A Fine-Grained Telemetry Stream for Security Services in 5G Open Radio Access Networks

Haohuang Wen
wen.423@osu.edu
The Ohio State University

Phillip Porras
porras@csl.sri.com
SRI International

Vinod Yegneswaran
vinod@csl.sri.com
SRI International

Zhiqiang Lin
zlin@cse.ohio-state.edu
The Ohio State University

ABSTRACT

The Open Radio Access Network (O-RAN) is an emerging paradigm for developing the next-generation radio access network (RAN) for 5G and beyond. Inspired by the principles from software-defined networks (SDNs), the key innovation of O-RAN is the disaggregation of control logic from the network data plane, by using a centralized RAN intelligent controller (RIC), with customized xApps and service models. The O-RAN's novel design transforms the traditional monolithic network infrastructure into an open, programmable, and interoperable RAN. These distinctive features make O-RAN ideal for deploying extensible security services against a wide range of prevalent threat vectors (e.g., malicious transmitters that spoof, interfere, or flood communications between mobile devices and the 5G RAN), which can compromise the security, privacy, and availability of mobile devices and the network itself, at very low cost. Unfortunately, we find that the existing exemplar xApp models of O-RAN and the underlying telemetry streams that drive these applications are insufficient for developing robust security countermeasures. In this paper, we propose MobiFlow, a fine-grained telemetry stream tailored for security analysis on O-RAN. We envision MobiFlow as an enabling building block upon which novel 5G services can be implemented, offering device and RAN-specific run-time security monitoring, intelligent RAN control, and security-focused AI/ML assisted applications.

CCS CONCEPTS

• **Security and privacy** → **Mobile and wireless security**; • **Networks** → **Mobile and wireless security**; **Programmable networks**.

KEYWORDS

5G, Open Radio Access Network, Cellular Network Security

ACM Reference Format:
Haohuang Wen, Phillip Porras, Vinod Yegneswaran, and Zhiqiang Lin. 2022. A Fine-Grained Telemetry Stream for Security Services in 5G Open Radio Access Networks. In *Emerging Topics in Wireless (EmergingWireless '22), December 9, 2022, Roma, Italy.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3565474.3569070

1 INTRODUCTION

The fifth-generation (5G) cellular network has significantly changed the way we connect to the world and provides many new opportunities. In addition to the generational bandwidth upgrade of mobile broadband, 5G also enables many novel applications, including but not limited to massive Internet-of-Things (IoT), smart transportation (e.g., autonomous and connected vehicles), and immersive augmented and virtual reality (AR/VR). A key revolution in 5G is the Open Radio Access Network (O-RAN) [2] proposed by the O-RAN Alliance, which emerges as a new paradigm for developing the next-generation cellular network. Such a design has been instantiated in many O-RAN-compliant implementations, and one notable example is the Software-Defined Radio Access Network (SD-RAN [5]) developed by the Open Networking Foundation (ONF), which is an open-source and cloud-native platform.

The O-RAN design is inspired by the software-defined network (SDN) principles to address the limitations of contemporary Radio Access Network (RAN) which is monolithic and closed-source. More specifically, the O-RAN achieves this breakthrough by separating the network control logic from the network implementation. To this end, O-RAN breaks down the all-in-one RAN infrastructure into multiple logical components, including the Central Unit (CU), Distributed Unit (DU), and Radio Unit (RU), which are inter-connected via standard interfaces. These O-RAN nodes are further connected and controlled by an Intelligent RAN Controller (RIC), which serves as a standalone programmable component to monitor and manage the network. Such a design enables network-overlay services to be developed as "plug-and-play" xApps and deployed on the RIC, regardless of the vendor-specific implementation of the RAN.

O-RAN brings unprecedented opportunities for not only network operators, but also other stakeholders (e.g., vendors and system integrators) and researchers to build extensible security services atop the programmable RIC to address various cellular threats. However, there is currently a dearth of security-focused xApps and underlying service models targeting the O-RAN. Although the O-RAN has provided many exemplar xApps, such as key performance indicator (KPI) monitoring, traffic steering, and anomaly detection [3], these

EmergingWireless '22, December 9, 2022, Roma, Italy

ACM ISBN 978-1-4503-9934-0/22/12…$15.00
https://doi.org/10.1145/3565474.3569070

xApp development models cannot support robust security defenses against a wide range of practical attacks. Specifically, we find that the telemetry collected and used by these xApps are too coarse-grained to support sophisticated security analysis.

Notably, attacks targeting cellular networks have risen in complexity and stealthiness over time, due to the availability of inexpensive commercial-off-the-shelf (COTS) software-defined radios (SDRs) and open-source cellular software stacks, that simplify execution of cellular-infrastructure attacks. For instance, it costs one as little as $500 to build a fully controllable attack device (e.g., a rogue UE or base station) by using a Lime SDR mini and a Raspberry Pi 4 [1]. It has been demonstrated that the RAN and its connected user equipment (UE) are vulnerable to numerous attacks from various classes of (active) adversaries, including malicious UEs [9, 17, 18], fake base stations [15, 17, 20], Man-in-the-Middle (MiTM) attackers [28, 29], and surgical signal injectors [19, 23, 32]. Most attacks are actually transparent to users and occur very early before the UE establishes a reliable data connection with the RAN (e.g., by injecting a sequence of malicious control signals), making them very difficult to detect in practice. These exploits on the RAN and UEs can trigger severe consequences, such as service outages, resource depletion, and privacy leakage.

To facilitate the design of extensible and robust security services on the O-RAN, we propose MobiFlow, a cellular network telemetry stream tailored for security analysis. It is inspired by *NetFlow*, which was initially introduced on Cisco routers to collect TCP/IP network traffic, and has been widely adopted for network monitoring [21]. MobiFlow converts the cellular traffic from various network entities (i.e., UEs, RAN, and the core network) into flow records, and exported to the upper control layer through the standard E2 interface. The exported MobiFlow records contain fine-grained statistics of the network entities in real-time, to further enable robust, intelligent, and real-time security analysis performed by the xApps on the near-real-time-RIC. We envision MobiFlow to serve as a foundational building block to drive many practical security applications, including but not limited to real-time monitoring, intrusion detection, intelligent control, and artificial intelligence (AI) or machine learning (ML) assisted security services.

The rest of the paper is organized as follows. We introduce the O-RAN background in §2, describe our motivation in §3, and illustrate MobiFlow's design in §4. Finally, we conclude in §5 and outline the potential research directions.

2 BACKGROUND

This section provides the pertinent background regarding O-RAN's unique features including its disaggregated architecture and its control-layer design, as depicted in Figure 1 [3].

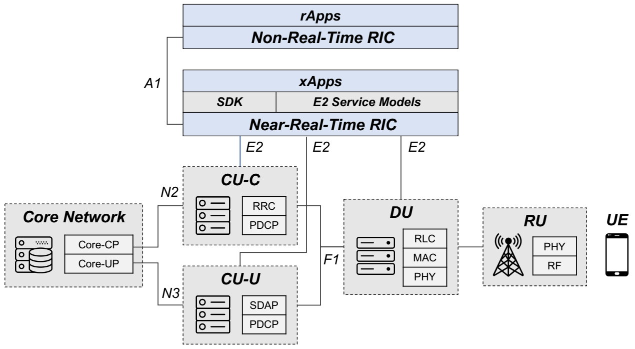

Figure 1: The O-RAN architecture.

O-RAN's Disaggregated Architecture. The O-RAN design follows the functional split defined in the 3GPP specification [6] and divides the 5G RAN infrastructure into logical nodes, each of which hosts specific functions of the protocol stack. Starting from the lower layers of the protocol stack, the Radio Unit (RU) is the typical radio hardware deployed at the front-haul network to process layer-1 physical radio signals from surrounding user equipment. Next, the Distributed Units (DU) and Central Units (CU) are logical components that can be hosted at edge clouds, and they handle the layer-2 and layer-3 functions of the network. To be more specific, the DU handles layer-2 functions such as Medium Access Control (MAC) and Radio Link Control (RLC). The CU mainly hosts layer-3 control functions such as Radio Resource Control (RRC), and partial layer-2 functions. A CU node is further split into the Control-Plane (CU-C) and User-Plane (CU-U) nodes for control signals and user data, respectively. For upper-layer protocols, such as the Non-Access Stratum (NAS) that communicates the UEs with the core network, and the Internet Protocol (IP), their traffic is also relayed by the CU. O-RAN's components are inter-connected through standard interfaces, such as the F1 interface that connects the DU and CU nodes, and the N2/N3 interfaces that connect the CU-C and CU-U to the core network [7]. The uplink traffic of subscribed UEs consecutively flows through the DU and CU at each stack layer, and is finally delivered to the core (vice versa for the downlink traffic).

O-RAN's Control-layer Design. Another key innovation in 5G O-RAN is the RAN Intelligent Controller (RIC), which decouples the network control plane (i.e., the layer that handles network control logic) from the data plane (i.e., the actual network implementation). The RIC is independent from the RAN nodes (i.e., CUs and DUs) and connects with them through the standard E2 interface [3]. Therefore, such a design provides the RAN with interoperability among various vendors, and offers valuable programmability to allow "plug-and-play" software components to be deployed at the control layer. To support such programmability, the O-RAN introduces the xApps, which are microservices deployed on the near-real-time RIC (nRT-RIC) for various customized control and monitoring functions over the RAN components. The O-RAN Alliance has demonstrated this capability with a few exemplar xApps such as Key Performance Measurement (KPM), RAN slicing management, and traffic steering [3]. The xApp development is facilitated by the software-development kit (SDK) and the E2 service model (E2SM) which acts as a contract to define the communication between the specific xApp and its subscribed RAN nodes. To meet the real-time requirement, the nRT-RIC runs a control loop at a timescale of 10-1000ms. For tasks with non-real-time requirement (e.g., ML model training), they are hosted as rApps executed on the non-real-time RIC with timescale larger than 1s.

3 MOTIVATION

3.1 Existing Cellular Attacks and Defenses

To motivate why there is a need for security in the cellular network, we describe existing cellular attacks and their impact on security and privacy. Based on the adversary types, we summarize the following six categories of adversary models below in details [8].

- **Adversarial UE**. An attacker can easily set up a malicious UE by using a valid subscriber network identity (e.g., SIM) and a cellular device (e.g., a smartphone). It is preferable for the attacker to build such a device with an open-source cellular stack [14] on an SDR to achieve fine-grained control over the protocol messages. Using such an adversarial UE, an attacker can perform DoS attacks targeting either the RAN (e.g., radio jamming) or another UE in the network (e.g., TMSI replay [18]).
- **Fake Base Station**. Due to the lack of verification of the authenticity of base stations, fake base station (FBS) has been an infamous issue, which is usually exploited to steal end-users' International Mobile Subscriber Identity (IMSI) in prior GSM (2G) network [11]. As in the new 5G release, this problem is mitigated by introducing IMSI encryption. However, FBS can also be set up for many other control message attacks, and may cause denial-of-service and privacy leakage in UEs [17].
- **MiTM Attacker**. A MiTM attack relay essentially impersonates a legitimate BS towards a victim UE and a legitimate UE towards a victim BS. As a result, two SDRs are required to launch the attack. An MiTM attacker can then replay, eavesdrop, or inject messages in the traffic. For instance, the aLTEr attack [28] exploits the lack of integrity protection for user-plane messages to redirect DNS requests of the victim UE. Fundamentally, MiTM attacks share the same root cause with FBSes.
- **Signal Injector**. Most recently, there is a new type of attack that allows adversaries to use an SDR to inject malicious signals into cellular traffic while maintaining high stealthiness (i.e., with slightly higher signal strength) [32]. This signal injection (or overshadowing) attacks can be further exploited to leak user's identity (e.g., IMSI in LTE network [19]) or DoS a specific UE [13].
- **Compromised Core Network**. Although not discussed extensively in the literature, a compromised core network node can bring security and privacy risks, such as exposing sensitive UE information or cryptographic keys. As in 5G, the core network may impose actual security threats due to 5G-specific design choices such as network slicing and Network Function Virtualization (NFV).
- **Passive Eavesdropper**. It has been shown that passive sniffer can eavesdrop on the uplink or downlink traffic for privacy attacks such as IMSI stealing and user tracking in LTE networks [16, 19]. This privacy information is inferred from the paging messages [30], the synchronization procedure [19] or side-channel information [16].

In response to these threats, existing defenses are deployed at (1) the protocol layer, (2) UEs, and (3) the RAN. Protocol-level defenses provide fundamental countermeasures but require modification to the cellular standard. For instance, digital-signature-based or cryptographic-based solutions can be applied to verify BS identities [8]. Unfortunately, these countermeasures not only require extensive changes of the specification and the infrastructure, but also introduce additional cost and performance overhead, making them hard to push out in the near future. UE-based defenses detect threats on local devices and provide the corresponding remedies [10, 12], but they have a limited view on only a single device. Network-based defenses have a global vision on subscribed UEs and networks (i.e., can detect attacks not only targeting UEs but also the RAN), and existing defenses mainly focus on FBS detection [22, 24, 25, 33]. While this type of defense is promising, the existing works are not based on the O-RAN.

3.2 Why O-RAN?

We believe that the O-RAN is ideal for deploying security services against various cellular threats, for the following three reasons. (1) *Programmability*: the disaggregation of the RAN control plane and data plane offers programmability that allows software-defined security solutions to easily integrate as extensible xApps on the nRT-RIC without modification of the RAN infrastructure. (2) *Openness*: the O-RAN provides open standards and interfaces, as well as open-source reference implementations, which allow security stakeholders to play a role in the development and vetting process. Contributions from the open-source community can also help improve the security of the O-RAN. (3) *Intelligence*: the O-RAN design enables data-driven solutions to be deployed. We anticipate that a more fine-grained security telemetry stream and dataset can drive AI/ML-based approaches to enhance the security of the O-RAN, such as intelligent RAN monitoring and control.

However, we still need to address some current limitations in the O-RAN in order to achieve the security goal. Specifically, we find that the existing exemplar xApp models [3] fail to support sophisticated security services, due to the following two main reasons. First, these xApps are not designed to achieve certain security goals. In particular, the anomaly detection xApp detects anomalous UEs in the network and does not aim to detect intrusion attacks. Second, these xApps rely on coarse-grained telemetry from the network which is not sufficient to support fine-grained security analysis. For instance, the KPI monitor xApp collects only a few high-level or aggregated statistics from the RAN nodes such as UE identifiers, locations, and cell status. However, as shown in previous work [12, 22], a robust security service requires more fine-grained telemetry, such as packet-level information and physical-layer metrics. As a result, it is clear that a network telemetry stream tailored for security is needed to facilitate the security services.

4 MOBIFLOW DESIGN

To fill the research gap, we propose MobiFlow, a fine-grained network telemetry stream tailored for security analysis, which provides fundamental building blocks to drive various security services on the O-RAN. MobiFlow's design is inspired from the *NetFlow*, a network stream that aggregates TCP/IP packets into flow records for upstream analysis tasks [21]. Similarly in MobiFlow, it gathers telemetry streams from different cellular network entities (i.e., UEs, RAN, and core networks) to support the security services of xApps on the nRT-RIC. The envisioned architecture of MobiFlow

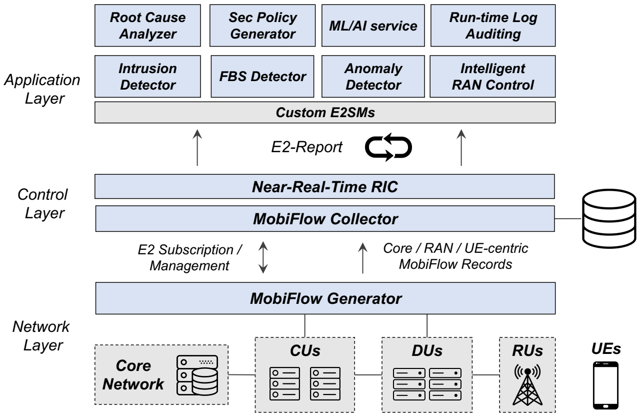

Figure 2: Envisioned architecture of MobiFlow and its driven applications atop O-RAN.

deployed atop O-RAN is illustrated in Figure 2, and is abstracted as three different layers: network, control, and application, which are further explained below.

4.1 Network Layer

The network layer aims to collect network telemetry from the cellular network components and export MobiFlow records to the control layer. Based on the architecture of the O-RAN described in §2, we design a MobiFlow generator to export MobiFlow records from the CU and DU nodes, which have direct connections to the RIC via the standard E2 interfaces. To comprehensively support the security analysis of adversarial threats mentioned in §3.1, we present three types of MobiFlow records described below.

4.1.1 UE-centric MobiFlow. This type of MobiFlow conveys the fine-grained statistics of a specific UE subscribed to the network. It can express the UE's meta information (e.g., identifiers and its subscribed network), states (e.g., RRC connection states and security conditions), control-plane traffic (e.g., RRC message sequences), user-plane traffic, security algorithms, error codes, timers, and other aggregated statistics (e.g., counters for connection failure and establishment).

Example Security Applications. The UE-centric MobiFlow enables the security xApps on the nRT-RIC to assess the security of UEs, and could also be useful in leveraging the UE's sensing capability to detect surrounding threats. In the following, we describe a number of concrete security applications driven by UE-centric MobiFlow records.

- **Malicious UE Detection**. As we have presented in §3.1, adversarial UEs can be easily established using COTS radio hardware and open-source software, which can further perform availability attacks on the RAN infrastructure and other UEs. Therefore, the fine-grained telemetry of specific UEs can help O-RAN operators identify potentially malicious UEs subscribed to the network, based on metrics such as abnormal traffic at the cellular control-plane or user-plane (e.g., numerous RRC connection request messages sent to DoS a base station [17]).
- **UE-targeted Attack Detection**. As a UE is vulnerable to many external vectors, the UE-centric MobiFlow could help detect malicious attacks towards a specific UE, perform diagnosis for abnormal scenarios, and identify the root causes. For instance, it is shown that UEs are subject to spoofing attacks which cause service outages [18]. As such, the MobiFlow records among various UEs in the network can be cross-checked. In addition, other problematic scenarios such as malicious paging [16] and roaming [9] could also be identified.
- **Rogue Base Station Detection**. It has been shown that crowdsourced measurement reports from UEs are useful in detecting nearby rogue base stations [22, 24]. Since these measurement reports include various measurements of BSes such as location, physical signal strengths, and broadcast information, they can serve as the inputs for heuristics-based detection algorithms, which identify abnormal statistics (e.g., unusual signal strengths and incorrect locations).

4.1.2 RAN-centric MobiFlow. The RAN-specific telemetry is converted into RAN-centric MobiFlow records. The data of interest could be the RAN-specific identifiers, aggregated statistics of subscribed UEs (e.g., counters tracking currently connected UEs and failure UEs in history), measurement statistics (e.g., signal strengths and broadcast parameters), RAN-based traffic (e.g., E2 interface packets), timers, and RAN node states.

Example Security Applications. We anticipate the RAN-centric MobiFlow could drive various RAN-specific security tasks, and we list two concrete security applications below.

- **Malicious RAN Node Detection**. The cloud-native deployment of O-RAN could bring security risks in RAN nodes for being compromised. The abnormal traffic of RAN nodes could help the nRT-RIC infer potential compromised RAN nodes that transmit malicious traffic to other RAN nodes or the RIC.
- **RAN-targeted Attack Detection**. The aggregated measurement statistics could help identify malicious attacks targeting the RAN. For instance, malicious UEs could launch BTS depletion attacks [18] to exhaust the resources of a single BS, further denying the service of all subscribed users. For other similar attacks on the cellular data-plane (e.g., malicious TCP/IP traffic) can be detected in a similar fashion.

4.1.3 Core-centric MobiFlow. Core-specific MobiFlow stream consists of aggregated RAN statistics (e.g., active and inactive RAN nodes), core node status, core network traffic, etc. As the core network does not have a direct connection to the RIC, the core-centric MobiFlow records need to be transmitted via the N2 and N3 interfaces (shown in Figure 1) and reported from the CU nodes.

Example Security Applications. The core-centric MobiFlow records are expected to drive various core-based security applications such as the detection of anomalous core network nodes and network diagnosis. For instance, core-related traffic can help diagnose anomalies such as authentication failure and security mode failure of a specific UE. Core-targeted attacks could also be detected (e.g., network DDoS and Diameter attacks [31]), and compromised core network nodes could be identified. With 5G introducing many new network features, such as network slicing [26], there could be potentially larger attack surfaces and new types of attacks (e.g., ML-based attacks) on the core network side.

4.2 Control Layer

The control layer resides at the nRT-RIC. In addition to MobiFlow record collection and storage, it manages the subscribed RAN nodes via the E2 interface, and distributes the MobiFlow records to the xApps. The control layer essentially serves as a proxy to manage subscriptions and transmission between the xApps and the RAN nodes. We describe the control layer's two main functions below.

Subscription Management. In the O-RAN, both CU and DU nodes are abstracted as E2 nodes and managed through a publish-subscribe model via the standard E2 interface [3]. Upon the start of an E2 node, it must first complete an E2 setup procedure with the nRT-RIC, which declares the E2 node's service and capability as *RAN Functions*. At the application layer, the xApp examines the setup E2 nodes for their supported report metrics and report styles, and subscribes to E2 nodes of interest. The subscription messages indicate the report metrics of the xApp's interest (can be a subset of all supported metrics of the E2 nodes). Afterwards, the communication between the E2 nodes and the xApps is activated and performed through the nRT-RIC.

MobiFlow Collection and Reporting. The collection and reporting of MobiFlow records needs to address two problems: (1) what telemetry to be collected in a MobiFlow record? and (2) how often to report a MobiFlow record to the xApps? For (1), the telemetry required by a xApp is declared in its E2 service model (E2SM), as mentioned in §2. An E2SM is defined based on four basic E2 services: *Report*, *Insert*, *Control*, and *Policy*. For instance, an intrusion detection xApp may periodically receive UE-centric MobiFlow records reported via the *E2-Report* operations, to subscribe to the real-time UE status. For (2), the current E2 interface supports either *periodic* or *trigger-based* report declared in the E2 setup messages. For example, a xApp may ask an E2 node to periodically report a RAN-centric MobiFlow record, and receive a UE-centric MobiFlow record upon state change of a UE (i.e., trigger-based events). The property of the nRT-RIC ensures that each operation cycle is performed at a near-real-time scale (i.e., 10-1000ms [3]).

4.3 Application Layer

As mentioned in §2, the upstream control layer logic is implemented at the xApps on the nRT-RIC, which are hosted at the application layer. Each xApp interacts with the nRT-RIC through the four basic E2 operations, such as receiving MobiFlow records via the *E2-Report* packets with flexible report frequencies. While we have already described a number of concrete security applications in §4.1, they can be integrated into various security services. In the following, we further summarize a set of envisioned security xApps enabled by MobiFlow on the application layer.

- **Real-time Monitoring**. MobiFlow telemetry and the property of the RIC enables real-time monitoring of the cellular network and UEs. Such monitoring can aim to detect network or UE anomalies, based on either aggregated statistics (e.g., failure counters and timers) or fine-grained states and packet-level information. One notable example is rogue base station detection [22, 24] as we have mentioned in §4.1.
- **Intrusion Detection**. In response to the various adversarial attacks in §3.1, an intrusion detection system can be deployed to detect incoming attacks and generate real-time alerts. The detection can be formulated using either heuristic-based rule sets or learning-based approaches (e.g., automata and machine learning), by leveraging the MobiFlow telemetry from the E2 nodes (e.g., control message sequences of UEs [12]). We have mentioned a few example applications such as detecting malicious UE attacks in §4.1.
- **Intelligent RAN Control**. To address the detected intrusions, an intelligent RAN controller can be deployed with security policies and rule sets. In addition to the *E2-Report* mechanism, the E2 interface supports *E2-Control* that allows xApps to proactively initiate control requests to the E2 nodes through the RIC. To this end, dynamic RAN control adjustment can be achieved, such as fine-tuning of control parameters (e.g., UE timeout thresholds), switching certain control functions on/off, and dropping a specific packet or connection.
- **AI/ML-assisted Security**. The O-RAN allows ML models to be deployed as either standalone containers or downloaded files. To this end, pre-trained ML models can be loaded to facilitate various security tasks (e.g., malicious UE prediction, attack detection, and anomalous traffic classification) [27]. MobiFlow records can be fed into various AI/ML models and algorithms (e.g., decision trees, support vector machines, and deep neural networks) for training and prediction.
- **Root-cause Analysis**. MobiFlow records can help infer the root causes of anomalous or failure events within the network, such as tracking data provenance, developing causal models, and visualizing data flow graphs. This could help network analysts quickly identify the core of the problems and provide resolutions.
- **Automated Security Policy Generation**. The management of the xApps and the network involves sophisticated security policies. Although manual configurations can be costly and error-prone, such a process can be automated to reduce security risks and provide more fine-grained control adapted to various scenarios and requirements.

5 CONCLUSION AND FUTURE WORK

We proposed MobiFlow, a NetFlow-based telemetry stream tailored for developing extensible security services on top of the emerging 5G O-RAN. We have illustrated an envisioned architecture of MobiFlow on the O-RAN, and show how extensible modules can be programmed as xApps on the nRT-RIC to drive various security applications, such as network monitoring, intrusion detection, and AI/ML-driven security services. MobiFlow can be instantiated in real networks or existing O-RAN compliant testbeds such as SD-RAN [5] and experimental RAN implementations such as OpenAirInterface [4]. In the following, we outline some potential research directions and challenges for future work.

The first challenge is to balance security and privacy. While MobiFlow provides a fine-grained data telemetry stream for security, it also imposes privacy risks to end-users. Data breaches could occur if the data transmission channels are not properly secured, or the RIC, CUs, and DUs deployed on regional clouds are compromised. Therefore, MobiFlow's practical design should consider how fine-grained data can be collected at a minimum while providing security guarantees and respecting user privacy. On the other hand, all O-RAN nodes and communication channels should be hardened to prevent data leakage (e.g., using end-to-end encryption).

Second, the O-RAN architecture brings many new threat vectors compared to the conventional RAN infrastructure, and thus a zero-trust security model should be adopted when designing O-RAN components. Specifically, a zero-trust model should assume that all communication sources are not trusted, and the authenticity of the senders should always be verified. For instance, new attack surfaces may emerge as compromised DUs and CUs, malicious E2 nodes, RIC servers, and xApps on the RIC, which may inject stealthily malicious packets for data poisoning and control interference.

Third, since the O-RAN embraces the open-source community, potential threats also exist among the publicly available open-source xApps and ML models imported to O-RAN. For instance, one notable threat vector may come from attackers distributing or contributing to open-source O-RAN software. Therefore, security vetting techniques should be applied to thoroughly audit these open-source programs before they are used. The isolation and access control among xApps should be properly enforced, and communication should be encrypted and integrity-protected to ensure untrusted xApps cannot break the security and privacy of existing xApps.

Finally, at the design and implementation level, practical design constraints will bring additional challenges for security. As MobiFlow requires fine-grained telemetry for security analysis, its high data granularity could likely impact network performance overhead metrics (e.g., throughput and latency). The O-RAN also requires the nRT-RIC to execute near-real-time tasks with relatively low latency. Therefore, efficient data structures, encoding schemes, and algorithms are important as security building blocks.

ACKNOWLEDGMENT

We thank the anonymous reviewers for their constructive feedback. This research was supported in part by NSF awards 2112471 and 2226443. Any opinions, findings, conclusions, or recommendations expressed are those of the authors and not necessarily of the NSF.

REFERENCES

[1] Hardware options - srsran. https://docs.srsran.com/en/latest/app_notes/source/hw_packs/source/index.html.

[2] O-ran alliance. https://www.o-ran.org/.

[3] O-ran specification. https://www.o-ran.org/specifications.

[4] oai / openairinterface5g. https://gitlab.eurecom.fr/oai/openairinterface5g.

[5] Sdran. https://www.sdran.org/.

[6] 3GPP. Ng-ran architecture description. http://www.3gpp.org/DynaReport/38401.htm.

[7] 3GPP. Ng-ran f1 application protocol (f1ap). http://www.3gpp.org/DynaReport/38473.htm.

[8] 3GPP. Study on 5g security enhancements against false base stations (fbs). http://www.3gpp.org/DynaReport/33809.htm.

[9] Evangelos Bitsikas and Christina Pöpper. Don't hand it over: Vulnerabilities in the handover procedure of cellular telecommunications. In *Annual Computer Security Applications Conference*, pages 900–915, 2021.

[10] Ravishankar Borgaonkar, Andrew Martin, Shinjo Park, Altaf Shaik, and Jean-Pierre Seifert. White-stingray: evaluating imsi catchers detection applications. USENIX, 2017.

[11] Adrian Dabrowski, Nicola Pianta, Thomas Klepp, Martin Mulazzani, and Edgar Weippl. Imsi-catch me if you can: Imsi-catcher-catchers. In *Proceedings of the 30th annual computer security applications Conference*, pages 246–255, 2014.

[12] Mitziu Echeverria, Zeeshan Ahmed, Bincheng Wang, M Fareed Arif, Syed Rafiul Hussain, and Omar Chowdhury. Phoenix: Device-centric cellular network protocol monitoring using runtime verification.

[13] Simon Erni, Martin Kotuliak, Patrick Leu, Marc Röschlin, and Srdjan Capkun. Adaptover: Adaptive overshadowing attacks in cellular networks. *arXiv*, pages 2106–05039, 2021.

[14] Ismael Gomez-Miguelez, Andres Garcia-Saavedra, Paul D Sutton, Pablo Serrano, Cristina Cano, and Doug J Leith. srslte: An open-source platform for lte evolution and experimentation. In *Proceedings of the Tenth ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation, and Characterization*, pages 25–32, 2016.

[15] Syed Hussain, Omar Chowdhury, Shagufta Mehnaz, and Elisa Bertino. Lteinspector: A systematic approach for adversarial testing of 4g lte. In *Network and Distributed Systems Security (NDSS) Symposium 2018*, 2018.

[16] 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) Symposium2019*, 2019.

[17] 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*, pages 669–684, 2019.

[18] Hongil Kim, Jiho Lee, Eunkyu Lee, and Yongdae Kim. Touching the untouchables: Dynamic security analysis of the lte control plane. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 1153–1168. IEEE, 2019.

[19] Martin Kotuliak, Simon Erni, Patrick Leu, Marc Roeschlin, and Srdjan Čapkun. {LTrack}: Stealthy tracking of mobile phones in {LTE}. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 1291–1306, 2022.

[20] Gyuhong Lee, Jihoon Lee, Jinsung Lee, Youngbin Im, Max Hollingsworth, Eric Wustrow, Dirk Grunwald, and Sangtae Ha. This is your president speaking: Spoofing alerts in 4g lte networks. In *Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services*, pages 404–416, 2019.

[21] Yuliang Li, Rui Miao, Changhoon Kim, and Minlan Yu. Flowradar: A better netflow for data centers. In *13th USENIX symposium on networked systems design and implementation (NSDI 16)*, pages 311–324, 2016.

[22] Zhenhua Li, Weiwei Wang, Christo Wilson, Jian Chen, Chen Qian, Taeho Jung, Lan Zhang, Kebin Liu, Xiangyang Li, and Yunhao Liu. Fbs-radar: Uncovering fake base stations at scale in the wild. In *NDSS*, 2017.

[23] Norbert Ludant and Guevara Noubir. Sigunder: a stealthy 5g low power attack and defenses. In *Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, pages 250–260, 2021.

[24] Prajwol Kumar Nakarmi, Mehmet Akif Ersoy, Elif Ustundag Soykan, and Karl Norrman. Murat: Multi-rat false base station detector. *arXiv preprint arXiv:2102.08780*, 2021.

[25] Peter Ney, Ian Smith, Gabriel Cadamuro, and Tadayoshi Kohno. Seaglass: Enabling city-wide imsi-catcher detection. *Proc. Priv. Enhancing Technol.*, 2017(3):39, 2017.

[26] Ruxandra F Olimid and Gianfranco Nencioni. 5g network slicing: A security overview. *IEEE Access*, 8:99999–100009, 2020.

[27] Michele Polese, Leonardo Bonati, Salvatore D'Oro, Stefano Basagni, and Tommaso Melodia. Understanding o-ran: Architecture, interfaces, algorithms, security, and research challenges. *arXiv preprint arXiv:2202.01032*, 2022.

[28] David Rupprecht, Katharina Kohls, Thorsten Holz, and Christina Pöpper. Breaking lte on layer two. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 1121–1136. IEEE, 2019.

[29] David Rupprecht, Katharina Kohls, Thorsten Holz, and Christina Pöpper. Imp4gt: Impersonation attacks in 4g networks. In *NDSS*, 2020.

[30] Altaf Shaik, Ravishankar Borgaonkar, N Asokan, Valtteri Niemi, and Jean-Pierre Seifert. Practical attacks against privacy and availability in 4g/lte mobile communication systems. 2016.

[31] Zhaowei Tan, Boyan Ding, Zhehui Zhang, Qianru Li, Yunqi Guo, and Songwu Lu. Device-centric detection and mitigation of diameter signaling attacks against mobile core. In *2021 IEEE Conference on Communications and Network Security (CNS)*, pages 29–37. IEEE, 2021.

[32] Hojoon Yang, Sangwook Bae, Mincheol Son, Hongil Kim, Song Min Kim, and Yongdae Kim. Hiding in plain signal: Physical signal overshadowing attack on {LTE}. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 55–72, 2019.

[33] Zhou Zhuang, Xiaoyu Ji, Taimin Zhang, Juchuan Zhang, Wenyuan Xu, Zhenhua Li, and Yunhao Liu. Fbsleuth: Fake base station forensics via radio frequency fingerprinting. In *Proceedings of the 2018 on Asia Conference on Computer and Communications Security*, pages 261–272, 2018.