Mitigating Security Challenges in 5G Wireless Networks

Friday Oodee Philip-Kpae [1,*], Agbotiname Lucky Imoize [2], Kennedy Chinedu Okafor [3,4,5,6,7], Kelvin Anoh [3]
[1] Department of Electrical and Electronic Engineering, Faculty of Engineering, Rivers State University, Port Harcourt, Nigeria
[2] Department of Electrical and Electronics Engineering, Faculty of Engineering, University of Lagos, Akoka, Lagos 100213, Nigeria
[3] School of Engineering, University of Chichester, Bognor Regis, PO21 1HR, U.K
[4] Department of Electrical and Electronic Engineering Science, University of Johannesburg, Johannesburg, South Africa
[5] Imperial College Business School – South Kensington Campus, Imperial College London, Exhibition Road, London SW7 2AZ, UK
[6] Federal University of Allied Health Sciences, Enugu – Dental Avenue, Trans-Ekulu, Enugu State, Nigeria
[7] Department of Engineering (Faculty of Science and Engineering), Manchester Metropolitan University – Dalton Building, Chester Street, Manchester M1 5GD, UK
philip-kpae.fo@ust.edu.ng; aimoize@unilag.edu.ng; kennedy.okafor@ieee.org; k.anoh@chi.ac.uk
*Correspondence: philip-kpae.fo@ust.edu.ng

Abstract*— 5G networks increasingly rely on key enabling technologies such as Software-Defined Networking (SDN), Network Function Virtualization (NFV), Multi-Access Edge Computing (MEC), and end-to-end network slicing to deliver heterogeneous services with strict quality-of-service (QoS) guarantees. However, programmability, multi-tenancy, and distributed edge–cloud operation significantly expand the attack surface. At the same time, traditional rule-based and reactive security mechanisms remain slow to adapt and may violate latency constraints during mitigation. This paper addresses the problem of QoS-compliant, closed-loop security control for sliced SDN/NFV infrastructures. We propose an AI-assisted, cross-layer security orchestration framework that integrates epoch-wise telemetry with ML-based risk estimation and formalizes mitigation as a Constrained Markov Decision Process (CMDP). The CMDP controller selects enforceable actions—slice isolation, rate limiting, traffic rerouting, and key reconfiguration—while explicitly satisfying latency/overhead constraints, and executes them via SDN flow-rule updates and NFV policy/VNF reconfiguration. Simulation results over 50 decision epochs demonstrate effective response to an injected high-risk event: risk spikes to 0.95 at epoch 15, after which the controller drives risk toward ≈0.10 while maintaining latency below the 40ms QoS bound (with a transient rise during mitigation and subsequent stabilization). The reward trajectory briefly degrades during disruption but recovers and converges to a positive long-term return, indicating stable constraint-aware operation. This work provides (i) a deployable cross-layer orchestration architecture for sliced networks, (ii) a QoS-constrained CMDP decision model that converts risk signals into actionable SDN/NFV controls, and (iii) empirical evidence that adaptive mitigation can reduce security risk without sacrificing service guarantees.

Index Terms—*5G Security, beyond 5G/6G, Network Slicing, SDN/NFV Security, Constrained Markov Decision Process (CMDP), AI-Enabled Mitigation, Mitigation Security, Intrusion Detection*.

I. Introduction

The evolution from 5G toward beyond-5G/6G is accelerating the shift to softwarized, cloud-native mobile networks built on Software-Defined Networking (SDN), Network Function Virtualization (NFV), Multi-Access Edge Computing (MEC), O-RAN, and end-to-end network slicing. These technologies enable rapid service rollout, elastic resource sharing, and differentiated QoS across diverse verticals, including ultra-reliable low-latency communications, massive IoT/IoMT, and vehicular networking. At the same time, increased programmability, multi-tenancy, and the distributed edge–cloud footprint enlarge the attack surface, complicate policy management, and intensify the impact of failures and security incidents. Recent studies, therefore, highlight a growing convergence between resilience and security in next-generation networks and anticipate ML-driven automation as a key enabler for 6G operations, while noting that practical closed-loop security control for operational-slicing environments remains immature [1]–[3]. Although security mechanisms for sliced and virtualized networks have advanced, traditional security approaches remain predominantly static and reactive.

Many deployments still rely on fixed rules, signature-based detection, and manual or partially automated response strategies that are difficult to scale and slow to adapt to rapidly evolving threats such as DDoS, cross-slice abuse, and control-plane exploitation. In softwarized 5G/6G settings, these approaches face three significant limitations. First, they often lack end-to-end visibility and coordinated response across the RAN–core–edge continuum, particularly in multi-domain slicing and federation scenarios. Second, they rarely ensure QoS feasibility during mitigation; actions such as throttling, rerouting, or isolation may inadvertently violate latency and service guarantees. Third, while several ML-based methods improve detection or resource coordination, they frequently stop short of delivering a network-enforceable, constraint-aware control policy that maps risk signals to real-time SDN/NFV actions [2], [4], [6], [12]. Consequently, an orchestration gap persists. Security orchestration is widely recognized as essential for zero-touch operation. Yet, the policy engine for selecting optimal mitigation actions under coupled security–QoS objectives remains insufficiently specified and validated [3], [10].

To address these challenges, this paper proposes a closed-loop, constraint-aware security orchestration framework for SDN/NFV-enabled network slicing in 5G/6G-oriented deployments. The proposed approach couples ML-based risk estimation with a Constrained Markov Decision Process (CMDP) controller that selects mitigation actions while explicitly satisfying latency/overhead constraints. Specifically, we present a layered security architecture spanning the sliced infrastructure, integrate monitoring with risk scoring, and execute mitigation through SDN/NFV control primitives (e.g., slice isolation, rate limiting, rerouting, and key reconfiguration). We then formulate orchestration as a CMDP, where the state captures evolving security risk, actions represent enforceable mitigation options, and constraints encode QoS requirements. This design enables systematic, constraint-respecting optimization of security responses under uncertainty, rather than ad hoc or purely heuristic mitigation. This paper makes the following contributions:

1. Development of a cross-layer security architecture for sliced SDN/NFV networks as an end-to-end orchestration framework spanning RAN–core–edge that integrates monitoring, risk estimation, and enforceable mitigation through SDN/NFV control.
2. Formulation of QoS-constrained security decision model for security orchestration as a CMDP, defining a concrete state–action space and an objective that balances risk reduction against operational cost while enforcing latency/overhead constraints.
3. Validation of constraint-feasible mitigation framework, showing the CMDP-driven policy responsiveness to high-risk episodes by triggering SDN/NFV mitigation actions while maintaining QoS feasibility (e.g., sustaining latency below a

predefined bound during mitigation and reducing risk thereafter).

The remainder of this paper is structured as follows: Section II reviews the background and related work on security, resilience, and orchestration in sliced 5G/6G systems. Section III presents the system model and the proposed cross-layer security architecture and details the CMDP formulation, including states, actions, reward design, and constraint definitions, and describes the solution workflow. Section IV presents the results and discussion from the simulations. Section V concludes the paper, outlining limitations and future research directions.

II. Related Work

Prior work motivates intelligent, automated, and resilient 5G/6G slicing but falls short of a deployable, QoS-aware security control loop. Specifically, [1] synthesizes resilience and failure-analysis concepts for next-generation networks and classifies failure drivers and countermeasures around resource depletion, security vulnerabilities, and availability, yet remains primarily taxonomic and does not translate risk indications into implementable mitigation decisions in softwarized slicing. Building on this motivation, the study [2] surveys the emerging role of ML in 6G. It highlights key challenges (e.g., trust, robustness, and deployability). Still, it does not instantiate a safety-constrained decision process that binds ML-derived risk to network-enforceable actions under service constraints. Complementarily, the work [3] positions Security Orchestration (SCO) as an essential automation plane for 5G/B5G and discusses its integration with SDN/NFV, slicing, O-RAN, and ZSM, while noting that practical deployments still lack sufficient closed-loop security automation; however, the policy engine for selecting optimal actions under competing security–QoS objectives remains unspecified. In parallel, the work [4] demonstrates that learning-based orchestration can outperform static protocols by leveraging multi-agent DRL to learn adaptive, interpretable inter-slice coordination for O-RAN resource management. Still, it does not address adversarial risk, security constraints, or mitigation actions (e.g., isolation or key reconfiguration). On the security side, the study [5] surveys threats and defenses for intelligent slicing and shows how multi-tenancy complicates isolation, yet does not provide runtime-adaptive optimization that jointly enforces QoS and security under uncertainty; similarly, the work [6] consolidates DDoS vulnerabilities and mitigation approaches for 5G/5G-A and beyond, but remains a catalogue rather than an integrated controller that selects mitigations subject to slice SLA constraints. Finally, the report [7] advances cloud-native slice federation across administrative domains through standardized interfaces and validated orchestration workflows, but does not realize an adaptive cross-domain security loop. In contrast, our paper bridges these gaps by operationalizing security vulnerability–induced failures as a constraint-aware CMDP implemented via SDN/NFV controllers, where ML-based risk signals drive a concrete mitigation action set (e.g., isolation, rate limiting, rerouting, and key reconfiguration) while explicitly satisfying latency/overhead constraints (e.g., sustaining latency below a 40ms bound during high-risk episodes and reducing risk thereafter).

In the cloud-native slicing context, the paper [8] consolidates SDN-based MEC-enabled slicing architectures and key enablers (virtualization, softwarized control, and edge–cloud coordination) and outlines open issues toward practical realization; however, security is addressed mainly as standalone mechanisms rather than as an end-to-end, closed-loop orchestration problem spanning RAN–core, edge. Accordingly, our work builds on this substrate by introducing a layered SDN/NFV-coordinated security architecture and optimizing adaptive mitigation so enforcement remains QoS-compliant in edge-intensive sliced deployments. Along similar lines, the paper [9] proposes Slice-block for secure slicing by coupling context-aware authentication handover with a DAG-blockchain trust layer in an edge-assisted SDN/NFV-6G environment, strengthening cross-slice authentication continuity; yet, it offers limited guidance on runtime, risk-driven mitigation decisions once anomalies or attacks emerge (e.g., when to isolate, reroute, or throttle under SLA bounds). Our paper complements Slice-block by adding a constraint-aware CMDP decision layer over concrete SDN/NFV actions (including key reconfiguration) that can be triggered by authentication/context signals while preserving latency constraints. Further, the work [10] broadens the discussion by detailing practical barriers to next-generation security orchestration—telemetry/monitoring, interface standardization, scalability, privacy, and multi-domain policy enforcement—thereby clarifying why operational deployment remains difficult; nevertheless, it does not formalize these requirements into a unified optimization/control model with enforceable constraints. We bridge this theory-to-practice gap by codifying orchestration as a CMDP-controlled, closed-loop mitigation process and demonstrating the feasibility of constraints and recovery under attack. Finally, the study [11] examines dynamic SFC deployment for SDN/NFV-based satellite–terrestrial integrated vehicular networks under highly time-varying connectivity and resource constraints, advancing learning-based orchestration for service continuity; however, security-aware control under QoS constraints is not central. Our contribution provides the security counterpart by embedding risk into the control objective via CMDP and ensuring latency/overhead feasibility—capabilities that can be layered atop SFC placement/orchestration pipelines to harden satellite–terrestrial slices against evolving threats.

In the detection layer, the work [12] develops and implements an anomaly-based intrusion detection system for IoMT networks, contributes a purpose-built dataset, and evaluates multiple ML/novelty (outlier) detection algorithms while demonstrating deployment practicality through runtime-cost reporting. However, the study remains focused on detection performance. It does not extend to a QoS-constrained, SDN/NFV-executable mitigation loop that translates risk into slice-level actions in softwarized 5G/6G infrastructures. Our paper addresses this gap by using ML-derived risk scores as the CMDP state signal and selecting enforceable SDN/NFV mitigation actions (e.g., slice isolation, rate limiting, rerouting, and key reconfiguration) under explicit latency/overhead constraints. From a methodological standpoint, the findings [13] frame robust autonomous-driving decision-making as a CMDP and introduce a safe-RL approach (RECPO) to improve robustness while respecting safety constraints, reinforcing CMDP/safe-RL as a principled paradigm for constrained optimization. Nevertheless, its scope is outside networking and does not model SDN/NFV slicing, attack-driven risk dynamics, or concrete mitigation primitives. Our work bridges this domain gap by transferring CMDP-based safe decision-making into 5G/6G security orchestration, where risk spikes drive SDN/NFV control actions while maintaining QoS feasibility (e.g., sustaining latency below a 40ms bound during high-risk episodes).

III. Methodology

This section presents the methodological workflow used to realize and evaluate the proposed closed-loop, QoS-constrained security orchestration for SDN/NFV-enabled network slicing. We describe the resources and modules that constitute the framework, the research design adopted, the simulation/experimental setup used to emulate a 5G/B5G slicing environment under attack, and the data collection technique used to quantify security–QoS trade-offs. The proposed framework is implemented as a set of tightly coupled functional modules aligned with softwarized 5G/B5G network operation. First, an emulated SDN/NFV-based slicing infrastructure provides a multi-slice service environment in which each slice exhibits its own traffic dynamics, QoS state, and resource utilization profile. A telemetry and monitoring component then collects epoch-wise measurements—such as latency, utilization, traffic indicators, and security alerts—which are ingested by an AI/ML risk estimation module to compute a quantitative risk score reflecting the likelihood and severity of ongoing threats. Based on this evolving state, a CMDP-driven decision engine selects mitigation actions while

enforcing QoS constraints. An SDN/NFV enforcement interface implements the chosen actions through SDN flow control and NFV policy/VNF configuration primitives, including slice isolation, rate limiting, rerouting, and key reconfiguration. Collectively, these components instantiate the paper's cross-layer architecture (data plane, SDN control plane, and NFV MANO/security functions) and support an end-to-end observe–decide–act closed-loop security orchestration workflow.

A simulation-driven experimental design is employed to examine how the proposed constraint-aware orchestration framework reacts to time-varying threats while preserving service guarantees. Within the simulation horizon, an abrupt attack event is injected to induce a pronounced increase in the estimated risk, thereby emulating the high-risk episode considered in this study. At each decision epoch, the closed-loop controller sequentially observes network telemetry, computes a risk estimate, selects an appropriate mitigation action, enforces it, and logs the resulting state transition for subsequent evaluation. Performance is then analyzed in terms of (i) the magnitude and speed of risk reduction, (ii) preservation of QoS under a predefined latency bound, (iii) the security–cost trade-off reflected in the reward trajectory, and (iv) long-term behavior using discounted-return and constraint-feasibility metrics. We emulate a softwarized 5G/B5G network slicing environment over a discrete horizon of T decision epochs, consistent with the temporal granularity used in the paper's figures and tables. The experimental setup considers multiple coexisting network slices with heterogeneous QoS requirements and enforces a latency constraint evaluated against a 40-ms threshold in the feasibility analysis. To stress the system, an attack event is injected during the run to induce a high-risk episode, reflected in a pronounced risk spike around epoch 15, after which the controller applies mitigation actions to restore acceptable risk levels while preserving QoS. The selected actions are assumed to be implementable using standard SDN/NFV primitives, including flow-rule updates, slice-isolation policies, routing adjustments, and security/key reconfiguration operations. All simulation parameters and epoch-wise observations (including risk, reward, and action snapshots) are logged and summarized in the corresponding tables and figures.

At each decision epoch, data are captured through structured logging of both controller outputs and environment responses. Specifically, state traces record the ML-estimated risk score ρ_t measured latency L_t, relevant QoS indicators, resource utilization, and alert/telemetry features. In parallel, action traces store the selected mitigation a_t together with its scope (per-slice or global). The framework further maintains outcome traces, including the resulting next state s_{t+1}, the computed reward $\alpha S_t - \beta D_t - \delta O_t - \eta V_t$, and indicators of mitigation overhead. From these logs, we derive evaluation metrics comprising (i) the risk trajectory and recovery time, (ii) the latency/QoS violation rate (constraint-feasibility ratio), (iii) the temporal trade-off among reward components, and (iv) the expected discounted return to reflect long-term stability and convergence. This data collection procedure supports reproducible assessment of whether the proposed CMDP-driven orchestration reduces security risk while maintaining QoS feasibility, consistent with the risk/latency/reward/feasibility analyses reported in the results section. The network parameters are given in Table 1, and the system parameters are presented in Table 2.

Table 1: Network parameters [14]

Parameter	Symbol	Value	Description
Time Epochs	T	50	Total decision steps in the loop.
Discount Factor	γ	0.95	Priority given to long-term security.
Security Weight	α	0.6	Importance of security improvement (S_t).
Latency Penalty	β	0.3	Impact of service degradation (D_t).
Overhead Weight	δ	0.1	Computational cost of mitigation (O_t).
Max Latency	$D_{\max}$	40 ms	The QoS constraint for slices.
Risk Score	ρ_t	0.1 to 1.0	0.0 is secure; 1.0 is a total compromise.
Stochastic Noise	w_t	$\mathcal{N}(0,\sigma^2)$	Random network fluctuations.

Table 2: System parameters for simulation

Epoch (t)	Risk (ρ_t)	Latency (D_k)	Reward (r_t)	Action (a_t)
1	0.7	25.0 ms	0.12	Observe/Monitor
15	0.95	28.2 ms	-0.45	Attack Detected
16	0.65	38.5 ms	0.08	Slice Isolation Applied
30	0.22	22.1 ms	0.42	Key Reconfiguration
50	0.15	18.4 ms	0.51	Optimal Steady State

A. *System model*

The network consists of user equipment, radio access network (RAN), core network, and edge/cloud infrastructure. Network slicing enables logical isolation of services with heterogeneous requirements. Adversaries may launch jamming, spoofing, DDoS, slice-isolation, VNF compromise, and AI-driven evasion attacks. We adopt a layered security architecture spanning physical, RAN, core, and edge/cloud layers, coordinated by SDN/NFV controllers. The system model is visualized in Figure 1.

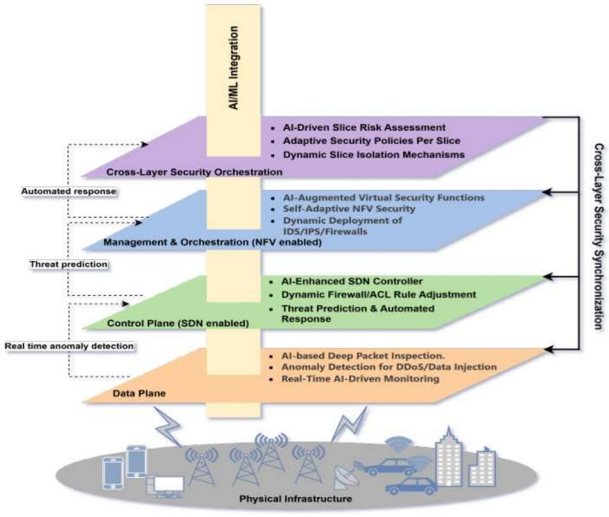

Figure 1: Cross-layer AI/ML-enabled security architecture for SDN/NFV-based 5G/6G network slicing. The framework integrates AI/ML across the data plane, SDN control plane, NFV management & orchestration layer, and cross-layer security orchestration to support real-time anomaly detection, threat prediction, and automated response, with end-to-end synchronization across the physical infrastructure [15].

B. *System setup*

We consider a wireless 5G and beyond communication network supporting heterogeneous services over a shared physical infrastructure. The network is softwarized using SDN and NFV and supports network slicing. Let $N = \{1, 2, \ldots, N\}$ denote the set of network nodes and $\mathcal{U} = \{1, 2, \ldots, U\}$ the set of user devices. The network supports slices is $S = \{s_1, s_2, \ldots, s_k\}$, each with slice-specific QoS and security requirements. Time is discretized into decision epochs where $t \in \{0, 1, 2, \ldots\}$.

C. *System state definition*

At each time step t, the global network state is defined as: $s_t = [L_t, B_t, Q_t, R_t, C_t, A_t]$, where $L_t, = \{L_k(t)\}_{k=1}^{k}$: traffic load and latency per slice, $B_t = \{B_k(t)\}_{k=1}^{T}$: bandwidth utilization, $Q_t = \{Q_k(t)\}_{k=1}^{k}$: QoS/QoE indicators, $R_t = \{\rho_k(t)\}_{k=1}^{k}$: estimate security risk scores, C_t: available compute and virtualization resources, and A_t: detected attack indicators and alerts. The state captures both network performance and security posture across slices.

D. *Action space*

At each decision epoch, the SDN/NFV controller selects a mitigation action: $a_t \in A$, where A = {slice isolation, rate limiting, traffic rerouting, key reconfiguration, monitor}. Actions may be applied globally or per slice.

E. *State transition dynamics*

Network dynamics evolve according to: $s_{t+1} = f(s_t, a_t) + w_t$, where $f(\cdot)$ models traffic evolution, mitigation impact, and attacker response, and w_t represents stochastic disturbances.

F. *Reward function*

The immediate reward is defined as a weighted trade-off: $r_t = \alpha S_t - \beta D_t - \delta O_t - \eta V_t$, where S_t: achieved security improvement, D_t: service degradation (latency, packet loss), O_t: operational and computational overhead, V_t: SLA or policy violations. Higher reward corresponds to effective threat mitigation with minimal performance loss.

G. *Optimization objective*

The controller seeks a policy π that maximizes the expected discounted return: $\max_{\pi} E_{\pi}[\sum_{t=o}^{T} \gamma^t r_t]$, where $\gamma \in (0, 1)$ is the discount factor.

H. *Constraints*

The optimization is subject to: $D_t(t) \leq D_k^{max} \quad \forall k \in S$ (1), $O_t \leq O^{max}$ (2), $a_t \in A_{safe}$ (3), $s_t \in S_{operational}$ (4). These constraints ensure: QoS guarantees, bounded overhead, compliance with security policies, and operational feasibility. The formulation defines a constrained Markov decision process (CMDP) suitable for AI-based adaptive security control. AI models estimate risk scores, cap R sub t, and guide action selection. Mitigation actions are enforced via SDN/NFV controllers. All datasets are publicly available and contain no personal data. Experiments follow data minimization and reproducibility best practices. The procedure is detailed in Algorithm 1.

Algorithm 1: Standard Adaptive Security Mitigation Loop (CMDP-Controlled)

Require: Time horizon T; slices S; max latency constraint L_{max}; risk estimator $f(\cdot)$; policy $\pi_\theta(a|s)$; action set A; controller interface $C_{\text{SDN/NFV}}$.

Enquire: Mitigation decisions $\{a_t\}_{t=1}^{T}$ that reduces risk while respecting QoS constraints

1: Initialize controller state, security baselines, and policy parameters θ
2: **for** $t = 1$ to T **do**
3: **Observe:** collect telemetry to form a global state s_t (traffic/QoS per slice, latency, resource usage, alerts)
4: **Estimate risks:** $\rho_t \leftarrow f(s_t)$
5: **Check QoS:** compute current latency L_t (per-slice or aggregate)
6: **Select action**
7: **if** $\rho_t < \rho_{th}$ **and** $L_t \leq L_{max}$ **then**
8: $a_t \leftarrow$ Monitor
9: **else**
10: Choose $a_t \sim \pi_\theta(a|s)$ subject to feasibility and constraints (e.g., prefer least-disruptive feasible mitigation)
11: **end if**
12: **Enforce:** apply a_t using $C_{\text{SDN/NFV}}$ (install flow rules, isolate slice, update VNF policy, rotate keys)
13: **Measure Outcome:** observe next state s_{t+1} and updated (ρ_{t+1}, L_{t+1})
14: **Reward (optional):** compute $r_t \leftarrow w_s \Delta\rho_t - w_t \cdot$ penalty L_{t+1} − Overhead (a_t)
15: **Policy Update (optional):** update θ using CMDP-safe learning (if enabled)
16: **end if**
17: **return** $\{a_t\}_{t=1}^{T}$

IV. Results And Discussions

This section reports the simulation results for the proposed CMDP-based security orchestration and discusses their implications. We examine the temporal behavior of risk, latency/constraint feasibility, selected mitigation actions, and rewards to evaluate how the controller responds to injected high-risk episodes while preserving QoS guarantees.

A. *Network State Evolution*

Initially, the system maintains a baseline risk of approximately 0.7 with a latency of around 25ms. A critical event occurs at epoch 15, when a simulated attack spikes the risk score to 0.95. In response, the adaptive security controller applies mitigation actions, evidenced by the blue latency curve rising to nearly 27ms as resources are diverted to defense. Over the subsequent 10 epochs, the risk is successfully mitigated back down to a stable level of 0.1. At the same time, the latency eventually settles at a consistent floor of 12ms, remaining well below the 40ms maximum threshold $D_{\max}$) indicated by the red dashed line. The network state evolution is given in Figure 2.

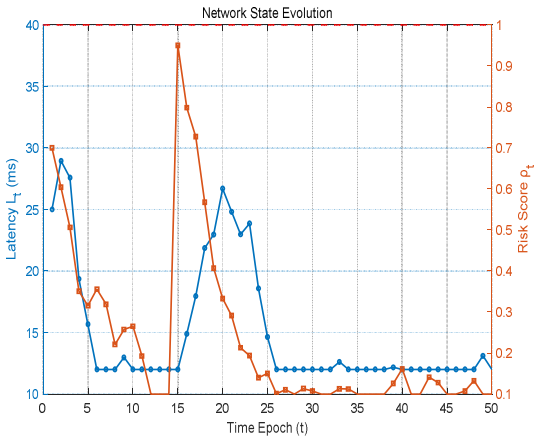

Figure 2: Network state evolution: it illustrates the dynamic relationship between latency (L_t) and security risk (ρ_t) over 50 epochs within a 5G network environment.

B. *Reward Components Trade-off*

Figure 3 illustrates the reward trade-off over 50 epochs. Initially, the system exhibits a rapid improvement, reaching a peak weighted reward value of 0.54 by epoch 14. At epoch 15, a sudden disruption causes the total reward to collapse to zero, indicating a critical system event or instability. Despite this abrupt decline, the model demonstrates strong recovery capability, quickly stabilizing and regaining a reward value of 0.54 by epoch 25. During this steady-state phase, Security Gain emerges as the dominant contributor, maintaining a consistent value of approximately 0.49. At the same time, the combined impact of Latency Penalty and Overhead is

significantly reduced to about 0.05. This behavior confirms the model's robustness, resilience, and ability to sustain efficient long-term performance despite transient disturbances.

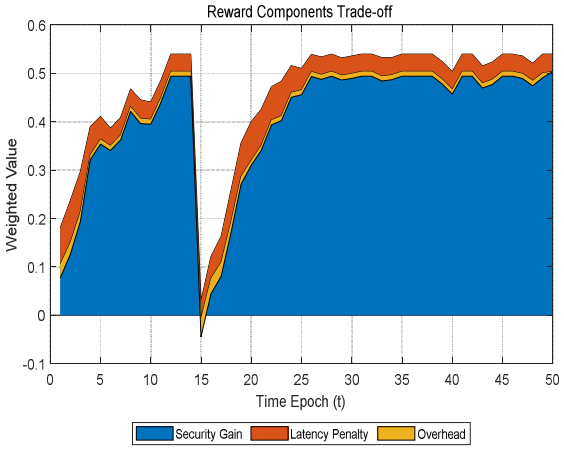

Figure 3: Reward Components Trade-off: visualizes the dynamic performance and security trade-offs of the proposed adaptive mitigation loop, and the graph shows a significant security breach at epoch 15.

C. *Constraint Feasibility Space Analysis*

The constraint feasibility space plot of Figure 4 illustrates the relationship between Latency $D_k(t)$ and Overhead O_t. Quantitative data show that all feasible solutions maintain latency below a critical threshold of approximately 32, as indicated by the pink-shaded region, which denotes a constraint violation. Within the possible range, overhead values remain relatively low, generally below 0.4. Multiple data points cluster at a low latency of 12 with overheads ranging from 0 to 0.38. As latency increases toward 29, overhead fluctuates between 0.1 and 0.3. The color gradient indicates an additional metric ranging from 0 to 0.45, correlating with the system's operational efficiency.

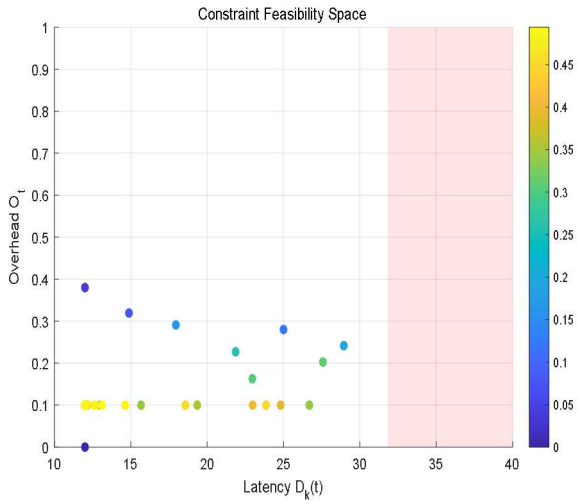

Figure 4: Constraint feasibility space

D. *Expected Discounted Return*

Figure 5 demonstrates the effectiveness of the adaptive security framework within a constrained Markov decision process. The Network State Evolution shows a critical attack at epoch 15, where the risk score ρ_t spikes to approximately 0.95. The controller immediately executes mitigation, causing latency L_t to rise to 27ms, successfully remaining below the 40 ms QoS constraint. This response is reflected in the Reward Components Trade-off, where the instantaneous reward rt briefly turns negative due to increased service degradation and overhead. Finally, the Expected Discounted Return illustrates robust long-term policy convergence, achieving a cumulative value of nearly 7 by epoch 50.

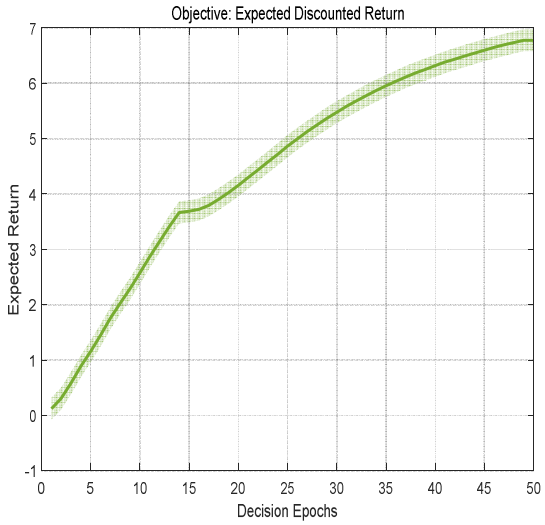

Figure 5: Expected discounted return: it demonstrates the adaptive security framework's effectiveness within a constrained Markov decision process.

Table 3 compares the proposed CMDP-based orchestration with a monitor-only baseline. It shows that after the attack peak (risk ≈ 0.95), the CMDP policy significantly reduces risk, improves reward, and maintains the latency constraint (<40ms). At the same time, the baseline remains elevated due to the lack of mitigation. The baseline corresponds to a monitor-only operation (no adaptive mitigation after attack detection). Risk remarks are percentage reductions relative to the baseline risk at the same epoch. Reward improvement is reported relative to the absolute magnitude of the baseline reward.

Table 3: Summary of Key Results and Baseline Comparison

Key result (metric)	Standard baseline (monitor-only / no mitigation)	Proposed CMDP orchestration	Remark (%)
Peak risk at attack (epoch 15), ρ	0.95	0.95	0.0% (reference)
Immediate post-mitigation risk (epoch 16), ρ	0.95 (persists)	0.65	31.6% risk reduction
Mid-run stabilized risk (epoch 30), ρ	0.95 (persists)	0.22	76.8% risk reduction
End-run risk (≈epoch 50), ρ	0.95 (persists)	≈0.10	89.5% risk reduction
Reward at attack (epoch 15), R	−0.45	−0.45	0.0% (reference)
Recovered reward (epoch 30), R	−0.45 (remains degraded)	0.42	+193.3% improvement

Latency constraint satisfaction	Not guaranteed (no control)	< 40ms (maintained)	100% compliance (observed)
Expected discounted return (by epoch 50)	N/A	≈7.0	N/A

V. Conclusion

This study demonstrates a principled, system-level approach to securing 5G and beyond wireless networks through an adaptive security architecture coordinated by SDN/NFV controllers. By formalizing threat mitigation as a Constrained Markov Decision Process (CMDP), the framework effectively balances stringent security requirements with service performance metrics, including latency and operational overhead. Simulation results confirm the robustness of the adaptive mitigation loop, particularly during high-risk events where security scores peaked at 0.95. The controller successfully reduces risk to 0.1 while maintaining latency well below the 40ms threshold, highlighting the model's practical feasibility for heterogeneous services. Furthermore, the convergence of the expected discounted return to a cumulative value of nearly 7 by epoch 50 validates the policy's long-term effectiveness in preserving operational stability under stochastic disturbances. Overall, this research establishes a scalable foundation for intelligent, policy-driven security in softwarized network infrastructures. The primary limitation of this study is that the evaluation is simulation-driven within an emulated SDN/NFV slicing environment, which necessarily abstracts several real-world factors such as noisy/incomplete telemetry, heterogeneous vendor interfaces, multi-domain policy conflicts, adversarial adaptation, and the operational overhead of mitigation execution in production 5G/6G deployments (e.g., O-RAN and cloud-native MANO). Future work will therefore focus on validating the proposed CMDP-driven security orchestration on an end-to-end experimental testbed or high-fidelity digital twin using realistic traffic mixes and attack traces, extending the formulation to multi-domain and federated slicing scenarios with explicit trust and policy constraints, integrating robust/adversarial learning to harden both risk estimation and decision-making, and improving scalability through online learning and explainable policy outputs that support standardized, zero-touch security orchestration.

References

1. S. Bi, X. Yuan, S. Hu, K. Li, W. Ni, E. Hossain, X. Wang, "Resilience and Failure Analysis in Next-Generation Communication Networks: A Contemporary Survey," in *IEEE Transactions on Network Science and Engineering*, vol. 13, pp. 2793-2821, 2026, doi: 10.1109/TNSE.2025.3620950.
2. M. Anis Oukebdane, A. F. M. Shahen Shah, A. Kalam Azad, J. Ekoru and M. Madahana, "Unraveling the Nexus of ML and 6G: Challenges, Opportunities, and Future Directions," in *IEEE Access*, vol. 13, pp. 114934-114958, 2025, doi: 10.1109/ACCESS.2025.3585051.
3. S. Batewela, M. Liyanage, E. Zeydan, M. Ylianttila and P. Ranaweera, "Security Orchestration in 5G and Beyond Smart Network Technologies," in *IEEE Open Journal of the Computer Society*, vol. 6, pp. 554-573, 2025, doi: 10.1109/OJCS.2025.3563619.
4. F. Rezazadeh H. Chergui, S. Siddiqui, J. Mangues, H. Song, W. Saad, M. Bennis, "Intelligible Protocol Learning for Resource Allocation in 6G O-RAN Slicing," in *IEEE Wireless Communications*, vol. 31, no. 5, pp. 192-199, October 2024, doi: 10.1109/MWC.015.2300552.
5. F. Salahdine, Q. Liu and T. Han, "Towards Secure and Intelligent Network Slicing for 5G Networks," in *IEEE Open Journal of the Computer Society*, vol. 3, pp. 23-38, 2022, doi: 10.1109/OJCS.2022.3161933.
6. S. Hoque, A. Aydeger, E. Zeydan and M. Liyanage, "A Survey on Distributed Denial-of-Service Attack Mitigation for 5G and Beyond," in *IEEE Open Journal of the Communications Society*, vol. 6, pp. 5840-5879, 2025, doi: 10.1109/OJCOMS.2025.3586199.
7. M. Dalgitsis, N. Cadenelli, M. A. Serrano, N. Bartzoudis, L. Alonso and A. Antonopoulos, "Cloud-Native Orchestration Framework for Network Slice Federation Across Administrative Domains in 5G/6G Mobile Networks," in *IEEE Transactions on Vehicular Technology*, vol. 73, no. 7, pp. 9306-9319, July 2024, doi: 10.1109/TVT.2024.3362583.
8. S. D. A. Shah, M. A. Gregory and S. Li, "Cloud-Native Network Slicing Using Software Defined Networking Based Multi-Access Edge Computing: A Survey," in *IEEE Access*, vol. 9, pp. 10903-10924, 2021, doi: 10.1109/ACCESS.2021.3050155.
9. I. H. Abdulqadder and S. Zhou, "SliceBlock: Context-Aware Authentication Handover and Secure Network Slicing Using DAG-Blockchain in Edge-Assisted SDN/NFV-6G Environment," in *IEEE Internet of Things Journal*, vol. 9, no. 18, pp. 18079-18097, 15 Sept.15, 2022, doi: 10.1109/JIOT.2022.3161838.
10. S. Batewela, P. Ranaweera, M. Liyanage, E. Zeydan and M. Ylianttila, "Addressing Security Orchestration Challenges in Next-Generation Networks: A Comprehensive Overview," in *IEEE Open Journal of the Computer Society*, vol. 6, pp. 669-687, 2025, doi: 10.1109/OJCS.2025.3564788.
11. D. Li *et al*., "Dynamic SFC Deployment for SDN/NFV-Based Satellite–Terrestrial Integrated Vehicular Networks," in *IEEE Transactions on Vehicular Technology*, vol. 74, no. 12, pp. 19568-19582, Dec. 2025, doi: 10.1109/TVT.2025.3586653.
12. G. Zachos, G. Mantas, K. Porfyrakis, J. Manuel Camões Sobral de Bastos and J. Rodriguez, "Anomaly-Based Intrusion Detection for IoMT Networks: Design, Implementation, Dataset Generation, and ML Algorithms Evaluation," in *IEEE Access*, vol. 13, pp. 41994-42028, 2025, doi: 10.1109/ACCESS.2025.3547572.
13. R. Zhao, Z. Chen, Y. Fan, Y. Li, and F. Gao, "Towards robust decision-making for autonomous highway driving based on safe reinforcement learning", *Sensors*, *24*(13), p.4140, 2024, https://doi.org/10.3390/s24134140.
14. 3GPP, "System architecture for the 5G System (5GS)," *3rd Generation Partnership Project (3GPP)*, Technical Specification (TS) 23.501, v17.0.0, Mar. 2021.
15. Z. Allaw, O. Zein, and A. M. Ahmad, "Cross-Layer Security for 5G/6G Network Slices: An SDN, NFV, and AI-Based Hybrid Framework". *Sensors*, *25*(11), p.3335, 2025, https://doi.org/10.3390/s25113335.