International Journal of AI, Big Data, Computational and Management Studies
Noble Scholar Research Group | ICAIDSCT26-Conference Proceeding
ISSN: 3050-9416 | https://doi.org/10.63282/3050-9416.ICAIDSCT26-137

AI-Driven Multi-Objective Optimization for Converged Private 5G and Wi-Fi 7 Industrial Networks

Pallavi Priya Patharlagadda
Independent Researcher – AI, Telecommunications & Cloud Systems.

Abstract: Industrial wireless networks are increasingly expected to support workloads that were traditionally confined to wired infrastructure, including real-time robotics, machine vision, and autonomous guided vehicles. The convergence of private 5G and Wi- Fi 7 (IEEE 802.11be) offers a promising path toward this goal, but it also introduces a complex and tightly coupled control space that is difficult to manage using static rules or isolated optimization strategies. This paper presents a practical, AI-driven, multi-objective optimization framework that combines reinforcement learning for radio and slice-level resource control, deep learning for encrypted traffic characterization, and anomaly detection for predictive network assurance. We describe a cloud–edge architecture and a digital twin-based evaluation environment that allow policies to be trained, tested, and deployed under representative industrial workloads. Experimental results show consistent reductions in tail latency and energy consumption, along with improved service-level agreement (SLA) compliance, when compared with conventional rule-based approaches.

Keywords: Wi-Fi 7, Private 5g, Artificial Intelligence, Multi-Objective Optimization, Industrial Networks, Sla Management, Edge Computing.

1. Introduction

The rapid adoption of Industry 4.0 technologies is reshaping the role of wireless connectivity on factory floors, in warehouses, and across large industrial campuses. Applications such as collaborative robotics, autonomous guided vehicles (AGVs), and high-resolution machine vision impose stringent requirements on latency, reliability, and mobility that were traditionally met only by dedicated wired networks. At the same time, operators are under pressure to reduce deployment and operational costs while maintaining predictable service quality across diverse device classes and traffic profiles.

Private 5G and Wi-Fi 7 have emerged as complementary technologies in this context. Wi-Fi 7 introduces multi-link operation (MLO) and ultra-wide channel support in the 6 GHz band, enabling high throughput and low contention for bandwidth-intensive workloads **[1]**. Private 5G, in contrast, provides fine-grained quality-of-service control through network slicing and ultra-reliable low-latency communication (URLLC) modes that are well suited to time-critical control traffic. When these two domains are deployed together, they create a heterogeneous access environment that can flexibly support a wide range of industrial services.

Managing such a converged environment, however, introduces a high-dimensional decision space. Control parameters span radio-level settings, traffic steering policies, slice configurations, and service priorities, all of which interact in non- linear ways under dynamic load and interference conditions. Heuristic or rule-based policies often fail to adapt quickly enough or to account for longer-term performance and cost trade-offs. This motivates the use of artificial intelligence as a unifying control mechanism. In this paper, we formulate cross-domain network management as a multi-objective optimization problem that explicitly balances technical performance metrics, such as latency and energy consumption, with business-oriented objectives expressed through SLAs. By combining cloud-based model training with low-latency edge inference, the proposed framework aims to deliver both global visibility and localized responsiveness in industrial wireless deployments.

This paper makes four primary contributions. First, it presents a converged cloud–edge–access reference architecture that unifies control across IEEE 802.11be (Wi-Fi 7) and private 5G domains. Second, it formalizes cross-domain network management as a weighted, multi-objective optimization problem that jointly considers latency, energy efficiency, SLA compliance, and operational cost. Third, it introduces a reinforcement learning-driven closed-loop control framework that integrates deep traffic characterization and anomaly detection with standardized network control interfaces. Finally, it proposes a reproducible digital twin environment for safe policy validation under representative industrial workloads.

2. Related Work

A substantial body of research has explored the application of machine learning in wireless networks, particularly in the areas of spectrum management, power control, and mobility optimization. Reinforcement learning-based approaches have been widely studied for dynamic channel selection, load balancing, and handover decisions in both cellular and Wi-Fi environments, demonstrating the potential for adaptive control policies to outperform static or rule- based configurations under time-varying traffic and interference conditions.

In parallel, deep learning techniques have been applied to the classification of encrypted traffic flows using statistical, temporal, and flow-level features, enabling application-aware network management without reliance on payload inspection. Such methods are especially relevant in industrial and enterprise environments, where stringent security and privacy requirements limit the feasibility of traditional deep packet inspection and content-based traffic analysis.

From a systems and architectural perspective, industry and standards initiatives, including the ETSI Zero-touch Network and Service Management (ZSM) framework and 3GPP studies on the integration of artificial intelligence and machine learning into 5G-Advanced and beyond, have emphasized closed-loop automation, policy-driven orchestration, and the use of standardized service-based interfaces for cross-layer coordination in autonomous networks **[2], [3]**.

Despite these advances, most existing work has focused on either cellular or Wi-Fi domains in isolation, or on single-objective optimization strategies targeting specific performance metrics. Relatively limited attention has been given to the joint, multi-objective optimization of converged Wi-Fi 7 and private 5G systems in industrial settings, particularly in conjunction with digital twin-based validation environments for safe and reproducible policy evaluation. This paper addresses this gap by presenting an end-to-end framework that integrates architectural design, learning-based cross- domain control, and systematic performance evaluation under representative industrial workloads.

3. Converged System Architecture

The proposed architecture adopts a three-tier cloud–edge–access model that separates long-term learning and global policy management from real-time, localized control. This separation enables the system to scale across large industrial deployments while preserving the low-latency responsiveness required for time-sensitive applications operating under dynamic radio and traffic conditions.

At the cloud layer, centralized services are used to train reinforcement learning models, maintain a digital twin of the network, and manage service-level agreement (SLA) and business policies. Historical telemetry is stored in time-series databases and feature stores, providing a foundation for offline training, model validation, and what-if analysis. Integration with operational and business support systems allows service priorities and cost constraints to be expressed as network-level objectives, ensuring that technical control decisions remain aligned with broader organizational goals.

The edge layer hosts the real-time intelligence of the system and is deployed on lightweight Kubernetes clusters co- located with access gateways. This layer executes inference pipelines for policy evaluation, deep learning-based traffic classification, and anomaly detection. By positioning these functions close to the access network, the architecture reduces control-loop latency and enables rapid adaptation to local conditions, which is critical for supporting industrial workloads such as robotics, autonomous vehicles, and high-resolution sensing.

The access layer comprises heterogeneous radio domains. The Wi-Fi 7 domain includes IEEE 802.11be access points and industrial wireless clients, such as robots, sensors, and vision systems, and exposes metrics including received signal strength indicator (RSSI), signal-to-noise ratio, packet error rate, airtime utilization, and multi-link operation (MLO) link statistics. The private 5G domain consists of base stations and core network functions that support ultra- reliable low-latency communication (URLLC) and enhanced mobile broadband slices for programmable logic controllers and AGVs, providing telemetry such as end-to-end latency, jitter, packet loss, and slice utilization.

Together, the telemetry and control interfaces exposed across these layers define the unified state and action space used by the learning-based, multi-objective optimization framework described in the following section.

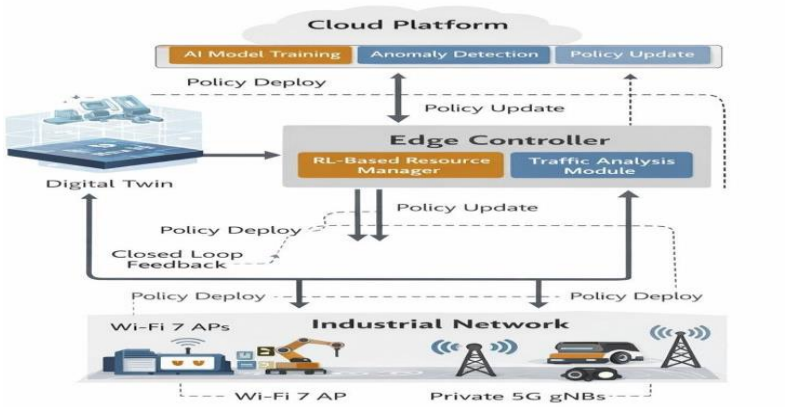

Figure 1: Illustrates the Bidirectional Flow of Policies, Models, and Telemetry across These Layers, Forming a Closed-Loop Control and Learning Framework

4. Multi-Objective Optimization Model

The state of the converged network at time t is represented by a unified feature vector, S(t), which aggregates radio- and traffic-level metrics from both Wi-Fi 7 and private 5G domains. This vector includes, among other parameters, received signal strength indicator (RSSI), signal-to-noise ratio (SNR), packet error rate, queue depth, end-to-end latency, jitter, and slice utilization.

The corresponding action vector, A(t), captures the set of control decisions available to the system. These include multi-link operation (MLO) link selection, channel width and transmit power adjustment in the Wi-Fi domain, as well as slice configuration and traffic steering in the 5G domain.

The optimization objective is expressed as a weighted cost function:

$$(t) = w_1 \cdot (t) + w_2 \cdot (t) + w_3 \cdot (t) + w_4 \cdot (t)$$

Where L(t) denotes end-to-end latency, V(t) represents the rate of service-level agreement (SLA) violations, E(t) captures energy consumption per transmitted bit, and C(t) reflects operational or resource cost. The weighting coefficients are defined by operator policy and service priorities, allowing the system to explore explicit trade-offs between performance, efficiency, and cost. This formulation defines the reward structure for the learning agent and naturally leads to a Pareto frontier of feasible operating points, which provides a structured way to reason about competing objectives in complex industrial deployments.

5. Algorithm Design

The control problem is modeled as a Markov Decision Process (MDP) **[5]**, in which the network environment evolves in response to both control actions and external traffic dynamics. A deep reinforcement learning agent is used to approximate a policy that maps observed network states to control actions in a manner that maximizes expected long- term reward under the multi-objective cost formulation.

At each decision interval, telemetry data is encoded into the current state vector and processed by a hybrid policy network. Actions are selected using an epsilon-greedy strategy that balances exploration of new configurations with exploitation of previously learned policies. The selected actions are enforced through standardized network control interfaces, including NetConf, gNMI, OpenFlow, and 3GPP service-based APIs.

The reward function penalizes high tail latency, SLA violations, and excessive energy consumption, while rewarding stable throughput and robust control behavior. Experience replay and target networks are employed to improve training stability and to reduce sensitivity to transient fluctuations in network conditions, enabling more reliable convergence under non-stationary industrial workloads.

6. Digital Twin and Cross-Domain Control Loop

To support reproducible evaluation and safe policy development, a digital twin of the converged network is constructed. The twin models radio propagation, device mobility, traffic generation, and interference patterns across both access technologies. This environment allows candidate policies to be tested under controlled conditions that closely approximate real industrial workloads before deployment in operational networks.

The operational system follows a closed-loop control process in which telemetry is continuously collected, transformed into feature vectors, and processed by the learning models. The resulting control actions are evaluated against SLA and business policies prior to enforcement in the network. Observed outcomes are then fed back into the learning process as reward signals, enabling continuous refinement of the control policy over time.

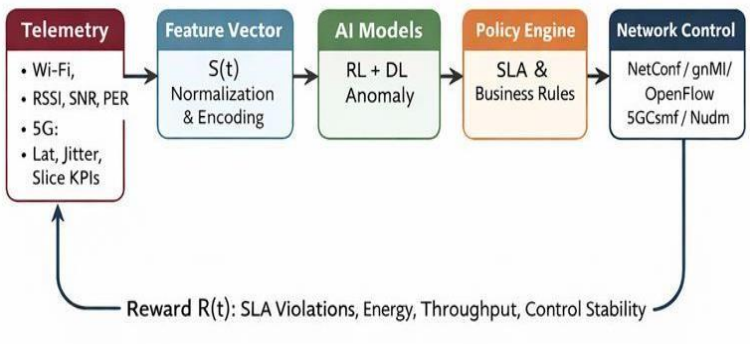

Figure 2: Depicts This Perception–Decision–Action Loop and Highlights the Role of Feedback in Driving Adaptive Optimization

7. Performance Evaluation and Pareto Analysis

The proposed framework is evaluated using a representative industrial workload mix that includes autonomous guided vehicles (AGVs), high-resolution machine vision pipelines, and programmable logic controllers, reflecting a range of latency-sensitive and bandwidth-intensive traffic profiles commonly found in industrial deployments. The learned control policy is compared against two baselines: a conventional rule-based configuration and an SLA-first heuristic that prioritizes service guarantees at the expense of energy efficiency and resource utilization.

Evaluation is conducted across a range of operating conditions, including varying traffic loads, device densities, and interference levels. Performance is assessed using tail latency, SLA compliance rate, and energy consumption per transmitted bit as primary metrics.

Across these scenarios, the AI-driven approach consistently achieves lower tail latency and improved energy efficiency relative to both baselines. The resulting Pareto frontier illustrates that the learned policy maintains high levels of SLA compliance while operating at lower overall cost, indicating more favorable trade-offs between performance, efficiency, and resource consumption in converged industrial network environments.

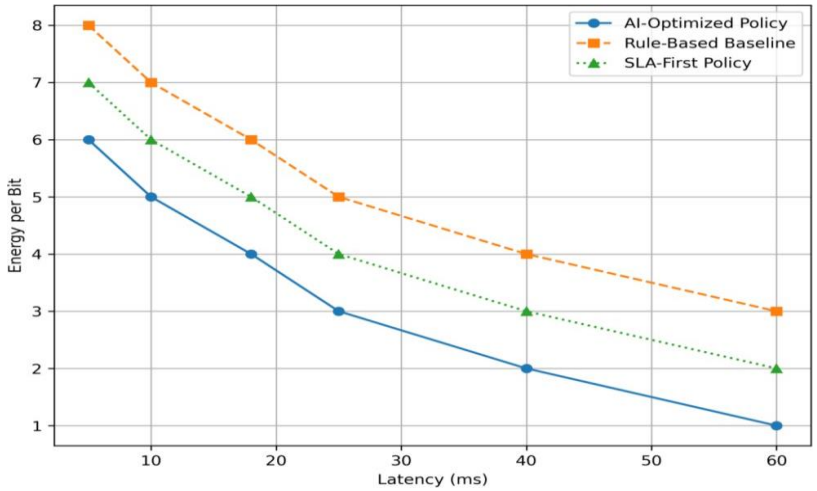

Figure 3: Presents the Latency–Energy Trade-Offs For the Different Policy Regimes and Highlights the Relative Dominance of the Learned Policy over Heuristic and Rule-Based Baselines

8. Security and Compliance

The introduction of an AI-driven control plane raises important security and compliance considerations, including the integrity of telemetry data, the protection of model artifacts, and the authorization of control actions across distributed network elements. The proposed framework addresses these concerns through the use of mutual Transport Layer Security (mTLS) on all telemetry and control channels, role-based access control for policy and model management, and cryptographic validation of deployed models and control policies.

In addition, anomaly detection components integrated into the edge layer are used to identify abnormal traffic patterns and potential attacks in near real time, enabling automated isolation or mitigation of compromised network elements and reducing the operational impact of security incidents.

9. Scalability and Deployment Considerations

All functional components of the framework are implemented as cloud-native microservices that can be independently scaled, upgraded, and orchestrated. Kubernetes-based deployment supports horizontal scaling in response to increases in telemetry volume or inference load, while service mesh technologies provide secure communication, observability, and fault tolerance across distributed deployments.

This design enables the system to be introduced incrementally, beginning with small-scale pilot installations and extending to large industrial campuses and multi-site environments without requiring fundamental changes to the underlying control architecture.

10. Conclusion

This paper has presented a practical, AI-driven framework for managing converged private 5G and Wi-Fi 7 networks in industrial environments. By combining global learning and policy management in the cloud with low-latency decision-making at the

edge, the proposed approach addresses the complexity of cross-domain control while maintaining predictable performance and service-level agreement (SLA) compliance. The integration of a digital twin-based evaluation methodology provides a systematic foundation for safe testing, validation, and future extensions toward more autonomous and self-optimizing industrial connectivity platforms.

References

1. IEEE 802.11beTask Group, Enhancements for Extremely High Throughput, IEEE Standards Association.
2. ETSI ZSM,Zero-touch Network and Service Management Framework.
3. 3GPP TR 23.700-99, Study on AI for Next Generation Networks.
4. ITU-T FG-ML5G, Machine Learning for Future Networks including 5G.
5. R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, MIT Press.