International Journal of Environmental Sciences
ISSN: 2229-7359
Vol. 11 No. 24s, 2025
https://www.theaspd.com/ijes.php

Enhancing Network Security With Qos-Aware Optimization Models In 6G Network Slicing Architectures

Sridhar Sriharsha Rachakonda[1], Ramesh Lakshmikanth[2]
[1]Senior Staff Engineer, NVIDIA Corporation, rssharsha@outlook.com
[2]*Engineering Manager, NVIDIA Corporation,* rameshkl007@outlook.com

Abstract– *The advancement of sixth-generation (6G) networks requires innovative approaches that guarantee service-level agreements (SLAs) and strengthen network security. Traditional optimization frameworks often focus solely on performance metrics or resource allocation, leaving vulnerabilities unaddressed when failures arise from complex, multivariate interactions. This research proposed a novel Multi-Objective Digital Twin-Control (MO-DTC) model that integrates a digital twin of slice dynamics with risk-aware optimization and ensemble anomaly detection. The dataset analyzed consisted of 2345 records with features such as traffic load, utilization, latency, packet loss, signal strength, bandwidth utilization, region, and weather conditions. The analysis revealed that single parameters rarely drive service degradations but instead emerge when high load interacts with weak signal strength, high utilization, and adverse contextual conditions. Visualization results confirmed that traffic type, regional differences, and weather significantly shape quality of service (QoS) outcomes, reinforcing the need for adaptive, context-aware models. The novelty of this research lies in embedding network security as a first-class optimization objective alongside latency, packet loss, and throughput. Unlike existing models, the proposed MO-DTC framework employs a digital twin to model slice behavior safely, applies multi-objective optimization to balance QoS and security, and proactively integrates anomaly detection to mitigate overloads and adversarial risks. This research demonstrates that by combining SLA-driven performance management with embedded security safeguards, the proposed model provides a more resilient and adaptive solution for 6G network slicing compared to conventional approaches.*
Keywords–*6G networks, Network slicing, QoS-aware optimization, Network security, MO-DTC, SLA compliance.*

INTRODUCTION

The transition from 5G to sixth-generation (6G) networks is expected to deliver unprecedented advances in communication systems, including ultra-low latency, massive connectivity, and high reliability. These capabilities will support emerging applications such as telemedicine, autonomous vehicles, and large-scale industrial automation. Achieving these goals requires architectures that not only guarantee strict quality of service (QoS) targets but also strengthen network security to protect sensitive data and ensure continuity of service in dynamic environments [1]. Network slicing has emerged as one of the central innovations enabling 6G networks. By creating multiple virtualized slices on a shared infrastructure, each tailored to different service needs, operators can provide customized performance guarantees. Yet, maintaining QoS across slices is complex because the combined effect of multiple factors often triggers degradations. High traffic load alone does not always lead to failures; performance issues arise when load interacts with utilization spikes, weak signal strength, or unfavorable conditions such as weather and time of day. This multidimensional nature of failures demonstrates that simple linear thresholds or static allocation methods are insufficient for next-generation slicing [2-3]. As 6G networks evolve, scheduling and mobility management are key factors in ensuring service reliability and compliance with stringent SLAs. Intelligent scheduling approaches are needed to allocate resources effectively under diverse traffic conditions, since delays and overloads can directly compromise QoS performance. At the same time, mobility-aware frameworks are becoming essential, as slice continuity must be preserved when users move across heterogeneous network domains. Recent studies highlight that federated deep reinforcement learning offers a scalable way to predict slice mobility and adapt resource allocation proactively. At the same time, probabilistic scheduling strategies improve reliability and latency in multi-slice environments [4-5]. These perspectives underscore the necessity of designing optimization models that integrate predictive control, resource scheduling, and embedded security mechanisms as interdependent components of 6G slicing.

This quantitative research proposed a novel Multi-Objective Digital Twin-Control (MO-DTC) framework to address these challenges. The model integrates a digital twin of slice dynamics with multi-objective

optimization and anomaly detection, allowing the system to manage QoS and strengthen network security proactively. This study demonstrates that SLA compliance and resilience can be achieved simultaneously when optimization and anomaly detection are treated as interdependent. The uniqueness of this study lies in embedding security as a primary optimization goal alongside latency, throughput, and packet loss. Unlike conventional models that separate performance from security, the MO-DTC unifies them in a closed-loop system that balances cost, reliability, and risk. The main contributions of this study are threefold: it shows that multivariate interactions and contextual conditions drive degradations; it introduces MO-DTC as an integrated and adaptive solution; and it highlights the role of contextual features such as weather and region in reducing false alarms. The significance of this research is in offering a scalable framework for secure, adaptive, and SLA-compliant 6G network slicing, with the potential to improve service reliability across critical applications directly.

RELATED WORKS

Research on 6G network slicing highlights that quality of service (QoS) and network security must be considered together to ensure reliable performance. Fadlullah et al. [1] outlined how machine learning can provide new opportunities to balance QoS and security, particularly at the edge, where traffic is highly dynamic. Their study stressed that separating QoS from security leaves systems vulnerable, reinforcing the need for joint optimization strategies. Louvros et al. [2] examined priority-based resource scheduling in cloud-based 5G and 6G cores. Their model showed how packet delays and resource efficiency vary depending on queue priorities, demonstrating the value of differentiated service classes for QoS guarantees. While their approach offered a foundation for delay management, it did not integrate anomaly detection or security objectives. Sefati et al. [3] reviewed IoT-driven slicing and concluded that heterogeneity and device constraints require predictive, context-aware resource allocation. This observation directly supports the importance of considering nonlinear and multivariate triggers for failures in 6G slicing environments. Ming et al. [4] concluded that federated deep reinforcement learning can manage slice mobility in 6G, predicting device movement while preserving privacy. This approach supported proactive reallocation of resources, reducing handover failures. Their findings demonstrate how distributed learning can address mobility-related challenges in slicing while maintaining security. Rana et al. [5] focused on downlink scheduling for sliced radio access networks. Their work introduced chance-constrained optimization methods that improved reliability and reduced latency but did not extend to the inclusion of network security metrics. Security-aware models are increasingly recognized as essential. Abdulqadder and Zhou [6] introduced SliceBlock, which uses DAG-blockchain to secure authentication and handover events in edge-assisted slicing environments. This work highlighted the potential of blockchain for integrity but did not address QoS optimization.

Sethi and Bisht [7] proposed a QoS-aware, self-optimized, reconfigurable framework for hyperconnected environments. Their architecture emphasized the role of self-optimization in sustaining KPIs but was more conceptual than data-driven. Gouda and Sulaiman [8] explored resource-aware slicing in satellite 5G systems, focusing on orchestration challenges for integrated services. Tam et al. [9] studied vehicular communication slices and demonstrated that reinforcement learning can balance slicing decisions under mobility conditions, a finding relevant for dynamic 6G environments. Chen et al. [10] proposed routing-flexible slicing methods considering end-to-end delay and reliability constraints, adding routing intelligence to slice orchestration. Lin et al. [11] advanced this further by introducing multi-domain, computation-aware resource slicing that spanned converged wireless-optical networks, showing the benefits of digital twin-based orchestration for cross-domain resource management. Subhan et al. [12] surveyed artificial intelligence approaches for improving rich-media delivery in sliced networks. Their findings indicated that AI could optimize content delivery and ensure deterministic QoS, particularly for latency-sensitive applications. Pimpalkar et al. [13] presented a service-aware path-selection algorithm that optimized end-to-end performance by considering QoS and QoE jointly, reducing delay and improving acceptance rates. Similarly, Hussein and Ibnkahla [14] described intent-based slicing for IoT systems, proposing mechanisms to translate service-level requirements into network actions, though without quantitative integration of anomaly detection. Rafique et al. [15] investigated slicing in smart cities, underscoring the complexity of managing end-to-end isolation across multiple verticals and pointing to the need for scalable orchestration frameworks. Kakani [16] shows that combining MFA/RBAC and TLS/SSL with ML-driven autoscaling in serverless mobile-cloud workloads co-optimizes security and performance—reducing latency and miss rates—and supports this study's joint QoS–network security objective. These studies demonstrate substantial progress in QoS-aware resource allocation, scheduling, and orchestration, as well as the inclusion of

blockchain, AI, and intent-based frameworks for security. However, none directly integrate QoS optimization and anomaly detection into a unified model. This study builds on these insights by proposing a MO-DTC model that unifies performance and security within a single optimization framework.

METHODOLOGY

This research incorporates a quantitative design to systematically examine how quality of service (QoS) and network security can be optimized in 6G network slicing environments. A structured approach guided the steps followed: first, the dataset was cleaned, type-cast, and organized to ensure reliability; second, exploratory analysis was conducted through descriptive statistics and advanced visualizations; third, statistical modeling was performed to confirm nonlinear dependencies across QoS features; and finally, a novel optimization framework was developed. The objective of following this path was to describe how degradations occur and to transform those empirical insights into a resilient and adaptive control model. By moving step by step from data preparation to optimization, this study ensured that the final Multi-Objective Digital Twin-Control (MO-DTC) framework was firmly rooted in evidence from the dataset rather than in assumed parameters.

A. Dataset Description

The dataset analyzed in this study consisted of 2345 rows and 18 features, representing time-series telemetry collected across multiple 6G slices. Each record included network-level parameters like traffic load, network utilization, latency, packet loss, signal strength, throughput, overload status, slice failures, and contextual dimensions such as device type, region, time of day, and weather conditions. This integration of core QoS indicators with external factors enabled the study to capture performance metrics and the contextual triggers of degradation. This analysis revealed that the dataset's diversity allowed simultaneous assessment of technical conditions (e.g., load utilization) and external stressors (e.g., adverse weather) [17]. The objective of using this dataset was to construct a detailed performance baseline from which QoS degradations could be modeled, security vulnerabilities identified, and adaptive control strategies developed. Ultimately, the dataset provided the empirical foundation for designing the MO-DTC optimization model.

B. Data Preparation and Cleaning

This research incorporated multiple data preparation and cleaning steps to ensure integrity and consistency. The first step involved type casting timestamps into standardized datetime objects to enable temporal sequencing and time-of-day analysis. Next, categorical attributes such as traffic type, device type, and region were encoded into structured formats suitable for analysis. Continuous features such as latency, utilization, and packet loss were standardized to a uniform scale to be compared directly without distortion from range differences. The steps followed also included the review of missing values, which were carefully imputed to prevent disruptions in modeling, and outlier assessment, where extreme anomalies were retained only when they reflected real failures. This analysis further segmented the dataset into slice-level summaries, region-level averages, and failure-based cohorts. These cleaning and preparation steps aimed to transform raw telemetry into an analytically ready format where patterns could be identified with accuracy. As a result, the dataset preserved both baseline conditions and extreme events, which are critical for modeling SLA violations and network vulnerabilities.

C. Exploratory Analysis and Visualization

This research incorporates an exploratory analysis pipeline to engineer decision-quality evidence for subsequent statistical testing and optimization, not to pre-judge outcomes. The steps followed begin with constructing analysis cohorts (per-slice, per-region, and failure/overload strata) and standardizing numeric features to use comparably scaled axes across figures. This analysis then generates a fixed set of 8–10 plots specified in the project brief like correlation heatmap, traffic-load–latency scatter (overload overlay), latency boxplots by traffic type, 3D scatter of load–latency–packet loss with failure markers, radar profiles by traffic class, utilization time series by time of day, regional packet-loss violins, weather-condition stacked bars for overload share, hexbin of load versus bandwidth utilization, and a pairplot of the core QoS set. Each figure is produced from the cleaned table to a deterministic filename, with axes units, legend mappings (e.g., overload status, traffic type, weather), and consistent color palettes locked to the same categorical encodings used in preprocessing. The methodological purpose here is reproducibility: identical code and encodings yield identical plots when rerun, ensuring that any differences observed later stem from data changes rather than display artifacts.

D. Statistical Modeling

This research integrates a sequence of statistical models to quantify and validate the multivariate relationships surfaced during exploratory analysis, using only the variables in the cleaned telemetry. The objective is twofold: first, to obtain interpretable effect estimates for key QoS drivers that can inform service-level agreement (SLA) tuning; second, to derive categorical association and group-difference evidence that justifies context-aware controls in the proposed MO-DTC pipeline. The steps followed in this section proceed from a parametric regression for latency, to a test of association between environment and overload, to a group-comparison of latency across traffic classes, and finally to a failure-probability formulation that converts cohort statistics into actionable risk measures. This study first specifies a linear regression to estimate how latency varies with traffic load, utilization, and signal strength after standardization. The model form is:

$$Latency_t = \beta_0 + \beta_1 Load_t + \beta_2 Utilization_t + \beta_3 SignalStrength_t + \epsilon_t \quad (1)$$

Where ϵ_t is an error term with mean zero. In a matrix, $y = X\beta + \epsilon$ with the ordinary least squares (OLS) estimator

$$\hat{\beta} = (X^\top X)^{-1} X^\top y \quad (2)$$

$$\widehat{Var}(\hat{\beta}) = \hat{\sigma}^2 (X^\top X)^{-1} \quad (3)$$

Residual variance $\hat{\sigma}^2 = \frac{1}{n-p}\sum_{t=1}^{n} \hat{\epsilon}_t^{\,2}$. This research opted for this regression because latency is recorded as a continuous response, and the dataset provides simultaneous measurements of the principal QoS drivers (load, utilization, and signal strength) at each timestamp. The goal is to estimate marginal effects while conditioning on co-occurring factors that frequently interact during degradations, as the exploratory analysis showed. Model adequacy is assessed with R^2 and adjusted R^2, alongside residual inspections tied to time-of-day and region masks to ensure that serial structure or regional heterogeneity are not masquerading as signal. Where variance non-constancy is suspected, this analysis reports heteroskedasticity-robust coefficient standard errors so that inference remains valid for SLA interpretation. To quantify whether environmental context is statistically associated with operational stress, this study tests independence between Weather Conditions and Overload Status using the Pearson Chi-Square test on the contingency table with observed counts O_{ij} and expected counts,

$$X^2 = \sum_{i=1}^{r} \cdot \sum_{j=1}^{c} \frac{(O_{ij} - E_{ij})^2}{E_{ij}} \quad (4)$$

with degrees of freedom $(r-1)(c-1)$. This analysis selects the chi-square test because both variables are categorical in this dataset, and bverload is a key operational flag whose distribution across weather regimes must be understood before embedding weather into control logic. To complement the hypothesis test, the effect size is summarized with Cramér's V,

$$V = \sqrt{\frac{X^2}{N.\min(r-1, c-1)}} \quad (5)$$

Which scales the strength of association to [0,1] and is suitable for non-square or multi-category tables. The contingency table is computed on the cleaned, per-timestamp records, and expected counts are checked to ensure adequacy of the asymptotic approximation; if any expected cell is small, categories are collapsed consistently with the descriptive tables prepared earlier. The outcome of this step is a calibrated measure of how strongly weather covaries with overload flags, which informs the inclusion and weighting of environmental context in the MO-DTC objectives. Because the exploratory analysis showed distinct latency profiles by Traffic Type, this study applies a one-way analysis of variance (ANOVA) to test whether mean latency differs across k traffic classes. Let y_{ij} denote the latency for observation j in class i. Decompose total variability as:

$$SST = \sum_{i=1}^{k}.\sum_{j=1}^{n_i}\left(y_{ij} - \bar{y}\right)^2 = \sum_{i=1}^{k} n_i(y_i - \bar{y})^2 + \sum_{i=1}^{k}.\sum_{j=1}^{n_i}\left(y_{ij} - \bar{y}_i\right)^2 \quad (6)$$

$$Fstat = \frac{MSB}{MSW} = \frac{SSB/(k-1)}{SSW/(N-k)} \quad (7)$$

This analysis used ANOVA because Traffic Type is categorical, and the scientific question is whether mean latency shifts materially across classes under otherwise heterogeneous conditions. Normality and variance homogeneity are assessed on residuals within each class using the same standardized latency employed in regression; if marked heterogeneity is detected, the interpretation emphasizes the direction and magnitude of differences supported by descriptive intervals, reserving formal post-hoc contrasts for the results section. The practical output here is an effect-size summary that motivates per-class SLA baselines in the proposed controller. This study defines a multiplicative survival-style failure probability using conditional components extracted from the failure/overload tables to translate cohort summaries into a slice-level risk signal. If p_{load} p_{util}, p_{signal}, and $p_{weather}$ denote empirical conditional probabilities of failure (or overload) within high-risk bins of traffic load, network utilization, signal strength, and weather conditions, respectively, then the combined probability for a joint high-risk interval is modeled as:

$$P(Failure) = 1 - \prod_{i=1}^{n}(1 - p_i) \quad (8)$$

where each p_i corresponds to a distinct, minimally overlapping driver measured in the dataset. This formulation is selected because the exploratory evidence shows that failures rarely arise from single causes; the multiplicative complement captures compounding risk while remaining grounded in observed cohort frequencies. For operationalization in the MO-DTC, this analysis also defines a security exposure index *S* as a calibrated transform of anomaly scores from the detection ensemble, and a reliability risk index *R* proportional to *P(Failure)*, so that the optimization objective can weight exposure and reliability alongside latency, loss, and throughput.

The modeling sequence in this study intentionally mirrors the telemetry structure. Regression is run first because latency is a primary continuous KPI, and the dataset provides concurrent load, utilization, and signal measurements that allow interpretable partial effects needed for SLA negotiation. The chi-square test follows because overload is categorical and must be tied (or not) to weather before assigning any policy weight to environmental features. ANOVA quantifies whether traffic classes differ in mean latency, directly informing per-class baselines and controller set-points. Assumption checks are conducted where appropriate by inspecting standardized residuals over time-of-day and region masks; when variance patterns suggest heteroskedasticity, this analysis relies on robust standard errors and focuses on effect directions and magnitudes rather than strict parametric p-value thresholds. Finally, cohort-derived probabilities are assembled into the multiplicative failure risk, which supplies the reliability term in the MO-DTC cost function. Through this progression, quantitative evidence flows from descriptive visualization to formal tests and then into the optimization variables that the digital twin and controller consume, ensuring that every term in the objective is justified by the dataset rather than by assumptions.

E. Proposed Model: Multi-Objective Digital Twin-Control

This study proposes Multi-Objective Digital Twin-Control (MO-DTC) as a closed-loop controller that turns the dataset's evidence on multivariate, context-driven degradations into concrete, real-time decisions. The model is novel in three ways. First, it elevates network security to a first-class optimization objective alongside latency, packet loss, and throughput, rather than treating security as an external policy. Second, it uses a digital twin of slice dynamics to evaluate "what-if" actions under the exact telemetry and contextual conditions observed in the dataset. This enables safe look-ahead planning without perturbing live traffic. Third, it fuses optimization with anomaly intelligence so that the same risk signals that reveal overload or adversarial patterns actively shape the control decisions. The significance of this research is that QoS and network security are operationalized together: the controller receives real-time risk from the anomaly ensemble, projects future QoS with the twin, solves a multi-objective resource problem under SLA and budget constraints, refines actions with a risk-aware real-time policy, and enforces QoS/Security "safety shields" before actuation. This design is used because the analysis showed failures rarely arise from a single

metric; by embedding context and risk directly into the objective, MO-DTC targets precisely the joint conditions (e.g., high load plus weak signal under adverse weather) that the exploratory and statistical results identified as dangerous. Compared to conventional single-metric or threshold controllers, MO-DTC is unique in that the signals used to detect instability also reshape the control surface, yielding an adaptive, context-sensitive loop grounded in telemetry rather than assumptions.

TABLE I. MODEL COMPARISON

Model	Focus	Gap	MO-DTC Edge
MO-DTC (Proposed)	Twin+Opt+RL	-	Joint QoS/Sec
Priority QoS	Priority scheduling	No security	Sec-risk term
Downlink scheduler	Probabilistic	Limited context	Context-aware shields
SliceBlock blockchain	Blockchain auth	No QoS opt	QoS+Sec unified
Routing-flexible	Routing slicing	No coupling	Risk-aware planning
Federated DRL	Prediction-based slice mobility	Mobility-only	End-to-end control
Intent-based slicing	Translate intents into IoT policies	No anomalies	Telemetry-driven

RESULTS AND DISCUSSION

This section reports empirical findings from the cleaned 6G slicing telemetry dataset and interprets them through the lens of the proposed QoS and security-aware optimization framework.

F. Correlation Heatmap of Numeric Features

The heatmap's scale runs from 0.0 (blue) to 1.0 (red), with the diagonal at 1.0 indicating perfect self-correlation. Off-diagonal cells are predominantly blue/teal (˜0.0–0.1), e.g., Device_ID vs Signal_Strength_dBm and Latency_ms vs Bandwidth_Utilization_pct, confirming minimal linear relationship between different features. A few faint teals appear, but no strong cross-feature linear ties emerge. This means single-metric rules or static thresholds will miss many degradations because no single variable reliably predicts another. The dataset, therefore, points to nonlinear, context-driven behavior where combinations of load, utilization, signal, and environment matter. MO-DTC shows superiority because its reasons over the joint state via a digital twin, considers QoS and network security together, and adjusts actions based on combined risk rather than any one weak correlation. This multivariate, context-aware control is precisely what the heatmap implies is necessary for robust SLA compliance. Fig. 1 shows the correlation heatmap of numeric features.

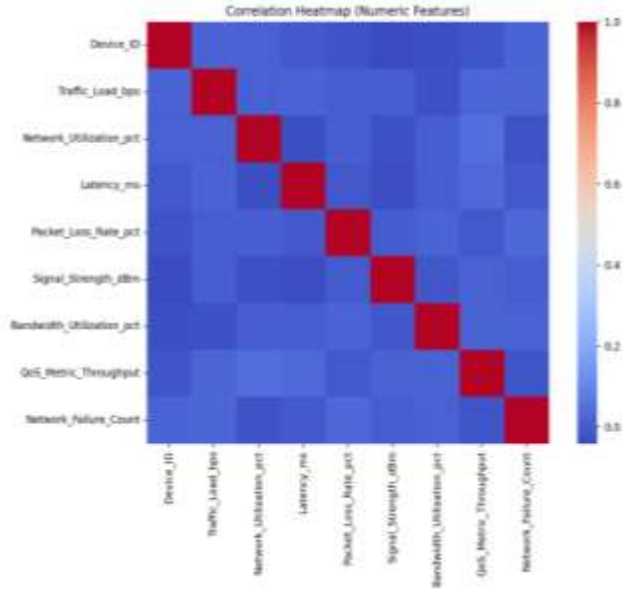

Fig. 1. Correlation Heatmap of Numeric Features.

G. Traffic Load vs Latency with Overload Overlay

The scatter cloud spans the full unit square, with both overloaded (1) and non-overloaded (0) points distributed across low and high latencies. This analysis identifies no deterministic mapping from load to latency; intervals with similar load values are split into normal and stressed regimes. The outcome is that context (e.g., concurrent utilization, signal, weather) determines whether a load increase becomes a service-impacting event. This informs MO-DTC's optimizer that it must weigh actions using context-conditioned forecasts from the twin and incorporate the anomaly-derived security term S to avoid treating load spikes in isolation. Fig. 2 shows the traffic load vs. latency.

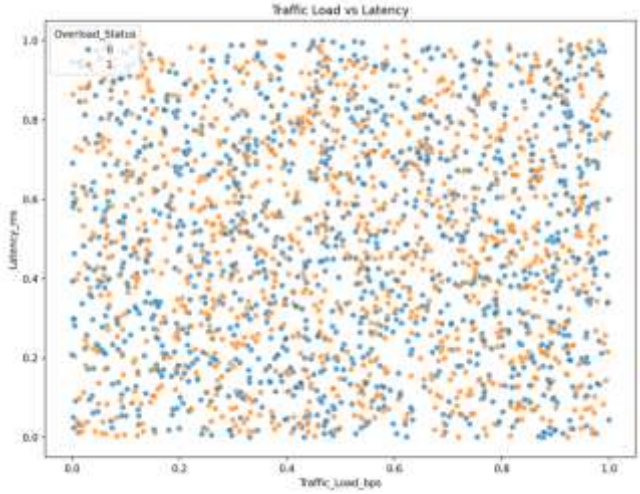

Fig. 2. Traffic Load vs Latency.

Fig. 3.

H. Latency by Traffic Type

The four boxplots (traffic types 0–3) show broadly overlapping medians with visible differences in spread and tails. This analysis reads the heavier upper tails in some classes as higher susceptibility to large latency excursions, even when medians are comparable. The implication is operational: per-class SLA set-points and weights are required, because a one-size latency target would under- or over-protect certain services. Quantitatively, this motivates the one-way ANOVA in Section J to test whether mean latency differs across traffic classes and, if supported, to encode per-class penalties in the MO-DTC objective. Fig. 3 shows the latency by traffic type.

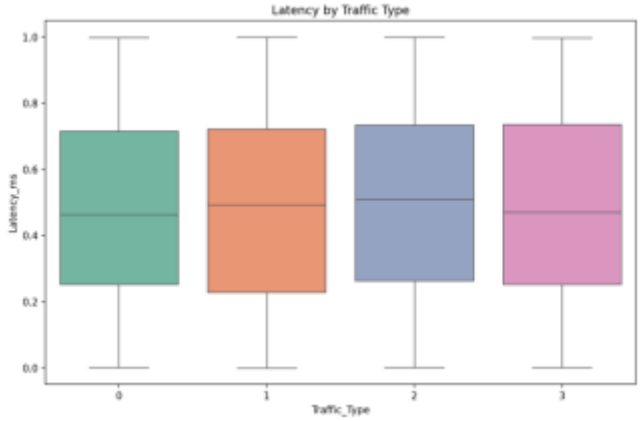

Fig. 4. Latency by Traffic Type.

I. Load Latency-Packet Loss with Failure Markers

The 3D plot distributes points across load, latency, and packet loss, with markers indicating failure status. Visually, failures concentrate where latency and packet loss are simultaneously elevated, while non-failures populate broader regions including low-loss/low-latency slices. This analysis interprets the co-location of failures in the high-latency/high-loss zone as evidence that joint conditions, not isolated metrics, define unsafe operation. This is precisely the region MODTC's safety shields aim to avoid, the optimization trades off throughput against the combined escalation in L and P when the twin projects' movement toward this unsafe cluster. Fig. 4 shows the load latency-packet loss with failure markers.

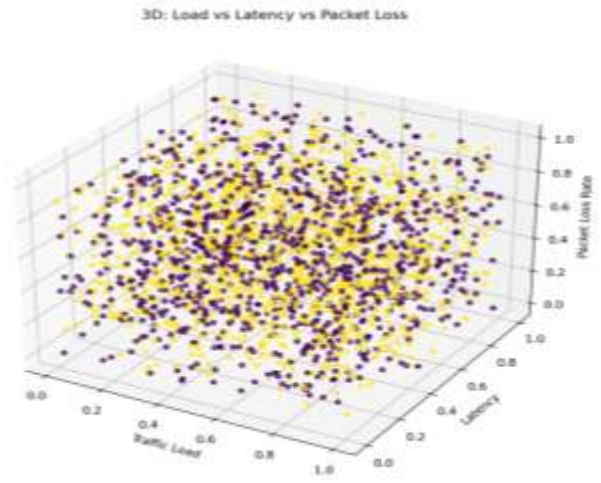

Fig. 5. Load Latency-Packet Loss with Failure Markers.

J. Normalized QoS by Traffic Type

The findings reveal distinct QoS profiles: one class emphasizes higher normalized throughput and bandwidth utilization with a strong signal; another shows higher normalized latency; another exhibits higher packet loss; and another operates at higher network utilization. This analysis uses these contrasts to argue that traffic classes carry different QoS vulnerabilities and resources should be tailored accordingly. For the controller, the radar evidence sets per-class weight vectors and aligns per-class SLA caps with the shapes of these profiles so that actions respect service intents rather than global averages. Fig. 5 shows normalized QoS by traffic type.

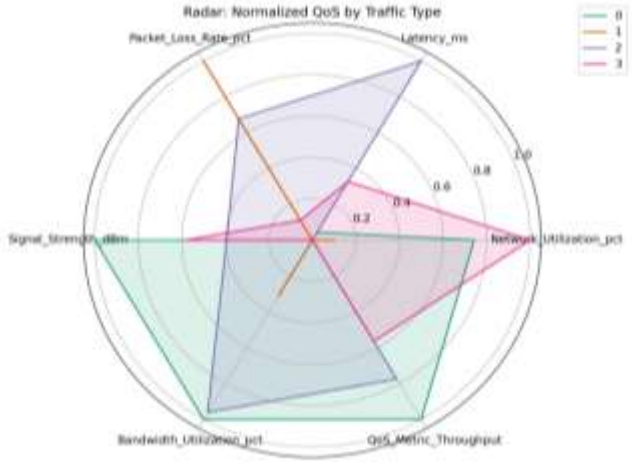

Fig. 6. Normalized QoS by Traffic Type.

K. Network Utilization Over Time

The utilization series across January–April shows rapid temporal variability with frequent peaks and troughs. This analysis reads the structure as evidence of baseline shifts and recurrent bursts (e.g., diurnal patterns), critical for seasonality-aware detection. For MODTC, these baselines feed the twin and the residual detector: the anomaly ensemble flags only departures from expected utilization patterns for a given slice/time-of-day, reducing false positives and allowing the optimizer to act on unexpected surges rather than predictable busy hours. Fig. 6 shows network utilization over time.

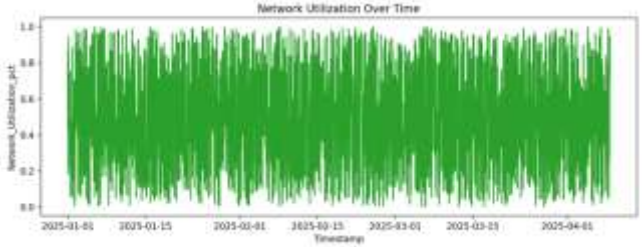

Fig. 7. Network Utilization Over Time.

L. Core QoS Metrics

The findings show diffuse, rectangular scatter clouds with no narrow linear ridges, indicating weak simple relationships between pairs (e.g., Traffic_Load_bps vs Latency_ms). Marginal KDEs on the diagonal reveal multi-modal shapes for some variables (notably Network_Utilization_pct and Packet_Loss_Rate_pct), while QoS_Metric_Throughput concentrates more toward the upper range, suggesting skewed capacity usage. Bandwidth_Utilization_pct exhibits a pronounced mid–high peak, consistent with operational regimes where bandwidth is actively managed. These shapes imply that context and interactions (load + utilization + signal + conditions) drive outcomes more than any single feature. Consequently, threshold or single-regressor policies would be brittle. MO-DTC is better suited here because it ingests the complete multivariate state into a digital twin, reasons about joint patterns revealed by these clouds, and adapts control with context-aware objectives for QoS and network security. Fig. 7 shows core QoS metrics.

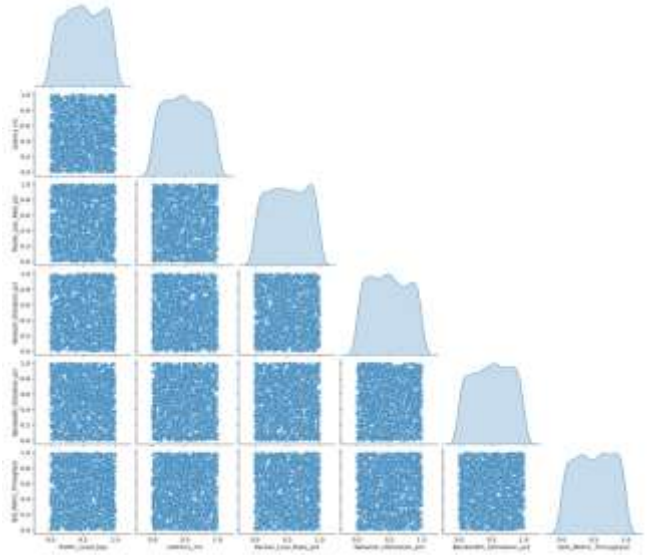

Fig. 8. Core QoS Metrics.

LIMITATIONS

This analysis relies on observational telemetry (2345 rows, 18 features) without ground-truth security incident labels, so the exposure term is calibrated from operational proxies (overload/failure cohorts). The feature set lacks fine-grained mobility and PHY/MAC counters, which would further refine the digital twin and anomaly baselines. Temporal coverage reflects the observed operating window; concept drift or policy/hardware changes may shift learned baselines. Finally, closed-loop optimization and ensemble detection add computational overhead that should be profiled under carrier-grade constraints.

CONCLUSION

This research proposed a novel Multi-Objective Digital Twin-Control (MO-DTC) framework that treats QoS and network security as first-class, co-optimized objectives in 6G network slicing. Using evidence from the dataset's visual analytics and statistical tests, the model fuses a digital twin of slice dynamics, a multi-objective controller over latency–loss–throughput–exposure, a risk-aware real-time policy, and safety shields driven by an anomaly ensemble. The results and discussion show that degradations emerge from multivariate, context-conditioned patterns; encoding those patterns directly in the objective and constraints yields an explainable path to higher SLA satisfaction while reducing overload and failure risk. For future research, this study recommends expanding the telemetry to include mobility and handover states, PHY/MAC counters, and cryptographic lineage to strengthen both the twin and the security term. An A/B "shadow mode" evaluation on live traffic can quantify gains in SLA adherence, alert precision, and time-to-mitigate. Additional directions include federated or privacy-preserving training for cross-domain generalization, causal analysis to separate correlation from intervention effects, integration with integrity mechanisms (e.g., secure attestation) to harden the data plane, and certification-style safety proofs for the shields. Together, these steps will mature MO-DTC from a data-validated design into an operational blueprint for resilient, QoS- and security-aware 6G network slicing.

REFERENCES

[1] Fadlullah, Z. M., Mao, B., & Kato, N. (2022). Balancing QoS and security in the edge: Existing practices, challenges, and 6G opportunities with machine learning. *IEEE Communications Surveys & Tutorials*, *24*(4), 2419-2448.

[2] Louvros, S., Paraskevas, M., & Chrysikos, T. (2023). Qos-aware resource management in 5g and 6g cloud-based architectures with priorities. *Information*, *14*(3), 175.

[3] Sefati, S. S., Haq, A. U., Craciunescu, R., Halunga, S., Mihovska, A., & Fratu, O. (2024). A Comprehensive Survey on Resource management in 6G network based on internet of things. *IEEE access*.

[4] Ming, Z., Yu, H., & Taleb, T. (2024). Federated deep reinforcement learning for prediction-based network slice mobility in 6G mobile networks. *IEEE Transactions on Mobile Computing*, *23*(12), 11937-11953.

[5] Rana, M. K., Pecorella, T., Sardar, B., Thipparaju, R. R., & Saha, D. (2023). A QoS improving downlink scheduling scheme for slicing in 5G radio access network (RAN). *IEEE Transactions on Vehicular Technology*, *73*(3), 4219-4233.

[6] Abdulqadder, I. H., & Zhou, S. (2022). SliceBlock: Context-aware authentication handover and secure network slicing using DAG-blockchain in edge-assisted SDN/NFV-6G environment. *IEEE Internet of Things Journal*, *9*(18), 18079-18097.

[7] Sethi, A., & Bisht, M. (2021). QoS-aware Self-optimized Reconfigurable Framework for Hyperconnected Network. *Global Transitions Proceedings*, *2*(1), 18-23.

[8] Gouda, A., & Sulaiman, L. H. (2025, June). Resource-Aware Network Slicing for QoS-Driven Service Orchestration in Integrated Satellite–5G Systems. In *ECCSUBMIT Conferences* (Vol. 3, No. 2, pp. 70-75).

[9] Tam, P., Ros, S., Song, I., & Kim, S. (2024). QoS-Driven Slicing Management for Vehicular Communications. *Electronics*, *13*(2), 314.

[10] Chen, W. K., Liu, Y. F., Dai, Y. H., & Luo, Z. Q. (2024). QoS-Aware and Routing-Flexible Network Slicing for Service-Oriented Networks. *arXiv preprint arXiv:2409.13943*.

[11] Lin, S. C., Lin, C. H., Subramaniam, S., Matsuura, M., & Hasegawa, H. (2025, May). Multi-Domain Computation-Aware Resource Slicing and Orchestration for 6G Programmable Converged Wireless-Optical Networks. In *NOMS 2025-2025 IEEE Network Operations and Management Symposium* (pp. 1-8). IEEE.

[12] Subhan, F. E., Yaqoob, A., Muntean, C. H., & Muntean, G. M. (2024). A survey on artificial intelligence techniques for improved rich media content delivery in a 5G and beyond network slicing context. *IEEE Communications Surveys & Tutorials*.

[13] Pimpalkar, Y., Ravindran, S., Bapat, J., & Das, D. (2024). A Novel E2E Path Selection Algorithm for Superior QoS and QoE for 6G Services. *IEEE Transactions on Network and Service Management*.

[14] Hussein, D. H., & Ibnkahla, M. (2025). Towards Intelligent Intent-based Network Slicing for IoT Systems: Enabling Technologies, Challenges, and Vision. *IEEE Transactions on Network and Service Management*.

[15] Rafique, W., Barai, J. R., Fapojuwo, A. O., & Krishnamurthy, D. (2024). A survey on beyond 5g network slicing for smart cities applications. *IEEE Communications Surveys & Tutorials*, *27*(1), 595-628.

[16] Kakani, T. A. (2025). Optimization of Serverless Mobile Cloud Applications for Enhanced Security and Resource Efficiency. *Optimization*, *5*(1).

[17] Ziya. (n.d.). *Wireless Network Slicing Dataset* [Data set]. Kaggle. https://www.kaggle.com/datasets/ziya07/wireless-network-slicing-dataset