

Machine Learning-Based Latency Prediction in Wavelength-Switched Data Center Interconnects

Marciano Arregui

Department of Computer Science, University of Houston, Houston, TX, USA.

marcianoarregui348@uh.edu

Yilmaz Kobelt

Department of Computer Science, Binghamton University, Binghamton, NY, USA.

contactyilmaz@binghamton.edu

Abstract

Wavelength-switched data center interconnects promise substantial gains in energy efficiency and bandwidth scalability by replacing or augmenting packet-switched electrical fabrics with reconfigurable optical paths. However, the operational benefits of such architectures depend heavily on their ability to meet stringent latency requirements under highly dynamic traffic conditions. This paper examines machine learning-based latency prediction as a system-level capability for wavelength-switched data center interconnects. It argues that latency prediction is not merely a performance optimization problem but a socio-technical infrastructure challenge involving architecture, scheduling, governance, robustness, fairness, and sustainability. The analysis begins with the sources of latency in optical circuit and wavelength-switched fabrics, including path reconfiguration delays, clock synchronization, buffering, contention, and workload-induced variability. It then explores how supervised learning, sequence modeling, and representation learning can estimate end-to-end latency from telemetry, traffic traces, and topology-aware features. The paper discusses deployment trade-offs between centralized and distributed inference, model freshness, observability overhead, and control-plane coupling. It further considers robustness under failures and adversarial conditions, fairness across tenants and services, and policy implications for service-level agreements in shared optical infrastructure. Evaluation methodology, operational monitoring, and cross-domain comparisons with electrical packet switching are also examined. The conclusion emphasizes that accurate latency prediction can enable proactive wavelength assignment, admission control, and workload placement, but only if prediction systems are governed as critical infrastructure with transparency, accountability, and sustainability constraints. The paper contributes a research agenda for integrating predictive intelligence into wavelength-switched data center interconnects without sacrificing reliability or equitable service delivery.

Keywords

wavelength-switched data center interconnects; latency prediction; machine learning; optical circuit switching; data center networks; resource scheduling; robustness; sustainability; governance; fairness.

1. Introduction

Data center interconnects have become a foundational layer of contemporary digital infrastructure, shaping the performance, cost, and energy profile of cloud services, artificial intelligence training, distributed storage, and real-time applications [1]. As workloads grow in

scale and synchronization intensity, electrical packet-switched fabrics face increasing pressure in terms of switching capacity, cabling complexity, and power consumption. Wavelength-switched and optical circuit-switched interconnects have therefore attracted sustained attention because they can offer high-bandwidth, low-power paths that are reconfigured at the optical layer rather than through per-packet electrical forwarding [2]. Yet optical reconfiguration introduces new temporal dynamics. A path may be fast once established, but the process of establishing, tuning, and synchronizing that path can introduce latency that is difficult to predict under bursty traffic and heterogeneous workload demands.

Latency prediction is central to making wavelength-switched interconnects operationally viable. In conventional packet-switched networks, congestion control, routing, and load balancing can react to changes at millisecond or sub-millisecond timescales, albeit with well-known limitations [3]. In wavelength-switched systems, the decision to assign a wavelength or establish a circuit may have longer-lasting consequences. If a scheduler overestimates the latency benefit of a reconfiguration, it may disrupt flows that require stable paths. If it underestimates latency, it may fail to use optical capacity when it is most valuable. Machine learning offers a way to model these nonlinear relationships by learning from telemetry, topology, traffic class, and workload context [7]. The challenge is not simply to fit a regression model but to build a prediction capability that is robust, governable, and aligned with the physical and economic realities of optical infrastructure.

This paper examines machine learning-based latency prediction in wavelength-switched data center interconnects from a system-level perspective. It focuses on structural trade-offs rather than model architecture alone. The first concern is architectural: where prediction resides, what data it consumes, and how it interacts with the control plane. The second concern is operational: how prediction supports wavelength assignment, admission control, workload placement, and failure recovery. The third concern is socio-technical: how latency prediction affects fairness among tenants, accountability for service-level agreements, and sustainability goals. The required reference at position six provides an important example of low-latency wavelength-switched intra-data center interconnects enabled by advanced fiber and synchronization techniques [6]. This work is not centered on that reference, but it illustrates why latency must be treated as a first-class design variable rather than a residual performance metric.

2. Architectural Context and Latency Sources

Wavelength-switched data center interconnects can be understood as a hybrid infrastructure in which optical paths provide high-capacity, low-power transport while electrical packet switching continues to handle fine-grained multiplexing and bursty traffic [2]. In such fabrics, latency arises from multiple layers. At the physical layer, propagation delay depends on fiber length and material properties. At the optical layer, wavelength assignment, switching time, amplifier transients, and synchronization mechanisms contribute to path establishment and stability. At the network layer, queueing, contention, and buffering occur when traffic aggregates compete for limited ports or wavelength resources. At the application layer, synchronization barriers, distributed consensus, and storage access patterns translate network latency into job completion time and user-visible responsiveness [10].

The distinction between static and dynamic latency is crucial. Static latency can be measured and documented through topology and hardware characterization. Dynamic latency emerges from traffic variability, failure conditions, thermal effects, and control-plane decisions. Wavelength-switched fabrics are particularly sensitive to dynamic latency because

reconfiguration decisions are often made at coarse timescales. A circuit that is established for a large data transfer may be efficient in throughput terms but harmful for latency-sensitive flows if it displaces a path that was already serving interactive traffic [5]. Prediction models therefore need to estimate not only average latency but also tail latency, because distributed systems often suffer disproportionately from rare but severe delays.

Data center traffic measurements show that workloads are highly heterogeneous, with mice flows, elephant flows, incast patterns, and synchronized bursts coexisting in the same physical infrastructure [10]. This heterogeneity means that a single latency model for all traffic classes is unlikely to be sufficient. A prediction system must condition on service class, tenant priority, flow size, deadline, dependency structure, and historical performance. In optical interconnects, the feasible action space also matters. Some paths may be reconfigured quickly, while others require synchronization or guard times. Some wavelengths may be unavailable due to maintenance or protection requirements. Prediction must therefore be coupled with a representation of what the infrastructure can actually do at a given moment [6].

Congestion control and transport protocols further complicate latency prediction. Protocols such as Data Center TCP and RDMA congestion control influence how traffic responds to delay and loss [11], [13]. If a prediction model treats traffic as an exogenous process, it may misestimate the feedback between latency and sender behavior. For example, a predicted increase in latency may cause a transport protocol to reduce its rate, which in turn changes the measured latency. Prediction systems must either model this feedback loop or operate at a timescale where feedback effects are limited. This is a systems problem as much as a learning problem, because the value of prediction depends on whether the control plane can act on it before conditions change [12].

3. Machine Learning Formulation and Feature Space

Machine learning for latency prediction can be framed as a supervised learning problem in which historical observations are used to estimate future latency for a given flow, path, tenant, or service class. The target variable may be mean latency, percentile latency, deadline violation probability, or a distribution over possible delays. The choice of target shapes the entire system. Predicting a single average value may be useful for capacity planning but insufficient for admission control, because tail latency determines whether interactive services remain viable. Predicting a full distribution or a set of quantiles provides richer information but increases the complexity of model training, validation, and operational interpretation [8].

Feature engineering remains important in networked systems because data is often noisy, sparse, and heterogeneous. Useful features may include topology distance, wavelength availability, switch reconfiguration history, queue occupancy, traffic matrix summaries, flow size distribution, tenant identity, time of day, and recent failure events [9]. In wavelength-switched interconnects, features related to synchronization state and clock stability are particularly relevant because these factors can influence both latency and jitter [6]. Features derived from the control plane, such as pending reconfiguration requests and protection switching state, can also improve prediction accuracy. The challenge is to avoid features that leak future information or encode operational biases that would make the model brittle in deployment.

A range of learning methods can be applied. Random forests and gradient-boosted trees are attractive for tabular telemetry because they handle nonlinear interactions and missing values reasonably well [14]. Sequence models such as recurrent neural networks and long short-term

memory networks are suited to time-series prediction when latency evolves over time and depends on historical traffic patterns [15]. Transformer-based models can capture long-range dependencies and multi-entity relationships, but they require substantial data and computational resources [16]. Deep learning more generally offers representation learning from raw or semi-raw telemetry, yet its opacity can create governance problems in critical infrastructure [17]. The appropriate choice depends on the operational regime, the availability of labeled data, and the tolerance for inference latency.

A key system-level issue is the distinction between offline training and online adaptation. Offline models can be trained on large historical datasets and validated against known scenarios. Online models can adapt to changing workload mixes, new tenants, and evolving hardware conditions. However, online adaptation introduces risks of instability, drift, and feedback loops. If the model learns from data generated by its own recommendations, it may reinforce suboptimal behavior. This is a classic problem in machine learning systems and has analogues in safety-critical artificial intelligence [18]. A wavelength-switched interconnect should therefore treat prediction as part of a controlled feedback system with fallback policies, uncertainty estimates, and human or automated oversight.

4. Prediction Architectures and Deployment Trade-offs

Where latency prediction runs is an architectural decision with significant consequences. Centralized prediction can aggregate telemetry across the entire fabric, enabling global optimization and consistent model updates. However, it creates a single point of failure, increases control-plane traffic, and may struggle with the tight latency budgets of intra-data center environments. Distributed prediction, by contrast, can run near switches or servers, reducing communication overhead and improving response time. It may also preserve locality and privacy by keeping tenant-specific data closer to its source. The trade-off is between global visibility and local responsiveness, and hybrid architectures may be necessary for large fabrics [1].

Model freshness is another critical dimension. A model trained on last month's traffic may fail during a new workload deployment or a topology change. Frequent retraining can improve accuracy but consumes computational and energy resources, which is a sustainability concern in large data centers. Incremental learning and transfer learning can reduce the cost of adaptation, but they require careful validation to prevent performance regression. In wavelength-switched interconnects, hardware changes such as new fiber types, amplifier settings, or synchronization regimes can shift the latency distribution. Prediction systems must therefore be coupled with change management processes that trigger revalidation and, when necessary, model rollback [7].

Observability overhead is often underestimated. To predict latency accurately, the system needs measurements from multiple layers: optical power levels, switch states, queue depths, transport metrics, and application-level traces. Collecting all of this data at high frequency can consume bandwidth, storage, and energy. There is a trade-off between measurement granularity and operational cost. Sparse sampling may reduce overhead but miss transient events that drive tail latency. Adaptive telemetry, which increases sampling when anomalies are detected, can be more efficient, but it introduces its own complexity and potential blind spots. The design of observability should be treated as part of the prediction architecture, not as an afterthought [8].

Control-plane integration determines whether predictions are actionable. If latency predictions are delivered as passive dashboards, they may inform operators but cannot improve scheduling automatically. If they are integrated into wavelength assignment, admission control, and workload placement, they can shape system behavior in real time. Yet tight integration increases the consequences of prediction errors. A false low-latency prediction might cause an interactive service to be placed on a path that later becomes congested. A false high-latency prediction might cause the scheduler to reject a request that could have been served. Safe integration requires uncertainty quantification, conservative fallback policies, and clear escalation rules. The system should know when it does not know [18].

5. Robustness, Fairness, Governance, and Sustainability

Robustness is a first-order requirement for latency prediction in optical infrastructure. Failures can occur at the fiber, amplifier, switch, transceiver, or control-plane level. Some failures are gradual, such as degradation in optical signal quality, while others are abrupt, such as a fiber cut or switch misconfiguration. A prediction model trained only on normal conditions may produce dangerously optimistic estimates during failures. Robustness can be improved through anomaly detection, ensemble methods, and explicit modeling of failure states. However, robustness also depends on the surrounding system: protection paths, fast rerouting, and graceful degradation must exist so that prediction errors do not become service outages [5].

Fairness is a subtle but important concern in shared wavelength-switched interconnects. If latency prediction is used to prioritize traffic, it may systematically favor tenants with predictable workloads or high-priority service classes. Tenants with bursty, unusual, or emerging workloads may receive worse predictions and therefore worse scheduling outcomes. Over time, this can create a feedback loop in which disadvantaged tenants receive fewer resources and generate less data, further degrading future predictions. Fairness-aware design may require separate models or calibration for different tenant classes, transparent reporting of prediction uncertainty, and policy constraints that prevent latency prediction from becoming an opaque mechanism of discrimination [18].

Governance and policy are essential because latency prediction affects service-level agreements, competitive dynamics, and trust in infrastructure. Who is accountable when a prediction leads to a deadline violation? How are model changes communicated to tenants? What audit trails are required to investigate incidents? These questions cannot be answered by technical accuracy metrics alone. Governance frameworks should define roles for model owners, operators, tenants, and regulators. They should require documentation of training data, known limitations, and performance across tenant groups. In regulated sectors such as finance and healthcare, latency prediction may influence compliance outcomes, making transparency and reproducibility especially important [8].

Sustainability adds another layer of complexity. Machine learning models consume energy during training and inference, and the data collection required for prediction also has an energy cost. In a data center interconnect, the primary sustainability benefit of wavelength switching is reduced power consumption compared with electrical switching [2]. If prediction systems are overly complex or continuously retrained on high-frequency telemetry, they may erode those gains. Sustainable design favors lightweight models, selective telemetry, and lifecycle assessment that accounts for both operational and embodied energy. It also favors using prediction to improve utilization, because higher utilization can reduce the number of overprovisioned paths and lower overall infrastructure footprint [6].

6. Evaluation, Operations, and Case Illustrations

Evaluating latency prediction requires metrics that reflect operational value, not just statistical fit. Mean absolute error and root mean squared error are useful but may hide poor performance on tail events. For interactive services, the probability of exceeding a deadline may be more important. For bulk transfers, throughput and completion time may dominate. For control-plane decisions, the cost of false positives and false negatives must be considered explicitly. A prediction that is slightly inaccurate but consistently conservative may be preferable to one that is accurate on average but occasionally dangerously optimistic. Evaluation should therefore include scenario-based testing, stress tests, and post-deployment monitoring [12].

Case illustrations can clarify the trade-offs. Consider a machine learning training cluster that periodically synchronizes gradients across nodes. If the interconnect can predict that a wavelength path will have low jitter during a synchronization window, the scheduler can allocate that path to the job. If prediction is uncertain, the scheduler may choose a more stable electrical path or introduce additional buffering. In another case, a latency-sensitive financial application may require predictable tail latency rather than maximum bandwidth. Prediction can help reserve a dedicated wavelength with known characteristics, but it must also account for failover and maintenance. In a multi-tenant cloud, prediction can inform admission control so that new workloads are placed only where their latency requirements can be met without harming existing tenants [10].

Operational deployment requires continuous monitoring and model maintenance. Prediction drift can occur when traffic patterns change, new hardware is installed, or software updates alter transport behavior. Operators need dashboards that show prediction error, uncertainty, and business impact over time. They also need automated alerts when model performance degrades beyond acceptable thresholds. A model that was accurate during a pilot may fail in production if the workload mix or scale changes. Therefore, deployment should be staged, with canary releases, shadow mode evaluation, and rollback plans. The prediction system should be treated as a production service with service-level objectives of its own [9].

Incident response is another operational concern. When latency spikes occur, operators must determine whether the cause is physical, algorithmic, or workload-related. If a prediction model is part of the control loop, it may obscure the root cause unless its inputs and outputs are logged. Explainability tools can help, but they are not a substitute for good telemetry and clear system boundaries. Incident reviews should examine whether prediction errors contributed to the event and whether governance processes were followed. Over time, these reviews can improve both the model and the surrounding infrastructure [18].

7. Cross-Domain Comparisons and Forward-Looking Perspectives

Comparisons with electrical packet-switched data center networks are instructive. In electrical fabrics, latency prediction often focuses on queueing and congestion because the physical topology is relatively stable. In wavelength-switched fabrics, prediction must also account for reconfiguration and synchronization. This makes the problem more similar to dynamic resource allocation in wireless networks or cloud scheduling than to traditional routing. However, optical interconnects also offer opportunities for deterministic latency because circuits can be isolated from competing traffic. The design challenge is to combine determinism where it is available with statistical prediction where it is not [3].

Cross-domain comparisons with other critical infrastructures highlight governance lessons. Power grids use forecasting to balance supply and demand, but they also maintain strict reliability standards and regulatory oversight. Transportation systems use machine learning for traffic prediction, yet they must handle rare events and ensure equitable access. Data center interconnects are not regulated in the same way, but they increasingly underpin essential services. As latency prediction becomes embedded in scheduling and admission control, it may deserve similar attention to transparency, accountability, and fail-safe design. The goal is not to slow innovation but to ensure that predictive intelligence does not undermine the reliability of the infrastructure it is meant to optimize [18].

Forward-looking perspectives suggest several research directions. First, uncertainty-aware prediction should become standard, because point estimates are insufficient for safe control. Second, prediction systems should be co-designed with optical hardware so that telemetry and reconfiguration capabilities align with learning needs [6]. Third, fairness and sustainability metrics should be integrated into model selection and deployment. Fourth, federated or privacy-preserving learning may allow multiple tenants or data centers to improve prediction without sharing sensitive traces. Fifth, policy frameworks should clarify liability and audit requirements for automated latency management. These directions require collaboration among network architects, machine learning researchers, operators, and policymakers [19], [20].

8. Conclusion

Machine learning-based latency prediction in wavelength-switched data center interconnects is a promising but demanding system capability. It can support proactive wavelength assignment, admission control, workload placement, and failure recovery, but only if it is designed with architectural realism and socio-technical responsibility. Accurate prediction depends on understanding the physical and control-plane sources of latency, selecting appropriate learning methods, and integrating prediction into operational workflows with uncertainty quantification and fallback policies. The required reference at position six illustrates the importance of low-latency optical and synchronization advances, but the broader lesson is that prediction must be part of a resilient, fair, and sustainable infrastructure rather than a standalone model. Future work should advance robust models, transparent governance, equitable scheduling, and energy-aware deployment so that wavelength-switched interconnects can deliver their promised performance without introducing new forms of fragility or exclusion.

References

1. Barroso, L. A., Clidaras, J., & Hölzle, U. (2013). *The datacenter as a computer: An introduction to the design of warehouse-scale machines* (2nd ed.). Morgan & Claypool.
2. Farrington, N., Porter, G., Radhakrishnan, S., Bazzaz, H. H., Subramanya, V., Fainman, Y., Papen, G., & Ford, J. (2010). Helios: A hybrid electrical/optical switch architecture for modular data centers. *ACM SIGCOMM Computer Communication Review*, 40(4), 339–350.
3. Al-Fares, M., Loukissas, A., & Vahdat, A. (2008). A scalable, commodity data center network architecture. *ACM SIGCOMM Computer Communication Review*, 38(4), 63–74.

4. Guo, C., Lu, G., Li, D., Wu, H., Zhang, X., Shi, Y., Tian, C., Zhang, Y., & Lu, S. (2009). BCube: A high performance, server-centric network architecture for modular data centers. *ACM SIGCOMM Computer Communication Review*, 39(4), 63–74.
5. Poutievski, L., Mashayekhi, O., Ong, J., Singh, A., Tariq, M., MacCauley, J., ... & Vahdat, A. (2022). Jupiter evolving: Transforming Google's datacenter network via optical circuit switches and software-defined networking. In *Proceedings of the ACM SIGCOMM 2022 Conference* (pp. 66–85). ACM.
6. Zhou, Z., Dzieciol, H., Clark, K., Luo, Y., Richardson, D., Poletti, F., ... & Liu, Z. (2023). Low-latency wavelength-switched clock-synchronized intra-data center interconnects enabled by hollow core fiber. *Optics Express*, 31(15), 24739-24748.
7. Musumeci, F., Rottondi, C., Nag, A., Macaluso, I., Zibar, D., Ruffini, M., & Tornatore, M. (2019). An overview on application of machine learning techniques in optical networks. *IEEE Communications Surveys & Tutorials*, 21(2), 1383–1408.
8. Boutaba, R., Salahuddin, M. A., Limam, N., Ayoubi, S., Shahriar, N., Estrada-Solano, F., & Caicedo, O. M. (2018). A comprehensive survey on machine learning for networking: Evolution, applications and research opportunities. *Journal of Internet Services and Applications*, 9(1), 16.
9. Wang, M., Cui, Y., Wang, X., Xiao, S., & Jiang, J. (2018). Machine learning for networking: Workflow, advances and opportunities. *IEEE Network*, 32(2), 92–99.
10. Kandula, S., Sengupta, S., Greenberg, A., Patel, P., & Chaiken, R. (2009). The nature of data center traffic: Measurements and analysis. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement* (pp. 202–208). ACM.
11. Alizadeh, M., Greenberg, A., Maltz, D. A., Padhye, J., Patel, P., Prabhakar, B., Sengupta, S., & Sridharan, M. (2010). Data center TCP (DCTCP). *ACM SIGCOMM Computer Communication Review*, 40(4), 63–74.
12. Mittal, R., Lam, V. T., Dukkipati, N., Blem, E., Wassel, H., Ghobadi, M., Vahdat, A., Wang, Y., Wetherall, D., & Zats, D. (2015). TIMELY: RTT-based congestion control for the datacenter. *ACM SIGCOMM Computer Communication Review*, 45(4), 537–550.
13. Zhu, Y., Eran, H., Firestone, D., Guo, C., Lipshteyn, M., Liron, Y., Padhye, J., Raindel, S., Yahia, M. H., & Zhang, M. (2015). Congestion control for large-scale RDMA deployments. *ACM SIGCOMM Computer Communication Review*, 45(4), 523–536.
14. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
15. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
16. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems* (Vol. 30, pp. 5998–6008). Curran Associates.
17. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
18. Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*.

19. Zilberman, N., Watts, P. M., & Rotsos, C. (2015). Optical interconnects for data centers. In Optical Fiber Communication Conference (pp. W3I.1). Optical Society of America.
20. Kachris, C., Kanonakis, K., & Tomkos, I. (2013). Optical interconnection networks in data centers: Recent trends and future challenges. *IEEE Communications Magazine*, 51(9), 39–45.