector $A : \mathbb{R}^p \rightarrow \mathbb{R}^n$ [35]. Together, these functions aim to minimize the reconstruction error, computed using the MSE loss in Equation 1, and optimized via back-propagation to capture the input data distribution. Vectors with high MSE values are flagged as anomalous, indicating potential outliers. In ARGOS, the anomaly threshold is set after training using the 99.9th percentile of MSE values from

the training dataset. This approach minimizes the likelihood of legitimate vectors being incorrectly flagged as anomalous (False Positives), while maintaining sensitivity to suspicious patterns.

$$\mathcal{L}_{\text{MSE}} = \frac{1}{n} \sum_{i=1}^n (x_i - \hat{x}_i)^2 \quad (1)$$

AEs are well-suited for the discussed problem, based on their ability to learn the underlying patterns of per-second measurements, capturing typical combinations of cells and their associated signal characteristics. By reconstructing these vectors, the AE effectively models normal telemetry behavior, enabling the detection of deviations indicative of anomalies. Given the presence of measurement noise due to reflections and other propagation effects, we extend the baseline AE to include a Denoising AE and a Variational AE, which improve generalization and mitigate overfitting by learning robust latent representations.

2) *Denoising-Autoencoder*: Denoising Autoencoders (DAEs) are more robust variants of AEs, used for error correction. In *ARGOS*, the DAE shares the same architecture as the standard AE, with the key difference being that input training vectors are corrupted with Gaussian noise. The model is then trained to reconstruct the original, noise-free vectors. The noise process is modeled by a function $T : X \rightarrow X$, where $T(x) = x + \epsilon$ and ϵ is sampled from a Gaussian distribution $\mu_T = \mathcal{N}(0, \sigma^2)$. This method assists the network in avoiding the memorization of the input, forcing it to learn the core features of the dataset.

3) *Variational-Autoencoder*: Similar to DAEs, Variational Autoencoders (VAEs) share the same architecture as standard AEs but are grounded in the mathematical framework of Variational Bayesian (VB) methods. In VAEs, the encoder maps each input vector to a Gaussian distribution in the latent space, parameterized by a mean vector μ and a standard deviation vector σ . A latent vector is then sampled from this distribution, and the decoder attempts to reconstruct the original input. The loss function combines a reconstruction loss (MSE) and a Kullback–Leibler (KL) divergence loss, as shown in Equation 2, which regularizes the latent space by encouraging it to match a prior distribution. This probabilistic formulation helps prevent overfitting and improves generalization.

$$D_{\text{KL}}(P \parallel Q) = \int_{-\infty}^{\infty} P(x) \log \left(\frac{P(x)}{Q(x)} \right) dx \quad (2)$$

4) *Isolation Forest*: Isolation Forests are a well-established anomaly detection algorithm based on binary trees. The core idea is that anomalies, being few and different, can be isolated with fewer partitions. The algorithm recursively builds *Isolation Trees* by randomly selecting an attribute and a split value between its minimum and maximum range. Anomaly scores are derived from the path length, as anomalies typically require fewer splits to be isolated. However, Isolation Forests assume anomalies are few and different in feature space, something

Metric	Equation
Accuracy	$(TP + TN)/(TP + FP + TN + FN)$
Precision	$TP/(TP + FP)$
Recall	$TP/(TP + FN)$
F1-Score	$2 \cdot \text{Precision} \cdot \text{Recall}/(\text{Precision} + \text{Recall})$

TABLE I
PERFORMANCE METRICS USED FOR ML MODEL EVALUATION.

that may limit performance in datasets with subtle or high-density anomalies.

V. EXPERIMENTAL EVALUATION

This section presents the experimental evaluation of the proposed framework, detailing the deployed software and hardware components, along with a performance analysis of *ARGOS* from both system and ML model perspectives.

A. *ARGOS O-RAN Compliant Testbed*

To evaluate the proposed system, a controlled O-RAN-compliant testbed, including 5G SA, LTE, 5G-CN/EPC and NearRT-RIC components, is deployed. All software-based components are deployed within the same x86_64 Ubuntu 22.04.4 LTS host, equipped with 8 11th Gen Intel Core i7-1195G7 @ 2.90 GHz, 32.0 GiB RAM and 1.0 TB disk capacity. The testbed leverages version 24.10 of srsRAN, version 23.11 of srsUE, latest version of Open5GS Release-17, and latest version of the br-flexric branch of FlexRIC to emulate real-world gNB (5G-RAN), eNB (LTE RAN), 5G-CN/EPC and NearRT-RIC behavior accordingly. The RU front end of the deployed networks is hosted within 2 Ettus Research Universal Software Radio Peripheral (USRP) X310 SDR devices. One Pixel 5 COTS UE is utilized, equipped with a sysmocom SIM card programmed with PLMN identifiers matching the 5G SA deployment. The core component of *ARGOS*, our ML-based xApp, is integrated into the NearRT-RIC through FlexRIC, supporting both Python and C implementations. Communication between CU and NearRT-RIC is established over the standardized E2 interface, while FlexRIC’s internal E42 interface facilitates communication between xApp and the RIC controller.

The testbed, configured using srsRAN-provided files [36], is used to demonstrate the feasibility of the RBS Downgrade attack. This setup enables active UE Measurement Reporting and supports handover, facilitating inter-cell movement closely replicating real-world deployments. Once the adversarial eNB becomes active, it impersonates the legitimate network’s PLMN ID as well as PCI while transmitting at a higher signal strength. The UE connects to the adversarial eNB and exposes its IMSI, resulting in the successful compromise of its identity, demonstrating both the feasibility and simplicity of the attack.

B. *Real-World Data Collection*

To ensure realism beyond controlled testbed conditions, the system is evaluated using real-world data collected directly

#	E-ARFCN	Physical Cell ID	Valid RSRP	Inst RSRP [dBm]	Inst RSRQ [dBm]							
0	66736	240	RSRP_SINR	-107	-107	86	-17	81	-17	81	-49	78
1	66736	187	RSRP_SINR	-113	-106	88	-20	50	-16	54	-83	44
2	66736	284	RSRP_SINR	-100	-111	76	-35	69	-21	44	-83	50
3	66736	184	RSRP_SINR	-114	-111	44	-21	34	-24	13	-83	50

Fig. 4. UE Modem Layer 1 (ML1) Cell Measurements captured via QXDM.

from commercial networks operated by different MNOs. Over a three-month period in 2025, 5G and LTE measurements were collected at various times across 10 urban areas in Boston and San Francisco. The dataset covers two major U.S. MNOs and was curated to capture real-world behavior of the Measurement Reporting mechanisms. Data collection was performed using rooted COTS UEs, including two Google Pixel 5 devices, an LG Velvet 5G, and a OnePlus 8 5G, all equipped with measurement tools such as Network Signal Guru (NSG) [37]. Devices were carried in motion through the areas while connected to commercial networks. Logged data was later analyzed using Qualcomm’s QXDM [38] to extract low-level Modem and Layer 1 metrics. In total, our dataset comprises 22,626 seconds of telemetry and 232,810 NSG-QXDM data points, with each point capturing detailed per-cell measurements, including PCI, RSRP, RSRQ, and SINR of neighboring cells.

In addition to standard Measurement Reports, we also incorporate *Modem Layer 1 (ML1) Cell Measurement Results*, obtained via QXDM, as depicted in Figure 4. Unlike traditional Measurement Reports, ML1 data provide high-frequency sampling of neighboring cell signal stability, offering a much finer temporal resolution. ML1 data are parsed similarly to standard Measurement Reports, as they share the same underlying structure. As a result, the telemetry vectors used by ARGOS maintain a consistent format while being enriched with additional measurements, thereby enhancing the dataset’s granularity. However, the main limitation of ML1 data is that they are never transmitted to the gNB over the air (OTA) and remain accessible only at the UE side. Since operator-configured Measurement Reports are often sparse or event-triggered, ML1 data significantly enhance the observability of UE behavior by providing a more continuous and detailed stream of measurements. We advocate for the integration of ML1 data into O-RAN control loops to support more accurate and timely anomaly detection.

C. Rogue Cell Inclusion

The performance of ARGOS is evaluated using both benign and malicious cellular network traffic. Since the dataset contains only legitimate data, two types of RBS strategies, *Adversary 1 (A1)* and *Adversary 2 (A2)*, are emulated by selecting a valid cell, excluding it during training, and reintroducing it during inference. This methodology allows for realistic evaluation of our system by leveraging authentic cell behavior, eliminating the need for synthetic data generation. A1 does not

replicate the PCI of an existing cell. As a result, the reintroduced cell during inference retains its identity, allowing us to assess the model’s ability to detect previously unrecognized cells. A2 carries out a more intricate attack by replicating the PCI of a legitimate cell, causing the reintroduced cell during inference to share the same PCI as an existing legitimate one. In the case of A2, the replicated PCI reappears across multiple per-second telemetry vectors where it was previously absent, creating unusual cell combinations. The model is evaluated both on abnormal co-occurrence patterns in Measurement Reports and on its ability to detect anomalies based on signal characteristics, as A2 instances, despite sharing the same PCI, will exhibit at least slight differences in power, relatively to learned power-levels in conjunction with neighboring cells.

D. Performance Evaluation

To evaluate our solution, we first compare the performance of the ML models integrated into ARGOS, followed by an assessment of their system-level impact on the RIC platform.

To assess the performance of the ML models, we use four standard classification metrics defined in Table I. In Table I, the value TP stands for True Positives, TN for True Negatives, FP for False Positives and FN for False Negatives. Accuracy provides an overall correctness measure, while Precision emphasizes the proportion of true anomalies among all flagged instances. Recall captures the model’s ability to detect all actual anomalies, and the F1-Score provides the harmonic mean between Recall and Precision, especially valuable for imbalanced datasets. It is important to mention at this stage that, under realistic conditions, if a rogue cell exists within a particular vicinity, regardless of its PCI or signal characteristics, the UE ML1 and Measurement Reports would reflect its presence with high frequency every second, and consequently, so would the per-second telemetry vectors sent to the xApp. As a result, we consider as anomalous those per-second vectors in which the reintroduced rogue cell appears more than a certain number of times, while the remaining vectors are considered legitimate. More specifically, as shown in Figure 5, we evaluate the ML models in terms of anomaly detection for seconds where the reintroduced cell appears at least 2, 3, or 4 times, a threshold we define as Per-Second Rogue Cell Count.

It is evident that across all AE variations, the best performance is achieved when the Per-Second Rogue Cell Count is ≥ 3 . As shown in Figure 5, the VAE attains the highest performance under this condition, reaching 99.5% Accuracy, 97.7% Precision, 99.5% Recall, and a 98.1% F1 Score, with a False Positive Rate (FPR) as low as 0.6%. Both the AE and DAE also demonstrate strong performance, achieving 98.6% and 98.3% Accuracy, respectively, while maintaining FPR values below 1.9%. In contrast, the Isolation Forest underperforms across all evaluation metrics, with a maximum Accuracy of 84.6%, indicating that the randomized attribute selection is suboptimal for capturing the temporal patterns in Measurement Report behavior.

Model	Train(s)	Infer(s)	CPU(%)	Memory(MB)
AE	109.50	0.28	65.20	527.20
DAE	158.68	0.28	67.16	527.24
VAE	81.99	0.27	67.25	528.72
IF	0.44	11.59	99.16	548.23

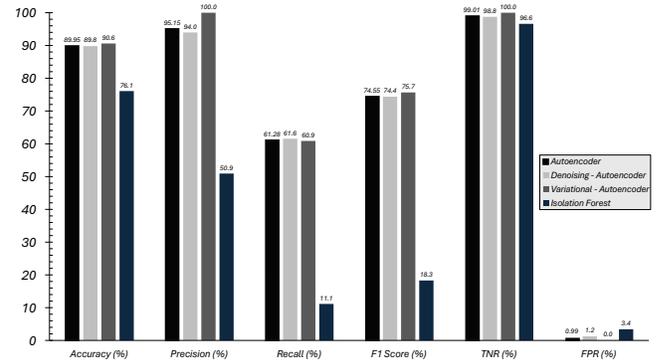
TABLE II
SYSTEM PERFORMANCE OF XAPP ML MODELS IN THE NEAR-RT RIC.

For a Per-Second Rogue Cell Count threshold of ≥ 2 , the models achieve their highest Precision, reaching 100% for the VAE. However, the performance of the remaining metrics declines, indicating that while the models are highly effective at avoiding false alarms, they struggle to correctly classify seconds with sparse rogue cell presence. This results in a drop in Accuracy, as such vectors are often misclassified as legitimate. Conversely, for a threshold of ≥ 4 , Recall reaches its peak, with all illegitimate seconds correctly identified as anomalous, but at the cost of a higher FPR. Based on this analysis, a threshold of ≥ 3 offers the most balanced performance, maintaining both high Accuracy and low FPR. For specific operational goals, the system can be configured with alternative thresholds, depending on the desired trade-off between Precision and Recall.

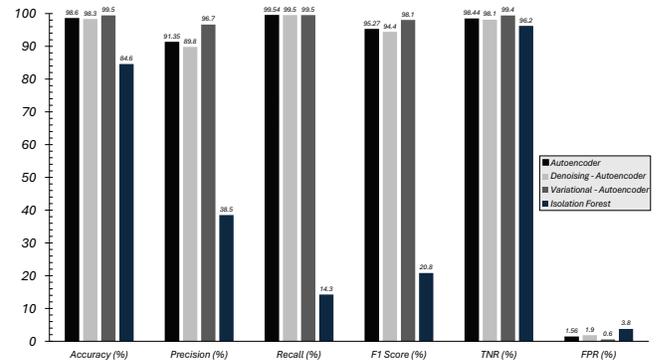
Additionally, we evaluate the impact of *ARGOS* on the control plane, specifically focusing on the NearRT-RIC. To this end, we assess the training and inference times, as well as the CPU and memory overhead, across all implemented ML models, as presented in Table II. The evaluation is performed using input datasets of 2000 seconds for training and 500 seconds for inference, drawn from the same area, and executed using a single CPU core on the host machine. As shown in Table II, the VAE achieves the lowest training (81.99 seconds) and inference (0.27 seconds) times among all AE variants, rendering it suitable for real-time systems. In contrast, the Isolation Forest yields the fastest training time (0.44 seconds) but the highest inference time (11.59 seconds), shifting its computational burden to inference. Regarding CPU utilization, all AEs occupy approximately 65–68% of a single core, whereas the Isolation Forest reaches up to 99.16%, accounting for its prolonged inference duration. Lastly, in terms of memory overhead, the Isolation Forest exhibits the highest usage at 548 MB, only slightly exceeding that of the AEs, which peak at 528 MB.

VI. ETHICAL CONSIDERATIONS

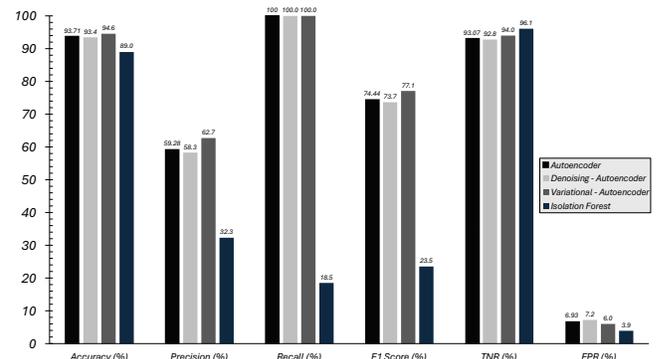
Due to ethical considerations and applicable legal frameworks, it is essential to clarify the methodology used in both our isolated testbed experiments and real-world measurements. All active RBS attack scenarios involving over-the-air transmissions were conducted exclusively within our isolated testbed environment, ensuring no interference with operational commercial networks. Specifically, all RF transmissions from srsRAN base stations and UEs were confined to a shielded anechoic chamber. Furthermore, all data collection procedures



(a) Per-Second Rogue Cell Count ≥ 2 .



(b) Per-Second Rogue Cell Count ≥ 3 .



(c) Per-Second Rogue Cell Count ≥ 4 .

Fig. 5. ML model inference performance across rogue cell appearance thresholds.

during outdoor measurements strictly adhered to ethical guidelines. Passive network monitoring and data collection were carried out solely using our own devices, each equipped with a valid SIM card. The process of connecting COTS UEs to live operator networks and logging RRC and PHY-layer messages reflects standard UE behavior and does not disrupt normal network operations. This study complies fully with the terms of service of the participating wireless carriers and does not raise any ethical concerns.

VII. CONCLUSIONS & FUTURE WORK

In this work, we present *ARGOS*, the first comprehensive O-RAN compliant system for detecting RBS that attempt downgrade attacks, deployed directly within the Near-RT RIC. *ARGOS* integrates an extended KPM Service Model and a custom xApp featuring ML-based anomaly detection. It enables real-time identification of RBS threats within a given area by classifying telemetry data based solely on UE ML1 logs and Measurement Reports. The proposed extension to the 3GPP KPM Service Model allows for richer UE and E2 node-derived telemetry, enhancing detection capabilities. To validate *ARGOS* under both controlled and real-world conditions, we built a dedicated testbed to verify the threat model and supplemented it with real-world measurements across commercial networks. Among the models evaluated, the Variational Autoencoder achieved the best performance, with 99.5% Accuracy and a False Positive Rate of just 0.6%. Additionally, our system demonstrates low CPU and memory overhead, making it practical for deployment in production O-RAN environments. As future work, we aim to extend the system to cover a wider range of attacks and further enhance the KPM Service Model to support richer telemetry. Beyond anomaly detection, we plan to explore ML-based resource optimization within the RAN and enable collaborative decision-making through integration with rApps at the SMO level.

REFERENCES

- [1] <https://www.ookla.com/articles/5g-global-reach-2025>.
- [2] <https://www.ericsson.com/en/reports-and-papers/mobility-report/reports/november-2024>.
- [3] Bedran Karakoc, Nils Fürste, David Rupprecht, and Katharina Kohls. Never let me down again: Bidding-down attacks and mitigations in 5g and 4g. *WiSec '23*, 2023.
- [4] Adrian Dabrowski, Nicola Pianta, Thomas Klepp, Martin Mulazzani, and Edgar Weippl. Imsi-catch me if you can: Imsi-catcher-catchers. *ACSAC '14*, 2014.
- [5] Andy Lilly. Imsi catchers: hacking mobile communications. *Network Security*, 2017.
- [6] Ivan Palamà, Francesco Gringoli, Giuseppe Bianchi, and Nicola Blefari-Melazzi. Imsi catchers in the wild: A real world 4g/5g assessment. *Computer Networks*, 2021.
- [7] Altaf Shaik, Ravishankar Borgaonkar, Shinjo Park, and Jean-Pierre Seifert. New vulnerabilities in 4g and 5g cellular access network protocols: exposing device capabilities. *WiSec '19*, 2019.
- [8] Ankush Singla, Rouzbeh Behnia, Syed Rafiul Hussain, Attila Yavuz, and Elisa Bertino. Look before you leap: Secure connection bootstrapping for 5g networks to defend against fake base-stations. *ASIA CCS '21*, 2021.
- [9] Open5GS. <https://open5gs.org>.
- [10] SRS. Software Radio Systems. Open source SDR 4G/5G software suite. <https://github.com/srsran/srsRAN>, 2020.
- [11] OpenAirInterface Software Alliance. Flexible RAN Intelligent Controller (FlexRIC). <https://github.com/lgs96/flexric>, 2021.
- [12] Kai Tu, Abdullah Al Ishtiaq, Syed Md Mukit Rashid, Yilu Dong, Weixuan Wang, Tianwei Wu, and Syed Rafiul Hussain. Logic gone astray: A security analysis framework for the control plane protocols of 5g basebands. In *33rd USENIX Security Symposium*, 2024.
- [13] Stavros Eleftherakis, Domenico Giustiniano, and Nicolas Kourtellis. Sok: Evaluating 5g protocols against legacy and emerging privacy and security attacks, 2024.
- [14] Xinxin Hu, Caixia Liu, Shuxin Liu, Wei You, Yingle Li, and Yu Zhao. A systematic analysis method for 5g non-access stratum signalling security. *IEEE Access*, 2019.
- [15] Syed Rafiul Hussain, Mitziu Echeverria, Omar Chowdhury, Ninghui Li, and Elisa Bertino. Privacy attacks to the 4g and 5g cellular paging protocols using side channel information. *Network and distributed systems security (NDSS)*, 2019.
- [16] Martin Kotuliak, Simon Erni, Patrick Leu, Marc Röschlin, and Srdjan Capkun. LTrack: Stealthy tracking of mobile phones in LTE. In *31st USENIX Security Symposium*, 2022.
- [17] 3rd Generation Partnership Project (3GPP). Security Architecture and Procedures for 5G System. Technical Specification TS 33.501.
- [18] 3rd Generation Partnership Project (3GPP). Study on 5G Security Enhancement against False Base Stations. Technical Report TR 33.809.
- [19] Stavros Eleftherakis, Timothy Otim, Giuseppe Santaromita, Almudena Díaz Zayas, Domenico Giustiniano, and Nicolas Kourtellis. Demystifying privacy in 5g stand alone networks. *ACM MobiCom '24*, 2024.
- [20] Syed Rafiul Hussain, Mitziu Echeverria, Imtiaz Karim, Omar Chowdhury, and Elisa Bertino. 5greasoner: A property-directed security and privacy analysis framework for 5g cellular network protocol. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019.
- [21] Haohuang Wen, Phillip Porras, Vinod Yegneswaran, Ashish Gehani, and Zhiqiang Lin. 5g-spector: An o-ran compliant layer-3 cellular attack detection service. In *Proceedings of the 31st Annual Network and Distributed System Security Symposium (NDSS'24)*, 2024.
- [22] Theodoros Tsourdinis, Nikos Makris, Thanasis Korakis, and Serge Fdida. Ai-driven network intrusion detection and resource allocation in real-world o-ran 5g networks. *ACM MobiCom '24*, 2024.
- [23] Jun-Hong Huang, Shin-Ming Cheng, Rafael Kaliski, and Cheng-Feng Hung. Developing xapps for rogue base station detection in sdr-enabled o-ran. In *IEEE Conference on Computer Communications Workshops*, 2023.
- [24] Emmanuel N. Amachaghi, Sulyman Abdulkareem, Sotiris Chatzimiltis, Mohammad Shojafar, and Chuan H. Foh. An efficient intrusion detection solution for near-real-time open-ran. In *2024 IEEE Symposium on Computers and Communications (ISCC)*, 2024.
- [25] Jung-Erh Chang, Yi-Chen Chiu, Yi-Wei Ma, Zhi-Xiang Li, and Cheng-Long Shao. Packet continuity ddos attack detection for open fronthaul in oran system. In *2024 IEEE Network Operations and Management Symposium*, 2024.
- [26] Alessio Scalingi, Salvatore D'Oro, Francesco Restuccia, Tommaso Melodia, and Domenico Giustiniano. Det-ran: Data-driven cross-layer real-time attack detection in 5g open rans. In *IEEE INFOCOM 2024*, 2024.
- [27] Pawel Kryszkiewicz and Marcin Hoffmann. Open ran for detection of a jamming attack in a 5g network. In *2023 IEEE 97th Vehicular Technology Conference*, 2023.
- [28] Bruno Missi Xavier, Merim Dzaferagic, Irene Vilà, Magnus Martinello, and Marco Ruffini. Cross-domain ai for early attack detection and defense against malicious flows in o-ran. In *ICC 2024 - IEEE International Conference on Communications*, 2024.
- [29] Haohuang Wen, Prakhar Sharma, Vinod Yegneswaran, Phillip Porras, Ashish Gehani, and Zhiqiang Lin. 6g-xsec: Explainable edge security for emerging openran architectures. *HotNets '24*, 2024.
- [30] 3rd Generation Partnership Project (3GPP). NR; Radio Resource Control (RRC); Protocol Specification. Technical Specification TS 38.331.
- [31] 3rd Generation Partnership Project (3GPP). NG-RAN; Architecture Description. Technical Specification TS 38.401.
- [32] O-RAN Work Group 3 (WG3). E2 Application Protocol (E2AP). Technical Specification O-RAN.WG3.TS.E2AP-R004-v07.00.
- [33] O-RAN Work Group 3 (WG3). Near-Real-Time RAN Intelligent Controller (RIC) Architecture. Technical Specification O-RAN.WG3.TS.RICARCH-R004-v07.00.
- [34] 3rd Generation Partnership Project (3GPP). Management and Orchestration; 5G Performance Measurements. Technical Specification TS 28.552.
- [35] Dor Bank, Noam Koenigstein, and Raja Giryes. *Autoencoders*. Springer International Publishing, 2023.
- [36] SRS. <https://docs.srsran.com/projects/project/en/latest/tutorials/source/handover/source/index.html>, 2020.
- [37] Qtrun Technologies. Network signal guru. <https://www.qtrun.com/eng/nsg/>, 2025.
- [38] Qualcomm. QxDM Professional Qualcomm eXtensible Diagnostic Monitor., 2022.