# Intelligent Wavelength Assignment for High-Speed Data Center Optical Networks

Mikaere Healey
Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.
mikaere.work@uab.edu

Shanna Gilsenan
Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
shanna258@buffalo.edu

**Abstract**

High-speed data center optical networks are increasingly expected to support extreme bandwidth, predictable latency, strict tenant isolation, and sustainable energy use under rapidly changing traffic conditions. Wavelength assignment is often treated as a narrow routing and resource allocation problem, yet in modern data centers it functions as a system-level governance mechanism that shapes capacity distribution, failure recovery, energy consumption, and fairness among workloads. This paper examines intelligent wavelength assignment as an architectural and socio-technical capability rather than a purely algorithmic task. It argues that effective wavelength assignment must coordinate optical switching, electronic packet routing, telemetry, control-plane policy, and physical-layer constraints. The discussion considers reconfigurable optical interconnects, hybrid electrical and optical fabrics, learning-based control, low-margin optical design, and hollow-core fiber based low-latency interconnects. It also analyzes trade-offs among capacity, latency, energy, robustness, fairness, and deployment cost. Particular attention is given to governance implications, including tenant service guarantees, carbon-aware scheduling, infrastructure modularity, and cross-domain lessons from telecommunications and high-performance computing. The paper concludes that intelligent wavelength assignment should be designed as a policy-aware, observable, and resilient control layer that balances technical efficiency with operational legitimacy. Such a layer can improve utilization and reduce energy waste, but only if it remains transparent, auditable, and adaptable to heterogeneous data center environments.

**Keywords**

optical data center networks; wavelength assignment; intelligent control; routing and wavelength assignment; optical switching; sustainability; governance; machine learning.

## 1. Introduction

The expansion of cloud computing, artificial intelligence training, storage replication, and real-time analytics has made the data center network a critical determinant of overall system performance. Global traffic forecasts have long indicated sustained growth in both east-west traffic among servers and north-south traffic between data centers and users, placing pressure on electrical packet switching fabrics that were not designed for current port densities and bandwidth demands [1]. Optical interconnection has emerged as a central response because wavelength division multiplexing can provide high-capacity links with lower per-bit energy consumption and longer reach than many electrical alternatives [2]. However, optical capacity

alone does not resolve the broader system problem. The network must decide which connections receive which wavelengths, when those wavelengths are reconfigured, and how those decisions interact with workload placement, failure recovery, and service-level commitments.

Wavelength assignment in data centers therefore cannot be reduced to a static graph-coloring exercise. It is a control problem embedded in an operational environment where traffic matrices shift rapidly, applications have heterogeneous latency requirements, and infrastructure is shared among tenants with different expectations for isolation and fairness. Early optical data center architectures demonstrated the value of hybrid electrical and optical switching, but they also revealed that reconfiguration speed, control complexity, and traffic predictability determine whether optical resources are actually useful [3], [4]. Reconfigurable interconnects such as ProjectToR further showed that optical circuits can improve performance when the control plane can identify and exploit long-lived traffic demands [5]. Intelligent wavelength assignment extends this insight by using telemetry, prediction, and policy constraints to decide how wavelengths should be allocated, migrated, and reclaimed.

## 2. Architectural Context and Requirements of Data Center Optical Networks

Modern data center optical networks occupy multiple levels, including intra-rack, inter-rack, intra-building, campus, and metro-scale interconnections. Each level presents different constraints. Intra-rack and inter-rack fabrics prioritize low latency, high radix, and rapid reconfiguration, while campus and metro links may emphasize reach, spectral efficiency, and cost per bit. Silicon photonics, integrated transceivers, and advanced modulation formats have improved the feasibility of dense optical fabrics, but they also introduce physical-layer impairments and margin requirements that affect wavelength assignment [6], [7]. A wavelength that is acceptable on one path may be unusable on another because of crosstalk, nonlinearity, amplifier noise, or limited regeneration. Intelligent assignment must therefore be aware of both logical topology and physical condition.

Architecturally, wavelength assignment intersects with switching granularity. A network may assign entire wavelengths to circuits, sub-wavelength flows through time or frequency slicing, or wavelength groups to bundles of traffic. The choice affects control overhead, buffering requirements, and failure semantics. Full-circuit switching can provide deterministic performance but may waste capacity when flows are bursty. Fine-grained assignment can improve utilization but demands more frequent signaling and tighter synchronization. Hybrid electrical and optical designs attempt to combine the statistical multiplexing of packet switching with the high capacity of optical circuits [4]. In such designs, wavelength assignment is not an isolated optical function; it is part of a hierarchical control problem that must decide which traffic stays in the electrical domain and which traffic is promoted to an optical path.

The data center as a computer perspective is useful here because it frames the network as part of a larger warehouse-scale machine rather than as an external utility [17]. From this view, wavelength assignment affects the placement of computation and storage, the cost of data movement, and the carbon intensity of operations. If optical paths are assigned without considering workload topology, the network may unintentionally amplify congestion or strand capacity. If assignment is too conservative, expensive optical resources remain underused. The architectural challenge is thus to create a control layer that is expressive enough to capture application intent, fast enough to respond to traffic shifts, and simple enough to be operated reliably at scale.

## 3. Wavelength Assignment as a System-Level Governance Problem

Routing and wavelength assignment has a long history in wavelength-routed optical networks, where the objective is typically to establish lightpaths while respecting wavelength continuity and capacity constraints [8], [9]. Classical formulations assume relatively stable traffic and centralized or distributed path computation. Data centers violate many of these assumptions. Traffic is highly dynamic, hierarchical, and often generated by distributed applications whose communication patterns are not known in advance. Consequently, intelligent wavelength assignment must operate as a governance mechanism that allocates a scarce optical resource among competing demands while preserving stability and accountability [10], [11]. Governance here refers not only to policy enforcement but also to the rules by which priorities are set, conflicts are resolved, and performance is audited.

A governance perspective highlights several structural trade-offs. First, centralized assignment can optimize global utilization but may become a bottleneck and a single point of failure. Distributed assignment can scale and survive controller outages but may produce conflicting decisions and oscillations. Second, reactive assignment can respond quickly to current demand but may thrash under bursty traffic, while predictive assignment can reserve wavelengths for anticipated flows but risks overcommitting capacity. Third, strict isolation guarantees can protect tenants but reduce opportunities for statistical multiplexing. These trade-offs are not merely technical; they determine whether the network can support differentiated services, regulatory expectations, and internal chargeback models.

The economic and operational context of cloud data centers reinforces the need for governance. Data center networks are built to support revenue-generating services, and their costs include capital expenditure, energy, cooling, maintenance, and lost productivity during failures [16]. Wavelength assignment decisions can shift traffic onto paths with different energy profiles, affect the lifetime of optical components, and change the feasibility of maintenance windows. If assignment policies are opaque, operators may struggle to explain performance anomalies or justify prioritization decisions. If they are rigid, they may fail to adapt to new hardware or changing tenant mixes. Intelligent assignment therefore requires explicit objectives, measurable constraints, and mechanisms for review.

## 4. Intelligent Control Architectures and Learning-Based Wavelength Assignment

Intelligent wavelength assignment depends on a control architecture that can observe, decide, and act across optical and electrical domains. Software-defined networking provides a useful foundation because it separates data-plane forwarding from control-plane logic, enabling programmable policies and centralized or hierarchical optimization. Telemetry from optical transceivers, switches, and servers can supply input on utilization, latency, errors, and queue depth. However, the control plane must also contend with incomplete observability, delayed feedback, and noisy measurements. A wavelength may appear available while physical-layer margins are deteriorating, or a path may appear congested while application-level demand is about to shift.

Machine learning and reinforcement learning have been applied to optical network optimization because they can learn policies from experience and adapt to non-stationary conditions [12]. In wavelength assignment, learning-based methods may predict traffic matrices, estimate path quality, or select wavelength configurations that maximize long-term reward. These methods are attractive when explicit models are difficult to construct, but they introduce risks. A learned policy may exploit simulation artifacts, overfit to historical traffic,

or produce decisions that are difficult to interpret. In safety-critical or financially significant environments, such policies must be constrained by hard rules, validated against digital twins, and monitored after deployment. Low-margin optical design further compresses the operating envelope, making accurate prediction and conservative fallback strategies important [13].

A robust intelligent control architecture should combine model-based reasoning with data-driven adaptation. Model-based components can enforce wavelength continuity, avoid known impairments, and respect capacity limits. Data-driven components can detect anomalies, forecast demand, and tune thresholds. Policy engines can encode tenant priorities, maintenance constraints, and carbon budgets. This layered approach allows the network to remain stable even when learning components are uncertain. It also supports human oversight, because operators can inspect the rules that generated a decision. The goal is not to replace engineering judgment but to augment it with faster and more comprehensive situational awareness.

## 5. Trade-offs in Capacity, Latency, Energy, and Robustness

Capacity is the most visible objective of wavelength assignment, but it is not the only one. Assigning wavelengths to maximize throughput can increase latency if traffic is routed through longer optical paths or if frequent reconfiguration disrupts flows. Latency-sensitive applications such as distributed machine learning, financial trading, and real-time control require predictable delay, not merely high average bandwidth. Recent work on low-latency wavelength-switched clock-synchronized intra-data center interconnects enabled by hollow core fiber illustrates how physical-layer advances can reduce latency and improve synchronization, but such benefits depend on careful wavelength coordination and clock distribution [14]. Intelligent assignment must therefore consider time, not only capacity, as a first-class resource.

Energy consumption is another central trade-off. Optical transmission can be more energy-efficient than electrical switching for high-capacity links, but optical components such as lasers, amplifiers, and cooling systems still consume power. Wavelength assignment can influence energy use by consolidating traffic onto fewer active paths, avoiding unnecessary regeneration, and enabling sleep modes for unused resources. However, consolidation may reduce redundancy and increase the impact of a single failure. The design must balance efficiency against resilience. A network that saves energy by overloading a small number of wavelengths may suffer severe degradation when one path fails. Conversely, a highly redundant assignment may waste energy on idle resources. Intelligent control can manage this balance dynamically by considering workload criticality and failure probability.

Robustness also encompasses physical-layer degradation and security. Wavelengths can be affected by aging components, connector contamination, fiber bends, and malicious interference. An intelligent assignment system should detect degradation early, migrate traffic before errors become service-impacting, and preserve forensic evidence for root-cause analysis. It should also consider adversarial scenarios in which an attacker attempts to exhaust wavelengths or induce crosstalk. Resilience is not simply redundancy; it is the ability to maintain acceptable service under uncertainty. Wavelength assignment policies that are aware of failure domains, maintenance schedules, and tenant risk profiles can improve availability without requiring excessive overprovisioning.

## 6. Fairness, Policy, and Operational Governance

Fairness in wavelength assignment is a multidimensional concept. At the infrastructure level, fairness may mean preventing one tenant from monopolizing high-quality wavelengths or low-latency paths. At the application level, fairness may mean ensuring that background storage replication does not starve interactive services. At the organizational level, fairness may mean aligning resource allocation with internal cost attribution and environmental commitments. These interpretations can conflict. A policy that gives priority to revenue-critical workloads may be economically rational but socially or legally contested if it disadvantages public-interest services. Intelligent assignment systems must therefore make value choices explicit rather than hiding them inside optimization weights.

Policy implications extend to transparency and accountability. If an operator uses machine learning to assign wavelengths, affected tenants may need to know whether service-level objectives are being met and whether prioritization follows published rules. Auditing is complicated by the fast timescale of optical reconfiguration and the volume of telemetry data. Nevertheless, logging decisions and their justifications is essential for dispute resolution and continuous improvement. Governance frameworks can require periodic review of assignment policies, independent testing of failure modes, and clear escalation paths when automated decisions cause harm. These measures are particularly important in multi-tenant data centers where tenants may have limited visibility into the shared optical fabric.

Operational governance must also address the human dimension. Network engineers need tools to understand why a wavelength was assigned, what alternatives were considered, and how a policy change will affect traffic. Automation that is opaque can erode trust and lead to workarounds that undermine global optimization. Conversely, tools that expose trade-offs and allow controlled overrides can combine algorithmic efficiency with human judgment. The most effective governance model is likely to be participatory, with input from network operations, application owners, security teams, and sustainability officers. Wavelength assignment then becomes a negotiated process supported by technical mechanisms, not a purely technical decision imposed on stakeholders.

## 7. Deployment Pathways, Sustainability, and Infrastructure Co-Evolution

Deploying intelligent wavelength assignment at scale requires an incremental pathway because data centers cannot be rebuilt overnight. Existing electrical fabrics, legacy transceivers, and heterogeneous hardware must coexist with new optical systems. A practical strategy is to introduce optical islands in high-demand areas, such as accelerator clusters or storage backends, and then expand as control-plane maturity improves. This approach limits risk and allows operators to learn from real traffic. It also requires interoperability standards and open interfaces so that wavelength assignment logic can work across vendors. Without such standards, intelligence remains trapped inside proprietary domains and cannot optimize end-to-end paths.

Sustainability adds another layer of complexity. Data centers already consume significant energy, and the embodied carbon of optical equipment is not negligible [18]. Wavelength assignment can support sustainability by improving utilization, extending equipment lifetime, and enabling carbon-aware routing when energy sources vary over time. For example, traffic could be shifted to paths or regions with lower carbon intensity, provided latency and contractual constraints allow. However, carbon-aware assignment must avoid unintended consequences such as shifting load to older equipment with higher failure rates or increasing network distance. A holistic assessment should consider operational energy, embodied carbon, cooling load, and service quality.

Infrastructure co-evolution is inevitable. As optical switching becomes faster and more granular, wavelength assignment will need to interact with compute scheduling, storage placement, and power management. The boundary between network and application may blur, with applications expressing intent such as required latency, tolerable jitter, and carbon budget. The network then translates that intent into wavelength assignments. This vision is promising but demanding. It requires common semantics, trustworthy telemetry, and governance mechanisms that can handle conflicting intents. The deployment path should therefore prioritize modularity, observability, and reversible decisions, so that early experiments do not lock operators into brittle architectures.

## 8. Case Illustrations and Cross-Domain Comparisons

Telecommunications networks provide a mature reference for wavelength assignment, including centralized path computation, distributed signaling, and protection switching [9], [10]. Data centers differ because traffic is more dynamic and the physical scale is smaller, but the governance lessons remain relevant. Telecom operators have learned that wavelength assignment must be integrated with fault management, service-level agreements, and capacity planning. Data center operators can adapt these practices while recognizing that tenant workloads are more heterogeneous. High-performance computing offers another comparison, where photonic networks have been explored for energy-efficient and bandwidth-reconfigurable interconnects [20]. HPC systems often have known communication patterns, which can make predictive assignment easier, but they also demand extreme reliability and low latency.

Cloud data centers face a different challenge because their workloads are multi-tenant and less predictable. Network architectures such as fat-tree and Clos fabrics were designed for uniform high bandwidth and fault tolerance [15]. Optical extensions must preserve these properties while adding reconfigurability. The emerging optical data center concept anticipates deeper integration of optics, but it also raises questions about control complexity and cost [19]. Cross-domain comparison suggests that there is no single best wavelength assignment strategy. The appropriate design depends on traffic predictability, physical topology, service guarantees, and organizational capacity for automation. A strategy that works for a research HPC cluster may not work for a commercial cloud with strict tenant isolation.

Case illustrations also reveal the importance of failure semantics. In a telecom network, a failed wavelength may trigger protection switching with well-defined restoration times. In a data center, a failed wavelength may affect many small flows, some of which are more critical than others. Intelligent assignment must therefore classify flows and apply differentiated recovery. It must also coordinate with higher-layer protocols that may retransmit or reroute independently. Without coordination, recovery actions can conflict and create cascading congestion. The lesson from cross-domain practice is that wavelength assignment is inseparable from resilience engineering and service management.

## 9. Conclusion

Intelligent wavelength assignment for high-speed data center optical networks is a system-level capability that connects physical-layer optics, network control, workload behavior, and organizational policy. It cannot be reduced to an optimization algorithm because the objectives are plural and sometimes conflicting. Capacity, latency, energy, robustness, fairness, and sustainability must be balanced through explicit governance. Learning-based

control and advanced telemetry can improve responsiveness, but they must be constrained by hard rules, transparent policies, and human oversight. Hollow-core fiber, silicon photonics, and reconfigurable switching expand the technical design space, yet they also raise the stakes for coordination and validation.

The path forward requires architectures that are modular, observable, and policy-aware. Wavelength assignment should be integrated with compute and storage scheduling where beneficial, but it should also preserve clear boundaries so that failures can be contained and responsibilities can be assigned. Operators should adopt incremental deployment strategies, open interfaces, and rigorous auditing. Researchers should continue to study fairness metrics, carbon-aware control, and failure semantics alongside performance optimization. If these dimensions are treated as first-class concerns, intelligent wavelength assignment can help data centers achieve higher utilization and lower energy use without sacrificing reliability or legitimacy. If they are ignored, optical capacity may increase while operational trust and sustainability outcomes stagnate.

**References**


1. Cisco. (2020). Cisco Annual Internet Report (2018–2023) white paper. Cisco Systems.

2. Kachris, C., & Tomkos, I. (2012). A survey on optical interconnects for data centers. IEEE Communications Surveys & Tutorials, 14(4), 1021–1036.

3. 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.

4. Farrington, N., Porter, G., Radhakrishnan, S., Bazzaz, H. H., Subramanya, V., Fainman, Y., Papen, G., & Vahdat, A. (2010). Helios: A hybrid electrical/optical switch architecture for modular data centers. ACM SIGCOMM Computer Communication Review, 40(4), 339–350.

5. Ghobadi, M., Mahajan, R., Phanishayee, A., Devanur, N. R., Kulkarni, J., Ranade, G., Blanche, P.-A., Rastegarfar, H., Kilper, D., & Bergman, K. (2016). ProjectToR: Agile reconfigurable data center interconnect. In Proceedings of the 2016 ACM SIGCOMM Conference (pp. 216–229). ACM.

6. Rumley, S., Nikolova, D., Hendry, R., Li, Q., Bergman, K., & Kilper, D. (2017). Silicon photonics for exascale systems. Journal of Lightwave Technology, 35(3), 429–438.

7. Cheng, Q., Bahadori, M., Glick, M., Rumley, S., & Bergman, K. (2018). Recent advances in optical technologies for data centers: A review. Optical Fiber Technology, 45, 329–341.

8. Chlamtac, I., Ganz, A., & Karmi, G. (1992). Lightpath communications: An approach to high bandwidth optical WANs. IEEE Transactions on Communications, 40(7), 1171–1182.

9. Zang, H., Jue, J. P., & Mukherjee, B. (2000). A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks. Optical Networks Magazine, 1(1), 47–60.

10. Ramaswami, R., & Sivarajan, K. N. (1995). Routing and wavelength assignment in all-optical networks. IEEE/ACM Transactions on Networking, 3(5), 489–500.

11. Banerjee, D., & Mukherjee, B. (1996). A practical approach for routing and wavelength assignment in large wavelength-routed optical networks. IEEE Journal on Selected Areas in Communications, 14(5), 903–908.

12. Suárez-Varela, J., Mestres, A., Yu, J., Kuang, L., Feng, S., Barlet-Ros, P., & Cabellos-Aparicio, A. (2019). Routing in optical transport networks with deep reinforcement learning. Journal of Optical Communications and Networking, 11(11), 547–558.

13. Pointurier, Y. (2017). Design of low-margin optical networks. Journal of Optical Communications and Networking, 9(1), A9–A17.

14. 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.

15. Al-Fares, M., Loukissas, A., & Vahdat, A. (2008). A scalable, commodity data center network architecture. ACM SIGCOMM Computer Communication Review, 38(4), 63–74.

16. Greenberg, A., Hamilton, J., Maltz, D. A., & Patel, P. (2008). The cost of a cloud: Research problems in data center networks. ACM SIGCOMM Computer Communication Review, 39(1), 68–73.

17. 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.

18. Soref, R. (2006). The past, present, and future of silicon photonics. IEEE Journal of Selected Topics in Quantum Electronics, 12(6), 1678–1687.

19. Vahdat, A., Liu, H., Zhao, X., & Johnson, C. (2011). The emerging optical data center. In Optical Fiber Communication Conference (p. OTuH2). Optical Society of America.

20. Kodi, A. K., & Louri, A. (2009). Energy-efficient and bandwidth-reconfigurable photonic networks for high-performance computing systems. IEEE Journal of Selected Topics in Quantum Electronics, 15(5), 1434–1443.