© 2025 IJRTI | Volume 10, Issue 8 August 2025 | ISSN: 2456-3315

ML-Enhanced Control Planes for Adaptive Distributed Systems

Ramkinker Singh
Independent Researcher,
Carnegie Mellon University, USA.

Abstract— The growth of adaptive distributed systems requires smarter and more resilient control planes that adapt to system behaviors and failures. Machine Learning (ML) is becoming an important enabler for advancing these control planes' capabilities with predictive analytics, failure detection, energy efficiency, and efficient data routing. This paper reviews recent research on ML-enabled control planes in adaptive distributed system contexts, including fault tolerance, data replication, energy management, optical networking, and security. The authors systematically reviewed some of the most significant research articles in these areas, and discuss how ML can be embedded within different layers of control plane architecture, and its consequent impact on performance, scalability, reliability, and so on.

The review points out important trends, issues, and technology that impact how control planes are developing and incorporates some aspect of reinforcement learning, active learning, and root cause analysis frameworks. The review also reinforces the ability of ML to make distributed systems resilient and context aware, as well as the important need for resilient, secure ML models. In contrast to other reviews that have pointed out the power of machine learning in disconnect domains, this paper takes an integrated view that correlates fault tolerance, the energy efficiency and security and optical networks with ML-enabled control planes—beginning to fill a significant gap in the literature with respect to adaptive distributed systems and supporting a holistic view. The insights also serve as a lead into direction for future research with respect to self-optimizing and secure machine learning-based control systems.

***Index Terms*— Machine Learning, Adaptive Systems, Distributed Control Planes, Fault Tolerance.**

1. Introduction and Background

Due to an increase in complexity and breadth of scope, adaptive distributed systems require highly developed and advanced control mechanisms to guarantee reliable, efficient, and scalable systems. The mechanisms that comprise the control plane traditionally either relied on rules that were either static, or one of a few pre-defined logics; rules which prescribe what the control plane should tell the other components of a distributed system to do. The emergence of dynamic and heterogeneous environments, such as cloud computing, IoT ecosystems and cyber-physical systems, has revealed the challenges control planes face when faced with dynamic changes to environmental conditions. Machine Learning (ML) can enhance the intelligence and responsiveness of the control plane by learning from data, predicting outcomes from those data, and adapting to changes in the environment.

There are a variety of notable milestones associated with the discussion of ML-based data replica approach to faults that are related to understanding the critical difference in traditional versus M-based replication in terms of good fault-tolerant process. Traditional data replication does provide fault tolerance and resilience options but is likely to be coupled with larger resource use and inefficiencies. ML has more requirements for modeling, having features that allow for dynamic predictions of node failure so that replication can occur with proper procedures to improve reliability and reduce resources. An intelligent system can anticipate a system's actions through analysis of known historical data, and coordinate proactive replication efforts so as to increase reliability with reduced downtime and minimize the chance of data loss [1].

Also, the incorporation of ML has now advanced into intelligent configuring Reconfigurable Intelligent Surfaces (RIS) varying degrees from introductory - entry-level distributed systems only using IoT hardware. In addition, high fidelity ML algorithms, in association with RIS technology, produce the ability to develop and configure intelligently adjusted, optimally managed communication channels and data transport paths resulting in more flexible and performant distributed systems. These types of enhancements facilitate a more adaptive and efficient control plane, which are critical aspects in IoT use cases that entail a plethora of heterogeneous devices and highly dynamic network conditions [2].

ML based learning mechanisms have also entered database systems, which effectively serve as a foundational layer of distributed applications. This may include the use of active learning, which has been leveraged for optimization of query performance and resource allocation by evaluating the most informative data and selectively determining what information to learn, to produce efficient data processing and an agile control plane [3].

To summarize, ML is starting to penetrate every aspect of the management of distributed systems - including replication control, channel reconfiguration, data queries and decision making. This shift from traditional rule-based control planes to a data-driven approach can potentially exponentially increase the robustness and autonomy of distributed systems. However, there will be challenges, such as achieving acceptable performance of ML models with respect to accuracy, secure the systems and integrate different sources of data.

As adaptive systems continue to evolve, the role of ML will play an increasingly important role in the development and realization of intelligent self-organizing control architectures. The remainder of this review paper discusses the specific aspects of ML-enhanced control planes, mainly in energy management, cyber-physical system security, root-cause analysis and optical networking, illustrating how recent research efforts and development demonstrate how ML is changing the fundamental operation of control planes in a timely and capable manner for next generation distributed systems [1] [2] [3].

What differentiates this review from previous work is its cross-sectional coverage reviewing ML-enabled control across many foundational pillars of distributed systems: fault tolerance; data replication; energy management; security in cyber-physical systems; root cause analysis; and optical network optimization. Existing surveys typically consider each domain independently, using limited scope, such as narrow analysis of a specific stack, or document a narrow set of application area technologies or use cases. The paper presents a new contribution by connecting these domains under a single overarching ML-driven control plane framework, revealing some common challenges, shared technologies, and unknown interdependencies. It is a comprehensive and wide-ranging classification of ML use cases that span across layers of control and identifies the range of research gaps that can be realized at the intersections of these domains. This integrated approach may inform future investigations going forward, and help develop distributed systems that are more integrated, intelligent, and resilient.

2. ML-Driven Energy Optimization and Resilience

This integrated approach is likely to inform future research directions, and enhance the development of distributed systems that are more integrated, intelligent, and resilient.

Energy management has emerged as an important top-level concern for distributed systems given the increasing density of edge devices, data centers, and IoT infrastructures that imply significant consumption of energy. The introduction of Machine Learning (ML) into the control planes of distributed systems has created new opportunities for adaptive & lower energy consuming processes. ML-enhanced control mechanisms demonstrate that power consumption can intelligently be managed with real-time monitoring, prediction, and optimization, to achieve energy consumed system-wide and improve resilience.

Modern EMS makes use of a variety of ML models to predict energy consumption behavior, identify waste, and design load-balancing activities. Each of these capabilities rely on the fact that being given limited energy, performance or trust can degrade into unacceptable product or service delivery where stateful executions occur. Time-series prediction models can be leveraged to predict peak demand periods, enabling the control plane to execute a revised workload or initial alternative power savings activities. For example, reinforcement learning has also been successful by enabling resource scheduling decisions to be made by learning optimal control policies in the long run [4].

Resilience in distributed control systems refers not just to recovering from faults but to being able to respond to energy constraints without performance loss. Certainly, ML accommodates both needs by providing the flexibility of real-time decision-making based on an overall appreciation of the system's operational condition. Supervised learning methods can be trained with historical operational data to identify an early indication of stress within the system and take preventative action in the form of mitigation plans, throttling non-essential services, or shifting loads across nodes [4].

A remaining threat challenges the security of these ML-based controls in networked cyber-physical systems (CPS) whereby physical processes are highly dependent upon computing environments. In a networked CPS, compromised controls could have a major delta in harm. That being said, CPS can gain resiliency through ML, particularly in the effective detection of anomalies. With anomaly detection, models can be trained on what is normal operation (using operational data) and flag variations in operational data as potential security threats or outright failures. Even so, ML models in and of themselves represent a new layer of threats, especially when adversarial attacks and data poisoning is not accounted for, allowing a wider range of attack surfaces to the control plane [5].

Security-oriented machine learning techniques are developing as a subset focused on both securing machine learning applications and using machine learning as a means to secure systems. Adversarial training, differential privacy, and robust learning algorithms can be embedded into the control plane architecture, which guarantees the intelligence gained via machine learning to be reliable and indeed secure. These methods not only protect the learning models but also lend resiliency to a system by offering redundancy in maintaining reliable availability and integrity of control tasks [5].

In addition, for high density compute environments like data centers, the management of enormous quantities of data traffic is primarily handled by optical networking. Machine learning has been increasingly used as part of the operations to optimize optical switch configuration, predict network congestion, and dynamically manage wavelength allocations. Many of these functions were developed as a more static switching action but have benefitted extensivley from the additional machine learning aptitude, allowing for energy savings, improvement in throughput, and reduction in latency [6].

The use of ML within energy management and resilience of energy systems will facilitate the transition to self-optimizing systems. As distributed systems become more complex, and energy demand continues to grow, ML powered control planes will emerge as a sophisticated and scalable way to retain efficient operation while maintaining reliability of the energy system. By utilizing predictive models, secure decision based action and dynamic load management, ML will support an environment where distributed systems are efficient, dependable and sustainable [4] [5] [6].

3. Root Cause Analysis and Incident Resolution

Within adaptive distributed systems, failings or performance degradation can result from numerous contexts such as hardware failure, misconfigurations, software bugs, and network issues. Often, these failings can result in a sequence of failures that make traditional troubleshooting problematic to remediate. Machine Learning (ML) has quickly established itself as a powerful method of root cause analysis (RCA) playing a crucial role in automating the control planes ability to detect, troubleshoot and respond with limited human involvement to system failings and anomalies. An ML supplemented RCA method ensures both resolution of incidents is accelerated and the opportunity for the system to avoid repeating faults. By creating a more intelligent and resilient control plane, not only can irrelevant human behavior be mitigated, the extent of incidents degrading operation can be reduced.

When utilizing a ML-enhanced RCA, historical data log events, performance metrics, and event sequences can be analyzed to build models that learn the correlations between root causes and symptoms. The models can find hidden dependencies and causal connections that a rule-based RCA would miss in many cases. For example, decision trees and support vector machines are commonly used for classifying incident causes based on system logs and alerts. After training, the model would return the most likely root causes of a failure, so that the control plane can take corrective actions autonomously [7].

In distributed systems that include numerous components, RCA is not only slow, it is also prone to human error. This is not the case with ML applications, as many issues can be diagnosed with a scalable and data-driven approach. Another major benefit of ML-based RCA is the notion of ongoing learning. The ML will update itself to adjust self when the system encounters new faults or performance bottlenecks, allowing for continuous learning. For adaptive systems operating in changing environments with varied workloads, this ability to learn over time becomes essential [7].

Furthermore, the control plane having RCA capabilities improves the operational efficiency of the entire system. By lowering MTTD (Mean Time to Detect) and MTTR (Mean Time to Repair), the RCA that's driven by ML attempts to minimize the downtime a system would experience in combination with more predictable SLAs (service-level agreements). Another benefit is the reduced dependence on system administrators and domain experts, which increased the consistency of incident management while requiring less resources.

One of the most common ways to implement ML-based RCA is through the correlation of the many log streams from distributed components. As that data of logs is parsed and converted into feature vectors, it becomes easier for the ML model to correlate to an abnormal pattern relative to the known classes of failure. Using unsupervised learning methods such as clustering and anomaly detection give better chances of identifying novel failure patterns that have yet to deviate from nominal usages based on previous incidents. This approach allows the control plane to provide proactive diagnostics, and allows the ability to remediate even on zero-day faults [7].

A prominent example of ML-enhanced RCA is in large-scale cloud infrastructure platforms that have many distributed applications producing a lot of operational telemetry data. The ML models can process this operational data to find root causes of failures, such as network congestion, I/O bottlenecks, memory leaks, etc. This operational insight is then used by the control plane to adjust system parameters or relocate scheduled workloads back to a normal operating state in a fast and efficient manner.

Thus, ML-enhanced RCA will eventually turn a control plane into a real-time diagnosis engine with new capabilities of failure detection and automatic corrections. As we drive toward new system paradigms with increased complexity and decentralization, we will need to incorporate RCA into the control plane with ML to ensure we maintain high levels of reliability, high availability, and agility [7].

4. ML in Optical Networking and Control Plane Optimization

Distributed systems are progressively evolving and creating a demand for higher-speed data movement and lower-latency communication. Due to this demand, optical networking is becoming a crucial infrastructure component for these needs. The integration of Machine Learning (ML) in the control plane of optical networks transformed the design and functionality of optical networks and enabled dynamic configuration, traffic management, and fault recovery. In a world of scaling systems, ML is essential to maintain optimal performance and resource consumption across diverse and dense networking environments.

Optical networks are very complex, consisting of many layers of interconnected photonic switching, Wavelength Division Multiplexing (WDM), and Optical Transport Networks (OTNs). The design thinking of control within these networks is also traditionally static. As a traffic experience becomes dynamic, the control within the operational interface becomes less than ideal acquired performance. The learning aspects of ML in the control plane offer the potential for feedback control based on historical trajectories, learning future demand through predictive analysis, and the ability to give dynamic control of paths and preventing congestion by using past experiences to forecast future demand to optimize recording or available bandwidth [5]. Achieving optimal throughput while minimizing latency [6]; this is one of the learning functionalities ML offers and could offer to optical networking.

A significant use of ML for use in optical networks can be found in photonic switching, where it's essential to possess a certain level of intelligence in order to control how signals are routed through different optical paths. ML algorithms can be utilized to investigate the potential switch configuration based on historical use and which parts of the network are actually being utilized, increasing the overall efficiency and reliability of data transferred over the optical network. Using ML can also identify when optical components are faulty before they degrade or impede the rollout of any service, offering a way to do proactive maintenance upon faulty components, and improving the reliability of the system overall [6].

ML can also support an optical network in assessing the performance of the system through the application of anomaly detection models that can monitor existing telemetry data on observing bit-error rates, "signal-to-noise" ratios, and delays. These models can detect very small deviations from historical observations that may not yet have data faults or performance issues, except to log that there is a need to monitor the existing telemetry data. When an incident is detected by an ML-enriched control plane, there are numerous control actions the system can take including but not limited to automatically rerouting traffic, adjusting the transmission set parameters, or turning on redundant links [8].

The configuration of the ML optical control plane and network is illustrated in the diagram below to visually demonstrate the impacts and integration of ML into an optical control plane.

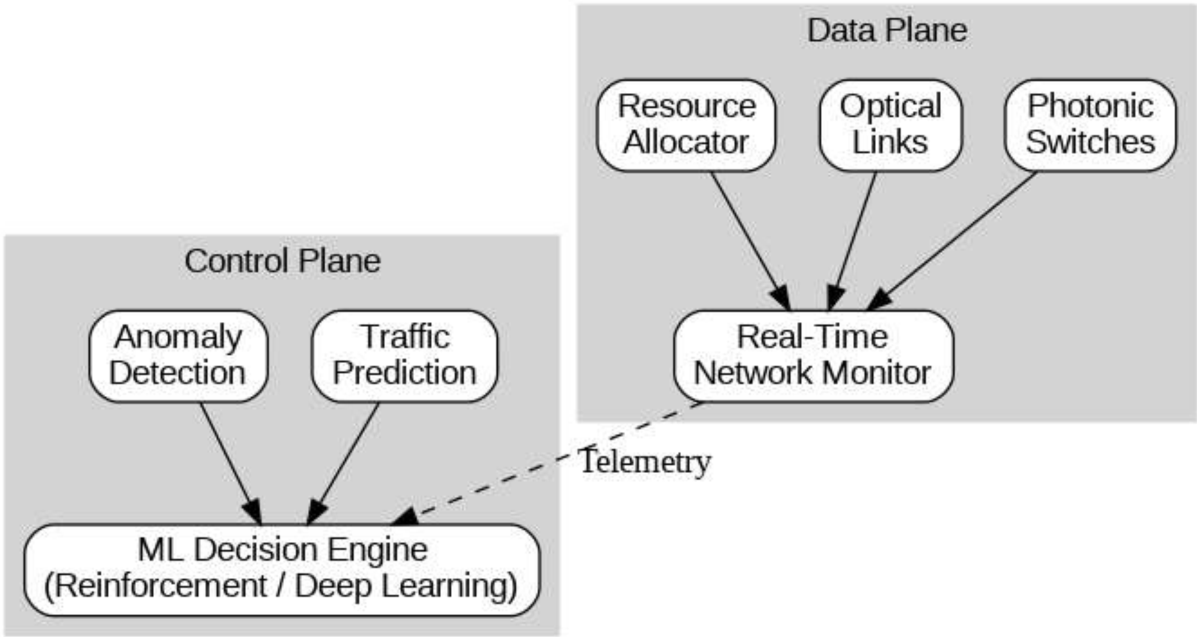

Figure 1: Machine Learning (ML) can be applied to both optical networking layers, i.e., the control and data planes, to support many real-time applications for things like traffic prediction, anomaly detection, and more intelligent routing.
Source: Adapted from [8]

But ML-based optimization has an even greater importance with optical access networks. These networks provide high-bandwidth connectivity to end-users, but must deal with issues stemming from constantly fluctuating traffic conditions, variable user demands, and fault management. In many cases, ML based solutions, such as supervised learning, can help these networks classify user demand profiles and predict when a high-utilization window is coming up, ultimately allowing the control plane to facilitate resource allocations, improve quality of service (QoS), and reduce latency on the data paths.[8]

Different ML models should also have the capability to assist and improve next-generation network design by simplifying tasks related to network planning, including capacity planning, topology optimization, and wavelength assignment. The next table illustrates some possible usages of ML techniques in optical networking.

Table 1: ML strategies and applications in optical networking [6] [8]

ML Technique	**Application Area**	**Benefit**
Supervised Learning	Traffic classification	Predict user demand, optimize routing
Unsupervised Learning	Anomaly detection	Identify network faults early
Reinforcement Learning	Dynamic wavelength assignment	Real-time resource optimization
Deep Learning	Photonic switch configuration	Enhance decision accuracy
Clustering	Fault pattern recognition	Improve fault management

To provide a better understanding of how ML can permit better performance from control planes in optical networking, the following image can be observed to show the difference in bandwidth utilization provided by an ML enabled system, as compared to the more traditional systems:

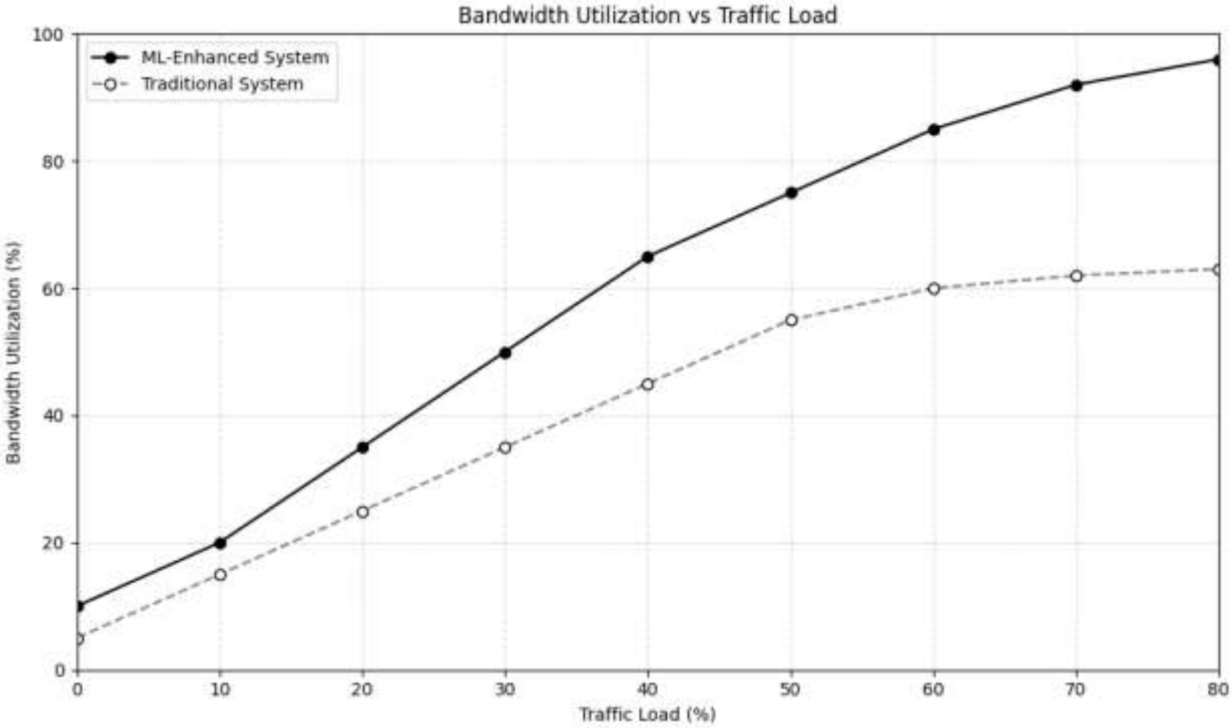

Figure 2: Comparative illustration of bandwidth utilization, again ML enabled systems provided (higher utilization and lower latency) than other systems during varied traffic conditions. Source: Derived from [8]

In conclusion, the inclusion of ML has changed the landscape in which even the control plane perspectives allow the optical networks to have capabilities unlike what has actually been possible prior. Whether capable of intelligence switching control planes or dynamically optimized capabilities or the advanced ability to provide predictive fault detection for modern the distributed applications requiring superior levels of performance and reliability, ML will change how optical networks will align for these demands. Furthermore, as modern future networks continue to scale in performance and complexities like artificial intelligence, machine control, and communications, ML will be a defining element for optimizing how control plane operations would react [6] [8].

5. Conclusion

Machine Learning (ML) has established itself as a critical aspect of the development process of control planes for adaptive distributed systems. The literature shows that ML can improve development processes taken for fault tolerance; data replication; energy management; root cause analysis; and optical network management. Each of those above activities have their own problems on how ML solved these issues utilizing predictive analytics; automated decision-making; and real-time adjustments.

Starting with the controllability of fault tolerance and data replication, ML provides systems the capability of preemptively identifying and addressing risks of system downtime while improving system reliability; allowing more adroit operating procedures. In an even more dynamic operational and communication manner, the same capabilities of adaptability expressed above are experienced in RIS supported IoT systems, where communication and network topologies are reevaluated, established, and performed in real-time thus providing even greater performance value within a multi-dimensional heterogeneous systems performance envelope. The modeling of response times with active learning techniques in database environments provide accuracy without retraining and increased efficiency of resources and timely query performance for control plane operations.

Energy management systems are able to leverage ML both for the forecast of usage trends and then for dynamically modifying state variables to reduce energy consumption while maintaining a given level of performance. In the context of being cyber-physical systems, the use of ML can enhance both resilience and security through anomaly detection and enhances the learning models against adversary-driven attempts. The clarification of root cause issues can be expedited by the ML capability to find causal relationships within the complex connected system leading to faster restoration of service, and improved availability of the system.

In optical networks, ML provides new capabilities such as automatically determining switch configurations, dynamically allocating bandwidth (on demand), and ensuring that service is communicated to the user with varying loads. This capacity enhances throughput, reduces overall latency, and improves fault handling. In summary, ML has begun to reshape our thinking about how modern networks will be designed, configured, and redeployed in the future.

As distributed systems continue to grow both vast and dynamic in nature, the urgency to embed ML into control planes will become a standard practice. Future developments will focus on deploying more capable models, making existing models more interoperable, and establishing a defense-in-depth model to support those practices. The big takeaway is that ML is not merely an upgrade, but a fundamental building block of the next generation of adaptive, intelligent, and resilient distributed systems.

5. References

[1] Barrio, L. (2023). Enhancing Fault Tolerance in Distributed Systems Through Machine Learning-Based Data Replication.

[2] Das, S. K., Benkhelifa, F., Sun, Y., Abumarshoud, H., Abbasi, Q. H., Imran, M. A., & Mohjazi, L. (2023). Comprehensive review on ML-based RIS-enhanced IoT systems: basics, research progress and future challenges. *Computer Networks*, *224*, 109581.

[3] Ma, L., Ding, B., Das, S., & Swaminathan, A. (2020, June). Active learning for ML enhanced database systems. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data* (pp. 175-191).

[4] Lee, S., Seon, J., Hwang, B., Kim, S., Sun, Y., & Kim, J. (2024). Recent trends and issues of energy management systems using machine learning. *Energies*, *17*(3), 624.

[5] Olowononi, F. O., Rawat, D. B., & Liu, C. (2020). Resilient machine learning for networked cyber physical systems: A survey for machine learning security to securing machine learning for CPS. *IEEE Communications Surveys & Tutorials*, *23*(1), 524-552.

[6] Yoo, S. B. (2023). New trends in photonic switching and optical networking architectures for data centers and computing systems. *Journal of Optical Communications and Networking*, *15*(8), C288-C298.

[7] Katragadda, S. R., Tanikonda, A., Peddinti, S. R., & Pandey, B. K. (2021). Machine Learning-Enhanced Root Cause Analysis for Accelerated Incident Resolution in Complex Systems. *Journal of Science & Technology*, *2*(4), 253-275.

[8] Wong, E., Mondal, S., & Ruan, L. (2023). Machine learning enhanced next-generation optical access networks—Challenges and emerging solutions [Invited Tutorial]. *Journal of Optical Communications and Networking*, *15*(2), A49-A62.