

SG-DEFENSE: A SPACE-GROUND DISTRIBUTED DEFENSE FRAMEWORK FOR DDoS MITIGATION IN O-RAN NTN

Cheng-Feng Hung , Po-Hung Chen , and Shin-Ming Cheng 

ABSTRACT

Non-Terrestrial Networks (NTNs) have been standardized by 3GPP to extend 5G services beyond terrestrial coverage. However, their satellite-based radio access components remain vulnerable to distributed denial-of-service (DDoS) attacks due to long propagation delays and limited on-board resources. Open Radio Access Network (O-RAN) introduces open interfaces and programmable control, creating new opportunities for standardized security enforcement in NTNs. This article presents SG-Defense, a space-ground distributed DDoS mitigation architecture designed for O-RAN-enabled NTNs. SG-Defense adopts a regenerative payload model in which both the centralized unit (CU) and distributed unit (DU) reside on the satellite, enabling low-latency detection and mitigation close to user equipment (UE). A CU-hosted distributed application (dApp) performs flow-level traffic analysis. It correlates multiple standardized UE identifiers to support precise blocking at the DU. At the same time, a ground-based extended application (xApp) on the Near-RT RIC continuously refines the detection model through an application-specific E2 service model that carries flow features and classifier parameters. Experimental results on an OpenAir-Interface-based testbed show that SG-Defense rapidly mitigates DDoS attacks. It also restores service quality for legitimate users.

INTRODUCTION

5G deployment has expanded the demand for connectivity beyond terrestrial coverage, leading the Third Generation Partnership Project (3GPP) to formalize Non-Terrestrial Networks (NTN) in Release 17 and enhance its specifications in Release 18 [1]. NTN extends coverage to maritime, aerial, and remote environments, supporting naval communications, aviation management, remote IoT deployments, and emergency response. In parallel, the mobile industry has shifted toward open, software-driven RAN architectures, with the Open RAN (O-RAN) Alliance defining interoperable interfaces and introducing the RAN Intelligent Controller (RIC) for programmable, data-driven network control. As NTN and O-RAN evolve along independent

standardization tracks, their integration has emerged as a key direction for building flexible, unified 5G systems. However, NTN's wide coverage and high access density enlarge the attack surface, where disruptions can affect entire regions and compromise mission-critical services. Therefore, ensuring the availability and security of O-RAN NTN deployments remains a key technical challenge.

In current NTN deployments, an attacker can exploit the high access capacity of satellite-connected user equipments (UEs) to inject high-rate uplink traffic, overwhelming the satellite link or the associated ground gateway and disrupting service across a wide region. Prior studies have examined this threat from different perspectives. Usman et al. [2] mitigated ICMP flooding in Geostationary Earth Orbit (GEO) satellite systems via a connection-limiting mechanism at the ground station. In Low Earth Orbit (LEO) constellations, Fratty et al. [3] introduced randomized routing to limit the attacker's ability to steer traffic toward specific inter-satellite links. Salim et al. [4] proposed a distributed federated learning framework enabling satellites and ground-edge nodes to collaboratively train intrusion-detection models. In O-RAN systems, Xavier et al. [5] showed that the Near-Real Time (RT) RIC can collect physical and Medium Access Control (MAC) features from a distributed unit (DU) and train machine-learning (ML) models to detect abnormal traffic. However, this centralized design struggles to react in real time in NTN, where the separation between satellite-based RAN nodes and ground-based RIC components introduces substantial latency. Existing approaches focus on specific attack types or deployment locations and lack a low-latency, scalable, O-RAN-compatible defense mechanism operable directly at the satellite segment.

To address these limitations, we propose SG-Defense, a space-ground distributed architecture for O-RAN NTNs. Rather than relying on a complex standalone detection algorithm, SG-Defense focuses on the architectural placement of defense functions by distributing detection, UE-level attribution, enforcement, and model refinement across the satellite and ground segments to form a closed-loop mitigation workflow. SG-Defense adopts the regenerative payload

This work was supported in part by the National Science and Technology Council (NSTC), Taiwan, under Grant NSTC 115-2221-E-011-101-MY3.

The authors are with the Department of Computer Science and Information Engineering, National Taiwan University of Science and Technology, Taipei 106335, Taiwan. Shin-Ming Cheng is corresponding author.

Digital Object Identifier: 10.1109/MCOMSTD.2026.3731778

concept in 3GPP TR 38.821 [6] and places both the centralized unit (CU) and DU on the satellite, enabling traffic analysis and connection control within the satellite-side RAN segment. The architecture aligns with the O-RAN Alliance's ongoing NTN-related work by adopting a coordinated distributed application (dApp) and eXtended application (xApp) model for programmable, extensible security functions. A dApp on the CU performs real-time inspection of GPRS Tunneling Protocol-User Plane (GTP-U) packets and control-plane signaling, correlating identifiers such as the Cell-Radio Network Temporary Identifier (C-RNTI), 5G Globally Unique Temporary Identifier (5G-GUTI), and Internet Protocol (IP) address to build a unified UE identity mapping that enables precise malicious-flow attribution and immediate DU-side blocking. A complementary xApp on the Near-RT RIC collects telemetry, trains the detection model, and sends updated parameters to the CU via an application-specific E2 procedures, allowing the defense logic to adapt to changing traffic patterns.

This placement separates immediate attack response from model refinement. Compared with centralized ground-side response paths, where telemetry and mitigation decisions traverse the satellite-ground link and incur additional transport delay, SG-Defense removes this delay from the time-critical response path by performing detection and UE identity correlation at the CU and enforcement at the DU before malicious

traffic reaches the ground core network or further depletes satellite RAN resources. The ground Near-RT RIC is reserved for model refinement, making SG-Defense a lightweight, updatable defense prototype aligned with the O-RAN control framework.

To validate SG-Defense, we implemented the proposed framework on a real O-RAN-based 5G testbed using OpenAirInterface (OAI), where the CU and DU emulate a regenerative-payload NTN. We evaluated controlled multi-UE flooding scenarios to assess real-time detection and mitigation under satellite-side DU resource saturation. The results confirm that SG-Defense promptly identifies malicious UEs and restores benign-user throughput under controlled high-rate UDP distributed denial-of-service (DDoS) attacks, validating the proposed space-ground mitigation workflow for volumetric uplink-saturation attacks.

BACKGROUND AND RELATED WORK

O-RAN-enabled NTN remains an emerging integration direction rather than a mature deployment model. Fig. 1 illustrates how 3GPP-defined NTN deployment models can be combined with O-RAN-based programmable control to coordinate satellite-side radio functions and ground-side intelligence through open interfaces.

NON-TERRESTRIAL NETWORK ARCHITECTURE

In 5G-NTN integration, platforms such as high-altitude platform stations and LEO/Medium Earth Orbit (MEO)/GEO satellites act as transparent relay nodes between UEs and the core network (CN), extending coverage beyond conventional terrestrial base stations. NTN provides connectivity to remote or infrastructure-scarce regions and support service continuity when terrestrial infrastructure is unavailable [7]. Coverage and latency strongly depend on platform altitude: GEO satellites offer global coverage with spot-beam diameters of 200–3500 km but incur one-way propagation delays exceeding 250 ms in bent-pipe architectures, whereas LEO satellites project smaller footprints of 50–1000 km with reduced latency of 5–20 ms [7], [8].

Table 1 lists NTN-related technical documents. Among them, 3GPP TR 38.811 [9] and TR 38.821 [6] define use cases and deployment scenarios for extending New Radio (NR) protocol support to NTN environments. According to TR 38.821 [6], NTN payloads are either transparent or regenerative: a transparent payload acts as a radio frequency (RF) repeater that forwards uplink signals to a ground gateway for 5G Node B (gNB) processing, whereas a regenerative payload hosts




FIGURE 1. Integrated O-RAN NTN Architecture

Document ID	Key Technical Scope	Relevance to NTN/O-RAN
3GPP TR 38.811 [9]	NTN channel models & deployment scenarios	Defines link budget and path loss models for satellite channels.
3GPP TR 38.821 [6]	NTN architecture & protocols (Rel-16)	Specifies regenerative payload architectures and split gNB options.
3GPP TS 38.101-5 [10]	UE RF requirements for satellite access	Sets Tx/Rx performance limits for NTN-capable UE.
3GPP TS 38.108 [11]	Satellite Node RF requirements	Defines RF characteristics for satellite.
O-RAN White Paper [8]	O-RAN-based NTN deployments	Outlines use cases for integrating RIC with satellite networks.

TABLE 1. Related Technical Documents.

part or all of the gNB functionality to reduce feeder-link load and shorten the control loop. For deployment flexibility, the gNB may be split into a DU and a CU, with the DU on the satellite and the CU on the ground, but this split introduces challenges for mobility management of the on-board DU and F1 connectivity across ground stations. When the gNB or DU resides on the satellite, the feeder link differs from terrestrial 5G mainly at the physical and data-link layers, where diverse transport-link specifications exist without a unified radio interface standard. Above the network layer, however, the user plane and control plane still follow the conventional 5G protocol stack. TS 38.101-5 [10] and TS 38.108 [11] further specify the minimum RF and performance requirements for UE and satellite transmission.

NON-TERRESTRIAL NETWORK COMBINED O-RAN ARCHITECTURE

O-RAN transforms the traditional closed RAN architecture by disaggregating the base station into the radio unit (RU), DU, and CU and introducing the Near-RT RIC and Non-RT RIC. This open, interoperable design enables flexible deployment, vendor diversity, and data-driven control through standardized interfaces [12]. These features are highly relevant to NTN systems, which must manage wide-area coverage, satellite mobility, and limited on-board resources. In this context, the RIC supports centralized analytics and control for beam orchestration, mobility management, and resource optimization: the Non-RT RIC manages long-term policy, while the Near-RT RIC provides near-real-time control.

The O-RAN Alliance has identified several NTN-specific use cases where RIC-enabled intelligence can enhance overall performance. The recent O-RAN whitepaper [8] highlights satellite handover support, coordinated beam management, and improved adaptation to rapidly changing coverage patterns as key opportunities. O-RAN is also extending the E2 interface to enable consistent control across DU/CU components in both terrestrial and satellite environments, supporting unified RAN management despite heterogeneous transport links. Recent research further explores dApps that execute intelligence directly at DU or CU endpoints [13]. By executing selected intelligence close to the RAN functions, these dApps reduce raw-telemetry transfer to the Near-RT RIC and preserve data locality [13], while xApps on the Near-RT RIC support broader analytics, policy refinement, and model updates. Under standardized E2 procedures, an xApp initiates a *RIC Subscription Request* to obtain telemetry from the CU or DU. Upon receiving the *RIC Subscription Response*, the DU/CU reports telemetry via periodic or event-triggered *RIC Indication messages*. The xApp can also deliver updated parameters or control decisions to the DU/CU via *RIC CONTROL REQUEST* messages, enabling the local dApp or control logic to apply these updates close to the RAN functions.

SECURITY CHALLENGES IN NTN

Existing studies show that satellite-ground communications face confidentiality and integrity risks because strict size, weight, and power constraints

limit the use of computationally intensive cryptographic protocols to protect control and user data [14]. Key management and access control are also challenging to maintain in NTN deployments, where long propagation delays and constantly changing network topologies hinder timely authentication and state synchronization.

LEO systems introduce additional exposure because rapid movement requires frequent handovers and signaling exchanges, which attackers may exploit to manipulate signaling, spoof nodes, or eavesdrop on control messages [8], [14]. DDoS attacks further amplify these vulnerabilities. In NTN, limited on-board processing power and scarce radio resources make even moderate-volume DDoS traffic sufficient to saturate the uplink or downlink and deny service to legitimate users [15].

RELATED WORK

In satellite communication systems, most DoS and DDoS mitigation studies focus on reducing congestion in specific operational scenarios. Usman et al. [2] examined ICMP flooding in GEO satellite networks and designed a threshold-based connection-limiting method at the ground station to filter abnormal traffic. Fratty et al. [3] introduced randomized routing within LEO constellations to reduce an attacker's ability to steer traffic toward targeted inter-satellite links. Salim et al. [4] employed a federated learning framework that allows satellites and ground-edge nodes to train intrusion-detection models while preserving data privacy, achieving strong detection performance across multiple attack classes.

These studies highlight the need for effective DoS defenses in NTN and reveal a growing trend toward distributed intelligence. However, most existing solutions remain tied to specific attack patterns or deployment assumptions and lack the flexibility required for the future integration of terrestrial and non-terrestrial networks. Table 2 compares these approaches in terms of system focus, attack type, mitigation concept, deployment site, and testbed configuration. The comparison illustrates that these methods offer limited portability and insufficient architectural scalability. As a result, NTN systems require defense frameworks that rely on open standards and can adapt to diverse network conditions.

In terrestrial O-RAN systems, Xavier et al. [5] showed that the Near-RT RIC can collect physical and MAC features from the DU and train ML models to detect malicious traffic. This design works well on terrestrial networks but depends heavily on a centralized controller. When applied to NTN, the long physical distance between satellite-based DU nodes and ground-based RIC components can increase E2 interface latency and place the satellite-ground link on the control loop. For NTN deployments, systems therefore need low-latency detection and mitigation mechanisms positioned close to the RAN nodes to compensate for the limitations of centralized control. Although DDoS mitigation mechanisms have been investigated in prior studies on terrestrial and satellite communication networks, prior work has not explicitly addressed how detection, UE-level attribution, enforcement, and model refinement should be distributed across the satellite and

Reference	System Focus	Attack Type	Mitigation Concept	Deployment Location	Testbed
Usman et al. [2]	GEO satellite network	ICMP Flooding	Connect limiting	Ground Station	MATLAB Simulation
Fratty et al. [3]	LEO satellite network	DDoS	Randomized Routing	Satellite Nodes	ICARUS Framework [15]
Salim et al. [4]	Satellite Communications	DoS/DDoS/Zero-Day	DFL IDS	Satellite/Ground Station	Simulation
Xavier et al. [5]	Open RAN	DoS/DDoS	ML using PHY/MAC	Near-RT RIC	srsRAN and USRPs
This article	O-RAN-enabled NTN	DDoS (UDP flooding)	Flow-level detection	Satellite/Ground Station	OAI

TABLE 2. Related work comparison.




FIGURE 2. SG-Defense Architecture and DDoS Mitigation Workflow.

ground segments in regenerative-payload NTNs with O-RAN-enabled programmable control.

SG-DEFENSE

To address the lack of real-time, scalable DDoS mitigation mechanisms in O-RAN NTNs, we introduce SG-Defense, a space-ground distributed framework that rapidly blocks attack traffic at the satellite side while supporting continuous model adaptation on the ground. As shown in Fig. 2, SG-Defense consists of a Space Segment and a Ground Segment that jointly form an integrated mitigation process tailored to NTN latency and resource constraints. SG-Defense is positioned as a lightweight operational prototype rather than an advanced detection engine, emphasizing low-complexity flow-level inference at the satellite CU, immediate mitigation at the satellite DU, and periodic model refinement at the ground-side Near-RT RIC.

The Space Segment deploys a dApp on the CU to execute real-time traffic detection. The dApp inspects GTP-U packets in the user plane and parses control-plane signaling to extract and correlate key UE identity information. These identifiers originate from different layers of the network: the DU assigns a C-RNTI to each UE at the radio interface, and the core network

authenticates the UE using the 5G-GUTI and allocates an IP address for GTP-U transmission. By observing the C-RNTI, 5G-GUTI, and IP address simultaneously, the dApp constructs a complete UE identity mapping table. Once the dApp identifies a malicious IP address, it immediately maps the IP to the corresponding 5G-GUTI and C-RNTI. It forwards the C-RNTI to the DU so the DU can block the associated connection at the earliest possible moment. The CU also stores the malicious 5G-GUTI to prevent the same UE from re-establishing connectivity during later registration attempts.

The Ground Segment deploys an xApp on the Near-RT RIC to perform model training and updates. The xApp continuously collects telemetry from the CU via the E2 interface and refines the model using this data. This feedback loop enables SG-Defense to adapt to evolving attack patterns while maintaining accurate, low-latency detection on satellite-side RAN nodes. In this model refinement loop, the satellite-side CU and ground-side xApp communicate through the RIC Indication and RIC Control procedures, corresponding to the REPORT and CONTROL RIC services, respectively. Both procedures are defined by a single application-specific E2SM advertised by the CU during E2 Setup. The

xApp initiates periodic reporting through a *RIC SUBSCRIPTION REQUEST*, after which the CU periodically sends flow identifiers and feature vectors extracted from locally generated flow records in *RIC Indication Messages*. After retraining, the xApp sends the normalization mean and scale vectors together with the logistic-regression coefficients and intercept to the CU through a *RIC CONTROL REQUEST*, allowing the CU to immediately apply the updated classifier during inference. Since E2SM-KPM supports standardized measurements and E2SM-RC defines existing control actions, neither supports this application-specific data exchange; therefore, we propose a dedicated application-specific E2SM.

DETECTION MODEL ON GROUND SEGMENT

Within the Ground Segment of SG-Defense, the traffic detection model uses logistic regression as its core classifier. Logistic regression provides a lightweight, interpretable solution for binary classification that meets the real-time, resource-aware requirements of the RIC environment. For data input, SG-Defense aggregates sniffed packets into flows and treats each flow as the basic unit of detection. Unlike the traditional 5-tuple definition (source IP, source port, protocol, destination IP, destination port), SG-Defense defines a flow based on the combination of source IP, protocol, and destination IP. The source and destination ports are not part of the flow key; instead, they contribute to derived statistical features.

This flow definition is motivated by the observation that, in many attack scenarios, the source or destination port changes frequently while the attacker and the targeted DU remain fixed. By grouping packets by source IP, protocol, and destination IP, SG-Defense aggregates all packets associated with the same malicious behavior into a single flow, helping the classifier more effectively distinguish regular traffic from anomalous patterns. For feature extraction, SG-Defense uses the following statistics: the flow total number of packets, packets per second, bits per second, the number of unique source ports and unique destination ports, and several statistical measures of packet sizes and inter-arrival times.

SPACE SEGMENT

In the Space Segment of SG-Defense, the system distributes security functions between the CU and the DU to form a hierarchical defense structure. The DU acts as a low-latency enforcement point and blocks malicious UEs as soon as it receives their associated C-RNTIs from the CU. The CU hosts the core detection logic. A dApp deployed on the CU runs the traffic detection model and processes both user-plane and control-plane data. On the user-plane side, the dApp decapsulates captured GTP-U packets to recover the inner IP header, aggregates the packets into the predefined flow representation, and extracts the corresponding features as input to the detection model.

In parallel, the dApp inspects control-plane signaling. By parsing F1 Application Protocol (F1AP) messages, such as the UE CONTEXT SETUP RESPONSE, it obtains the C-RNTI assigned by the DU. To link user-plane traffic to individual UEs, the dApp monitors NG Application Protocol (NGAP) Packet Data Unit (PDU) session management

procedures. During PDU session establishment, it associates each UE's traffic, identified by its GTP Tunnel Endpoint Identifier (TEID), with NGAP-level identifiers such as the PDU session ID and the UE NGAP ID. The dApp also monitors initial registration procedures to obtain a relatively persistent identifier such as the 5G-GUTI or its temporary form, 5G-S-Temporary Mobile Subscriber Identity (5G-S-TMSI). By correlating these identifiers, the dApp constructs a mapping that relates inner IP addresses and TEIDs to the corresponding PDU session ID, UE NGAP ID, and the pair (5G-GUTI, C-RNTI).

When the model identifies a malicious inner IP address, the dApp uses this mapping to retrieve the corresponding C-RNTI and instructs the DU to terminate the associated radio connection. At the same time, it records the malicious UE's 5G-GUTI, enabling the system to reject subsequent registration or service requests from the same UE and thereby strengthening the defense's persistence.

EXPERIMENTAL SETUP AND EVALUATION

The experimental environment uses OpenAir-Interface (OAI) to emulate a 5G network with a standard terrestrial configuration rather than NTN-specific radio parameters because integrated open-source O-RAN NTN solutions are not yet available and standardized connection procedures between a terrestrial Near-RT RIC and a satellite segment remain at an early conceptual stage. Since SG-Defense assumes a regenerative-payload deployment with the dApp, CU, and DU co-located on the satellite, the evaluation focuses on the local detection-and-enforcement loop, including CU-side packet processing, UE identity correlation, and DU-side blocking. However, the setup does not model NTN-specific radio and mobility effects, including Doppler compensation, satellite mobility, timing uncertainty, beam or satellite handover, feeder-link intermittency, and on-board hardware constraints. Therefore, the results validate the local detection-and-enforcement loop under controlled emulation rather than a full NTN deployment. All system components run in a Kubernetes (K8s) cluster using OAI 2025. w32 container images. The CU, DU, dApp, and supporting network functions are co-located in the same local K8s environment, so inter-component communication occurs over a local virtual network rather than an actual satellite-ground transport path. Consequently, the implementation is not intended for quantitative end-to-end control-loop latency evaluation or comparison with a conventional ground-RIC-only mitigation architecture. Instead, the evaluation verifies that detection, UE identity correlation, and DU-side enforcement execute within a local processing loop, eliminating the satellite-ground round trip during attack response. Six UEs attach to the DU via rfsimulator mode, and the DU connects to the CU, which interfaces with the core network. The emulation platform runs on a workstation with an Intel Core i7-13700F 24-core processor, 40 GB of RAM, an NVIDIA GeForce GTX 960 GPU, and 1 TB of storage under Ubuntu 22.04 LTS.

SG-DEFENSE IMPLEMENTATION

The core detection functionality of SG-Defense runs as a dApp deployed on the CU. The dApp integrates two tightly coupled processing pipelines.




FIGURE 3. Confusion matrix of the flow-level detection model used in SG-Defense.

The first pipeline performs UE identity correlation via control-plane signaling, and the second pipeline performs flow-level traffic detection using user-plane packets. For identity correlation, the dApp captures signaling traffic on the CU through the Python pcap library and decodes protocol messages using pycrate. It extracts the C-RNTI from the Initial Uplink RRC message carried in F1AP, obtains the 5G-GUTI from uplink NAS messages in NGAP, and derives the UE IP address from NGAP PDU Session Resource Setup procedures. The dApp correlates these identifiers in real time and stores them in a UE identity mapping table, enabling the DU to enforce C-RNTI-based blocking when the system detects malicious behavior.

For user-plane detection, the dApp continuously captures GTP-U packets on the CU and processes the raw packet data in Python to remove the GTP-U tunnel header before writing the decapsulated IP packets to an in-memory buffer. It aggregates the packets into flow records at fixed intervals using the source IP address, protocol, and destination IP address as the primary key, while the source and destination ports serve as statistical features rather than flow identifiers. To aggregate bidirectional traffic into a single flow, the dApp orders the endpoint IP addresses and records the original direction of the first observed packet as a flow attribute. After computing flow-level statistics, the dApp feeds the resulting feature vectors into a pre-trained classifier for real-time inference. An xApp running on the Near-RT RIC continuously refines the detection model. The xApp collects telemetry and aggregated flow-level features from the CU through *RIC Indication Message*, retrain the model locally, and delivers updated model parameters to the CU via *RIC Control Message*, enabling uninterrupted classifier updates at the dApp.

MODEL TRAINING AND EVALUATION

To construct the dataset for model training and evaluation, we generated traffic directly in the

OAI testbed rather than using an offline public dataset. This testbed-driven approach is adopted because SG-Defense is evaluated as a real-time mitigation workflow, in which detection results must be translated into UE identity correlation and DU-side blocking actions. The evaluation currently focuses on UDP flooding as an initial volumetric uplink-saturation scenario that creates reproducible resource pressure on the satellite DU. To generate malicious traffic, we used the hping3 tool with diverse configurations, including varying packet transmission intervals, packet sizes, and destination port patterns. Both random and fixed destination ports were considered to capture different UDP-flooding behaviors. For benign traffic, we emulated user activities such as social media browsing and video streaming to collect normal traffic under different load conditions. All packets were sniffed at the CU, where the system removed the GTP-U encapsulation and converted the inner IP packets into flow records based on the flow definition.

At the packet level, the dataset contained 55,053 benign packets and 4,582,072 malicious packets, exhibiting a highly imbalanced distribution typical of flooding attacks. After conversion to flow-level data, this imbalance was significantly mitigated. The final dataset consisted of 1,407 benign flows and 733 malicious attack flows, reflecting the high-density nature of DDoS traffic, in which a large volume of attack packets aggregates into a relatively small number of distinct flows. This packet-to-flow transformation not only aligns with the characteristics of flooding attacks but also reduces the data volume processed during inference, enabling faster detection and lower computational overhead on resource-constrained satellite-side CUs.

The flow dataset was randomly shuffled and split into training and testing sets using an 80:20 ratio. We implemented a Logistic Regression classifier using the Python scikit-learn library. Before training and inference, both datasets were standardized using StandardScaler to prevent feature-scale disparities from biasing the learned weights and to improve generalization. The trained model achieved an accuracy of 0.9976, precision of 0.9937, recall of 1.0, and F1-score of 0.9968 on the test set. As shown in Fig. 3, the confusion matrix indicates that only one benign flow was misclassified as malicious among 429 test flows. In contrast, all attack flows were correctly identified. Within the evaluated volumetric UDP-flooding scenario, these results show that the proposed flow-based model can accurately distinguish attack flows from benign flows with zero false negatives and a low false positive rate.

ATTACK SCENARIO AND SIMULATION CONFIGURATION

In NTN environments, volumetric uplink attacks can degrade service availability by exhausting RAN-side processing and transport resources before traffic reaches the ground core network. Therefore, this study adopts high-rate UDP flooding as an initial uplink-saturation threat model for evaluating the real-time SG-Defense mitigation loop. To generate malicious traffic under reproducible conditions, we use hping3 within the OAI emulation environment. Specifically, each compromised UE is configured to launch high-rate




FIGURE 4. Network throughput comparison.

UDP flood traffic with a packet size of 1000 bytes and an inter-packet interval of 10 microseconds. This intentionally aggressive configuration is used to reproduce uplink saturation in the OAI emulation environment and should be viewed as a controlled stress-test setting rather than a model of nominal NTN UE throughput, NTN-specific scheduler behavior, or satellite resource constraints. Uplink traffic is directed to the OAI Traffic Server, which serves as the anchor point for the UE data path.

To ensure synchronized attack initiation and prevent timing inconsistencies among UEs, we implemented an automation script that continuously monitored each UE's state and verified that the virtual interface `oaitun_ue` had been created and that the UE had successfully obtained an IP address. Only after all attacking UEs satisfied these conditions did the script simultaneously trigger the flood traffic. As illustrated in Fig. 4, the attack was launched at $t = 22$ s, providing a clearly defined transition point for evaluating SG-Defense under high-rate DoS conditions. This controlled stress-test reproduces DU-side processing pressure, allowing us to evaluate whether SG-Defense detects malicious uplink traffic, correlates flows with UE identities, and triggers DU-side blocking. Although the current setup focuses on UDP-based volumetric flooding and does not cover all DoS variants, it provides an initial stress test to validate the proposed closed-loop mitigation workflow.

RESULT ANALYSIS

Fig. 4 shows the throughput performance of the benign UE under four experimental conditions. In all experiments, six UEs were attached to the same DU. One UE served as the benign user whose TCP throughput we monitored, while the remaining five UEs acted as either additional benign users or attackers, depending on the scenario. In the single-UE baseline, only the benign UE was active, and its throughput remained stable at approximately 76.8 Mb/s. When all six UEs generated benign traffic, the target UE's average throughput decreased to 16.6 Mb/s due to normal resource contention on the DU. After the five

attacker UEs launched the UDP-flood attack at $t = 22$ s, the benign UE's throughput degraded sharply, fluctuating between 0 and 5.3 Mb/s with an average of 2.9 Mb/s. This behavior matches the expected uplink-saturation effect produced by high-rate flooding traffic directed at the satellite-side DU.

SG-Defense detected malicious flows at around $t = 27$ s, approximately 5 s after the attack onset in Fig. 4, and then instructed the DU to block the corresponding UEs based on their C-RNTIs. Within roughly 8 s after detection, the benign UE's throughput increased from below 5 Mb/s. It recovered toward the single-UE baseline, which indicates that the proposed workflow reclaimed DU radio resources previously monopolized by the attackers. These results highlight the effectiveness of CU-side attack identification combined with F1AP-derived C-RNTI enforcement at the DU, which preserves QoS for legitimate users and mission-critical traffic before the attack escalates.

CONCLUSION

This article introduced SG-Defense, a space-ground distributed defense framework for DDoS mitigation in O-RAN NTNs. By combining satellite-side detection and UE-level identity correlation with ground-side model refinement, SG-Defense establishes a defense workflow that addresses the latency and bandwidth constraints of NTN systems. OAI-based emulation results show that SG-Defense rapidly identifies and isolates malicious UEs while restoring benign-user throughput under a controlled high-rate UDP-flooding scenario, demonstrating the feasibility of the proposed closed-loop mitigation workflow for volumetric uplink-saturation attacks. Future work will extend SG-Defense beyond controlled emulation through improved feature engineering, model selection, and packet-level statistics for lower-latency detection of diverse attacks. Evaluation under NTN-specific radio and mobility conditions remains the next key step, particularly under beam and satellite handover, where satellite or gateway transitions may interrupt the still-immature E2 association with the ground Near-RT RIC and make telemetry collection

from satellite-side E2 nodes intermittent. In addition, future studies will investigate the effects of Doppler shifts, long propagation delays, additional telemetry sources, and multi-satellite deployments to further strengthen the proposed SG-Defense framework for O-RAN NTN.

REFERENCES

- [1] M. A. Jamshed et al., "Non-terrestrial networks for 6G: Integrated, intelligent, and ubiquitous connectivity," *IEEE Commun. Standards Mag.*, vol. 9, no. 3, pp. 86–93, Sep. 2025, doi: 10.1109/MCOMSTD.2025.3572635.
- [2] M. Usman et al., "Mitigating distributed denial of service attacks in satellite networks," *Trans. Emerg. Telecommun. Technol.*, vol. 31, no. 6, p. 3936, Mar. 2020, doi: 10.1002/ett.3936.
- [3] R. Fratty, Y. Saar, R. Kumar, and S. Arnon, "Random routing algorithm for enhancing the cybersecurity of LEO satellite networks," *Electronics*, vol. 12, no. 3, p. 518, Jan. 2023, doi: 10.3390/electronics12030518.
- [4] S. Salim et al., "Deep-federated-learning-based threat detection model for extreme satellite communications," *IEEE Internet Things J.*, vol. 11, no. 3, pp. 3853–3867, Feb. 2024, doi: 10.1109/JIOT.2023.3301626.
- [5] B. M. Xavier, M. Dzaferagic, D. Collins, G. Comarella, M. Martinello, and M. Ruffini, "Machine learning-based early attack detection using open RAN intelligent controller," in *Proc. IEEE Int. Conf. Commun.*, Rome, Italy, May 2023, pp. 1856–1861, doi: 10.1109/icc45041.2023.10279349.
- [6] *Solutions for NR To Support Non-terrestrial Networks (NTN)*, document (TR) 38.821 V16.2.0, RAN3 Working Group, Apr. 2023.
- [7] J. Baranda et al., "On the architectural split and radio intelligence controller placement in integrated O-RAN-enabled non-terrestrial networks," *IEEE Open J. Commun. Soc.*, vol. 6, pp. 8571–8590, 2025, doi: 10.1109/OJCOMS.2025.3616383.
- [8] O-RAN Alliance. (Feb. 2025). *Deployments of O-RAN-based Non-terrestrial Networks*. [Online]. Available: https://mediastorage.o-ran.org/ecosystemresources/O-RAN-2025.04.02.WP.O-RAN_NTN_Deployments-v08.4.pdf
- [9] *Study on New Radio (NR) To Support Non-terrestrial Networks*, document (TR) 38.811 V15.4.0, RAN Working Group, Oct. 2020.
- [10] *NR; User Equipment (UE) Radio Transmission and Reception; Part 5: Satellite Access Radio Frequency (RF) and Performance Requirements*, document 38.101-5 V18.11.0, 3rd Generation Partnership Project (3GPP) RAN4 Working Group, Oct. 2025.
- [11] *NR; Satellite Access Node Radio Transmission and Reception*, document 38.108 V18.8.0, Oct. 2025.
- [12] M. Polese et al., "Understanding O-RAN: Architecture, interfaces, algorithms, security, and research challenges," *IEEE Commun. Surveys Tuts.*, vol. 25, no. 2, pp. 1376–1411, Jan. 2023, doi: 10.1109/COMST.2023.3239220.
- [13] A. Lacava et al., "DApps: Enabling real-time AI-based open RAN control," *Comput. Netw.*, vol. 269, Sep. 2025, Art. no. 111342, doi: 10.1016/j.comnet.2025.111342.
- [14] I. Ahmad et al., "Security of satellite-terrestrial communications: Challenges and potential solutions," *IEEE Access*, vol. 10, pp. 96038–96052, 2022, doi: 10.1109/ACCESS.2022.3205426.
- [15] G. Giuliani et al., "ICARUS: Attacking low earth orbit satellite networks," in *Proc. USENIX ATC*, Jul. 2021, pp. 317–331.

BIOGRAPHIES

CHENG-FENG HUNG (Member, IEEE) (d10915002@mail.ntust.edu.tw) received the M.E. degree in information technology and applications from the College of Science and Engineering, National Quemoy University, Kinmen, Taiwan, in 2019, and the Ph.D. degree from the Department of Computer Science and Information Engineering, National Taiwan University of Science and Technology, Taipei, in 2025. He visited Warsaw University of Technology in 2022 and 2025. His research interests are O-RAN security and AI security. He has received a Student Travel Grant at the IEEE 2023 WPMC.

PO-HUNG CHEN (m11415021@mail.ntust.edu.tw) received the B.S. degree in computer science and information engineering from the National Taiwan University of Science and Technology, Taiwan, in 2025, where he is currently pursuing the M.S. degree in computer science and information engineering. His research interests include open radio access network security and mobile networks.

SHIN-MING CHENG (Member, IEEE) (smcheng@mail.ntust.edu.tw) received the B.S. and Ph.D. degrees in computer science and information engineering from National Taiwan University, Taipei, Taiwan, in 2000 and 2007, respectively. Since 2012, he has been on the Faculty of the Department of Computer Science and Information Engineering, National Taiwan University of Science and Technology, Taipei, where he is currently a Professor. He served as the Deputy Director-General of Administration of Cyber Security, Ministry of Digital Affairs, from 2022 to 2025, and as the Acting Director-General in 2025. His current interests are mobile network security, IoT system security, malware analysis, and AI robustness. He has received the IEEE Trustcom 2020 Best Paper Awards.