

Joint Optimization of Bandwidth and Latency in Software-Defined Optical Networks

Irmtraud Hanel

Department of Computer Science, George Mason University, Fairfax, VA, USA.
ihanel@gmu.edu

Mairead MacInerney

Department of Computer Science, University of Houston, Houston, TX, USA.
contactmairead@uh.edu

Tonicha R. Kelleher

Department of Computer Science, University of North Texas, Denton, TX, USA.
tonicha1974@unt.edu

Abstract

Software-defined optical networks are increasingly positioned as the programmable substrate for cloud computing, artificial intelligence, distributed data analytics, and latency-sensitive applications. Their central promise is not only flexible bandwidth allocation but also coordinated control of latency, resilience, and cost across optical and packet layers. This paper examines the joint optimization of bandwidth and latency in software-defined optical networks from a system-level perspective. It argues that bandwidth and latency cannot be treated as independent variables because capacity provisioning, routing, spectrum assignment, regeneration, buffering, and control-plane signaling create coupled trade-offs. The analysis synthesizes architectural foundations, elastic optical networking, traffic engineering, telemetry, orchestration, energy management, resilience, fairness, and governance. It highlights that joint optimization requires multi-objective control loops that incorporate physical-layer constraints, service-level objectives, policy requirements, and sustainability metrics. The paper also discusses deployment barriers in brownfield infrastructures, the risk of reinforcing market concentration, and the need for transparent governance of programmable optical resources. Case illustrations from cloud wide-area networks, data centers, and research networks show that different domains prioritize bandwidth and latency differently, complicating universal optimization models. The conclusion calls for interdisciplinary benchmarks, policy-aware orchestration, and robust mechanisms that protect fairness and resilience while enabling efficient use of optical infrastructure. Future work should integrate digital twins, open telemetry, and multi-stakeholder governance to ensure that software-defined optical network optimization serves broad social and economic goals rather than narrow efficiency alone.

Keywords

software-defined optical networks; bandwidth optimization; latency; network control; optical infrastructure; governance; resilience; sustainability; fairness.

1. Introduction

The evolution of optical transport has moved from static, conservatively engineered capacity toward programmable infrastructures that can adapt to traffic, service, and policy conditions.

Early software-defined networking research established the separation of control and data planes as a means of enabling innovation and centralized policy enforcement [1], and subsequent surveys articulated the architectural breadth of programmable networking across domains [2]. In optical networks, this logic extends to transceivers, reconfigurable optical add-drop multiplexers, flexible-grid spectrum assignment, and multi-layer orchestration. The resulting software-defined optical network is not merely an optical system with a controller attached; it is a socio-technical infrastructure in which bandwidth allocation, latency guarantees, energy use, and governance arrangements interact.

The joint optimization of bandwidth and latency is difficult because these objectives are coupled. Additional bandwidth can reduce queuing delay and permit rerouting around congestion, but it may also increase energy consumption, spectrum fragmentation, and capital expenditure. Conversely, aggressive latency minimization can require dedicated paths, frequent regeneration, or overprovisioned capacity that reduces statistical multiplexing gains. Elastic optical networking introduced finer-grained spectrum allocation and distance-adaptive transmission, expanding the design space for such trade-offs [4], [5]. Flexible optical networking further emphasized the need to coordinate physical-layer adaptability with control and management functions [6]. At the same time, recent physical-layer advances demonstrate that guaranteed bandwidth and low latency can be pursued together through frequency-referenced multiplexing, indicating that the optical substrate itself can support joint objectives when control and data planes are co-designed [7].

This paper argues that joint optimization should be understood as a system-level governance and architecture problem rather than a narrow routing or spectrum-assignment problem. The following sections examine architectural foundations, the nature of the bandwidth-latency trade-off, control-plane and telemetry requirements, deployment and sustainability pressures, resilience and fairness concerns, governance and policy implications, case illustrations, and future research directions. The central claim is that technically efficient optimization can become socially or economically fragile if it ignores infrastructure ownership, service equity, environmental externalities, and the institutional arrangements that shape access to programmable optical capacity.

2. Architectural Foundations of Software-Defined Optical Networks

Software-defined optical networks build on the separation of forwarding and control functions, but they also require abstraction across optical and packet layers. The control plane may include one or more software-defined networking controllers, path computation elements, orchestration systems, and domain-specific management platforms. Standards such as the Open Networking Foundation architecture describe layered abstractions for applications, controllers, and network devices [19], while traffic-engineering frameworks such as ACTN provide models for coordinating multi-domain and multi-layer transport resources [20]. These abstractions are essential because optical resources are not uniform: spectrum, modulation format, reach, regeneration, and port availability differ across vendors and domains.

The architectural promise of software-defined optical networks lies in their ability to expose programmable primitives that can be composed into services. A cloud provider might request a high-bandwidth virtual link with bounded latency between two data centers; a research network might request deterministic capacity for scientific instruments; a mobile operator might request dynamic backhaul for a temporary event. Each request implies different weights on bandwidth, latency, resilience, and cost. The controller must translate these intents into optical paths, spectrum assignments, and protection or restoration policies. This

translation is where joint optimization becomes consequential, because the selected configuration determines not only immediate performance but also future flexibility, failure exposure, and energy profile.

Architectural choices also create structural trade-offs. Centralized control can improve global optimization and policy consistency, but it introduces scalability, availability, and trust concerns. Distributed or hierarchical control can improve resilience and domain autonomy, but it may produce suboptimal end-to-end outcomes when domains optimize locally. Similarly, a tightly integrated optical-packet control plane can coordinate latency and bandwidth more precisely, yet it may reduce vendor diversity and increase operational complexity. The architecture therefore shapes what kinds of optimization are possible, which stakeholders can participate, and how accountability is assigned when service levels fail.

3. The Joint Bandwidth-Latency Optimization Problem

Bandwidth and latency are often treated as separate service attributes, but in optical networks they are linked by physics, economics, and control dynamics. Propagation delay depends on distance and the speed of light in fiber, while transmission delay depends on serialization and link rate. Queuing delay depends on traffic bursts and buffer management. Reconfiguration delay depends on control-plane signaling, hardware switching, and synchronization. Joint optimization therefore must consider the entire service path, including access, metro, core, and data-center segments. A path that offers high bandwidth in the core may still violate latency objectives if it traverses congested or circuitous metro segments.

Elastic optical networking expands the feasible set of bandwidth allocations by dividing spectrum into finer slots and adapting transmission parameters to reach and demand [4], [5]. Flexible optical networking further introduces reconfigurable elements that can respond to changing traffic and physical conditions [6]. These capabilities make it possible to imagine controllers that jointly choose route, spectrum, modulation, regeneration, and protection. However, the search space is large, and the objectives are not naturally commensurable. A slight reduction in latency may require a substantial increase in spectrum or energy, while a modest bandwidth increase may be achieved with negligible latency impact in one topology but severe impact in another.

Physical-layer advances can shift the trade-off frontier. Frequency-referenced multiplexing, for example, has been presented as a way to provide guaranteed bandwidth and low latency simultaneously, suggesting that tighter coordination between optical frequency references and network control can reduce the need for conservative overprovisioning [7]. Such advances do not eliminate trade-offs, but they change where the frontier lies. They also raise integration questions: controllers must understand new transceiver capabilities, telemetry must capture physical-layer state, and service-level agreements must be expressed in terms that can be mapped onto optical primitives.

The optimization problem is also temporal. Traffic matrices vary by time of day, application phase, and failure events. Latency requirements may be strict for interactive services and looser for batch transfers. A controller that optimizes for instantaneous throughput may create instability if it reconfigures paths too frequently. A controller that optimizes for stability may leave capacity unused or latency unnecessarily high. Joint optimization therefore requires forecasting, hysteresis, and policy constraints that prevent oscillation and preserve operational predictability.

4. Control Plane Design, Telemetry, and Orchestration

Control-plane design determines whether joint optimization remains theoretical or becomes operational. Software-defined optical controllers must maintain an abstract topology that reflects physical connectivity, spectrum availability, impairment constraints, and administrative boundaries. They must also interact with packet-layer controllers, cloud orchestrators, and service management systems. Standards such as ACTN provide a conceptual framework for abstraction and control of traffic-engineered networks across domains [20], while general software-defined networking architectures emphasize programmability and layered control [19]. These frameworks are necessary but not sufficient, because joint bandwidth-latency optimization requires fine-grained telemetry and rapid feedback.

Telemetry is the sensory system of the optimization loop. It must include link utilization, queue depth, packet delay, optical signal quality, amplifier conditions, spectrum occupancy, and control-plane latency. Without such data, controllers rely on static models that cannot capture congestion, aging infrastructure, or transient impairments. With abundant telemetry, controllers face data-volume, privacy, and correlation challenges. Telemetry streams may reveal traffic patterns that are commercially sensitive, and cross-domain telemetry may require trust agreements. The design of telemetry therefore implicates governance as much as engineering.

Orchestration connects optimization to service delivery. An orchestrator receives intents from applications or business systems, translates them into resource requests, and coordinates controllers across layers and domains. In cloud wide-area networks, software-driven traffic engineering has demonstrated the value of centralized optimization for improving utilization while respecting service constraints [9], [10]. Those experiences suggest that joint optimization in optical networks will require similar closed-loop control, but with additional physical-layer constraints. The orchestrator must decide when to reconfigure, which domains to involve, and how to handle conflicts between latency-sensitive and bandwidth-intensive services.

Artificial intelligence and machine learning can improve prediction, anomaly detection, and resource allocation, but they also introduce opacity and fragility. A model trained on historical traffic may fail during rare events or adversarial conditions. A model optimized for average latency may disadvantage tail flows. Control-plane design must therefore combine learning with formal constraints, fallback policies, and human oversight. The goal is not full autonomy but accountable adaptivity.

5. Deployment, Infrastructure, and Sustainability

Deployment realities shape the feasibility of joint optimization. Many optical networks are brownfield infrastructures with legacy equipment, heterogeneous vendors, and long depreciation cycles. Data-center networks have evolved through architectures designed for scalability and cost efficiency [11], while cloud providers have confronted fundamental cost and topology problems in data-center networking [12]. Measurements of data-center traffic have shown diverse patterns, including bursts and incast-like behavior, that affect latency and bandwidth management [13], [14]. These findings imply that optimization models must account for real traffic distributions rather than assuming smooth, predictable flows.

Sustainability adds another dimension. Optical transport is generally more energy efficient per bit than many electronic alternatives, but total energy consumption still matters because traffic grows rapidly. Research on minimizing internet-service-provider network energy cost

has shown that routing, capacity adaptation, and sleep modes can reduce consumption under low demand [15]. Joint bandwidth-latency optimization can support sustainability by consolidating traffic onto fewer links, yet latency constraints may prevent aggregation. Similarly, dedicated low-latency paths may increase energy use if they require overprovisioning or redundant equipment. A sustainable controller must therefore treat energy as a first-class objective alongside bandwidth and latency.

Deployment also involves physical constraints such as conduit availability, rights-of-way, spectrum licensing, and data-center siting. These factors are not purely technical, and they influence which optimization outcomes are achievable. A controller may identify an optimal path in graph terms, but the path may be unavailable because of municipal permits, environmental restrictions, or ownership boundaries. Infrastructure governance therefore directly affects technical performance. Sustainability metrics should include embodied carbon, equipment lifetime, and e-waste, not only operational energy.

6. Robustness, Resilience, and Fairness

Robustness is essential because optical networks face fiber cuts, amplifier failures, power outages, and cyberattacks. Resilience research has developed principles for survivability, including diversity, redundancy, and graceful degradation [16]. Joint bandwidth-latency optimization must incorporate these principles. A path that offers minimal latency under normal conditions may be highly vulnerable if it shares a conduit with other critical links. A bandwidth allocation that maximizes utilization may leave insufficient headroom for restoration after a failure. The controller must therefore optimize not only for expected performance but also for failure scenarios.

Resilience and latency can conflict. Fast rerouting may require precomputed backup paths that consume spectrum, while restoration after failure may introduce delay. Protection schemes can guarantee recovery but increase cost. The controller must decide which services require deterministic recovery and which can tolerate degradation. These decisions are inherently political and economic, not merely technical, because they determine who receives reliable service during disruptions. A hospital network, financial exchange, and video streaming service may all demand low latency, but their tolerance for failure differs.

Fairness is another under-examined dimension. If programmable optical capacity is optimized primarily for large cloud providers or latency-sensitive financial applications, other users may experience degraded service. Rural and underserved regions may lack the infrastructure to benefit from advanced optimization. Network neutrality debates have long recognized that infrastructure owners can shape traffic and market outcomes [17], and infrastructure scholarship emphasizes the social value of shared resources [18]. In software-defined optical networks, fairness can be supported through policy constraints, reserved capacity, transparent service classes, and inclusive governance. Without such measures, joint optimization may increase efficiency while widening digital inequality.

7. Governance, Policy, and Market Implications

Governance determines how optimization objectives are set and who has standing to contest them. Technical standards bodies, open-source communities, regulators, and operators all influence the design of software-defined optical networks. The tussle framework in internet governance shows that technical architectures often embody competing interests rather than a single rational design [17]. In optical networks, conflicts may arise over spectrum access, interconnection terms, telemetry sharing, and control-plane jurisdiction. A domain may be

unwilling to expose internal topology or real-time state to an external orchestrator, even if such exposure would improve end-to-end optimization.

Policy implications extend to competition and universal service. If advanced optimization requires proprietary telemetry, artificial intelligence models, or integrated hardware-software platforms, large incumbents may gain advantages. Small operators may struggle to participate, reducing market diversity and resilience. Regulators may need to mandate interoperability, open interfaces, or nondiscriminatory access to essential optical infrastructure. The social value of infrastructure suggests that access rules should be evaluated not only by efficiency but also by distributional effects [18].

Security and privacy are also governance concerns. A programmable optical control plane is an attractive target. An attacker who compromises the controller could redirect traffic, degrade latency, or disable capacity. Telemetry sharing can leak sensitive information about traffic patterns. Governance frameworks must therefore require authentication, authorization, auditing, and incident reporting. They must also clarify liability when optimization algorithms cause service failures. These issues are not peripheral; they determine whether stakeholders trust software-defined optical networks enough to delegate critical services to them.

8. Case Illustrations and Cross-Domain Comparisons

Cloud wide-area networks provide a useful illustration. Software-driven traffic engineering has been used to improve utilization across large backbone networks while maintaining service requirements [9], [10]. These systems demonstrate the value of centralized optimization, but they also show the importance of operational constraints, failure handling, and incremental deployment. Extending such approaches to optical layers would require joint control of spectrum, routing, and packet scheduling. The benefits could be substantial, but the complexity and governance requirements would also increase.

Data-center networks offer another case. Scalable architectures such as fat-tree designs address bandwidth and fault tolerance through topology [11], while cloud cost analyses highlight the economic pressures that shape network design [12]. Measurements of data-center traffic reveal that latency and bandwidth demands are highly variable [13], [14]. A software-defined optical network could provide reconfigurable capacity between racks or pods, but it must integrate with existing packet fabrics. The optimization problem is therefore not only optical but also organizational, because data-center operators may prioritize predictability and simplicity over maximal efficiency.

Cross-domain comparisons can illuminate institutional dimensions. Electric power grids use market mechanisms, reliability standards, and public regulation to balance efficiency, resilience, and equity. Transportation systems use pricing, zoning, and public investment to manage congestion and access. Optical networks lack a single global regulator, but they face similar tensions between private optimization and public interest. Research and education networks often provide a counterexample by prioritizing open access, experimentation, and community governance. These differences suggest that joint bandwidth-latency optimization should be designed with domain-specific governance models rather than assuming a universal market framework.

9. Future Research Directions

Future research should develop multi-objective benchmarks that evaluate bandwidth, latency, energy, resilience, and fairness together. Such benchmarks would require realistic topologies,

traffic traces, failure models, and policy constraints. They should be open enough to permit reproducibility while protecting sensitive operational data. Digital twins of optical networks could support safe experimentation by simulating control decisions before deployment. Open telemetry standards could reduce integration barriers, but they must include privacy and security safeguards.

Another direction is policy-aware orchestration. Controllers should be able to encode regulatory constraints, service-level equity requirements, and environmental limits as explicit objectives or constraints. This requires collaboration between engineers, economists, legal scholars, and policymakers. Artificial intelligence can assist with prediction and optimization, but it must be explainable and auditable. Human operators should retain authority over decisions that affect public safety, competition, or civil rights.

Finally, research should examine the life-cycle sustainability of programmable optical infrastructure. Energy savings from dynamic operation may be offset by the embodied carbon of additional transceivers, controllers, and cooling systems. Fairness and resilience should be evaluated over time, not only at deployment. Longitudinal studies of software-defined optical networks could reveal whether optimization concentrates benefits or distributes them broadly. Such research would help ensure that joint optimization contributes to sustainable, inclusive, and trustworthy digital infrastructures.

10. Conclusion

Joint optimization of bandwidth and latency in software-defined optical networks is a systems problem that spans physical-layer capabilities, control-plane design, orchestration, energy use, resilience, fairness, and governance. Bandwidth and latency are coupled through propagation, queuing, reconfiguration, and economic constraints, so treating them as independent service attributes leads to incomplete or fragile designs. Elastic and flexible optical networking expand the feasible design space, and advances such as frequency-referenced multiplexing show that physical-layer innovation can shift the trade-off frontier [7]. Yet technical capability alone cannot resolve competing objectives.

The paper has argued that joint optimization should be embedded in accountable architectures that combine telemetry, forecasting, policy constraints, and human oversight. Deployment in brownfield environments, sustainability pressures, and market concentration require regulators and operators to consider distributional and environmental effects. Resilience and fairness must be designed into optimization objectives rather than added after failures or inequities appear. Case illustrations from cloud wide-area networks, data centers, and research networks show that context matters and that universal optimization models are unlikely to succeed.

A forward-looking agenda should prioritize open benchmarks, digital twins, policy-aware orchestration, and interdisciplinary evaluation. If software-defined optical networks are to serve as critical infrastructure for artificial intelligence, cloud services, and public digital systems, their optimization must be judged not only by throughput and delay but also by sustainability, robustness, equity, and legitimacy. The challenge is therefore as much institutional as it is technical, and the most durable solutions will be those that align programmable optical control with the broader public values that infrastructures are expected to uphold.

References

1. McKeown, N., Anderson, T., Balakrishnan, H., Parulkar, G., Peterson, L., Rexford, J., Shenker, S., & Turner, J. (2008). OpenFlow: Enabling innovation in campus networks. *ACM SIGCOMM Computer Communication Review*, 38(2), 69–74. <https://doi.org/10.1145/1355734.1355746>
2. Kreutz, D., Ramos, F. M. V., Verissimo, P. E., Rothenberg, C. E., Azodolmolky, S., & Uhlig, S. (2015). Software-defined networking: A comprehensive survey. *Proceedings of the IEEE*, 103(1), 14–76. <https://doi.org/10.1109/JPROC.2014.2371999>
3. Simeonidou, D., Nejabati, R., & Channegowda, M. (2013). Software-defined optical networks technology and infrastructure: Enabling software-defined optical network operations. *Journal of Optical Communications and Networking*, 5(10), A274–A282. <https://doi.org/10.1364/JOCN.5.00A274>
4. Gerstel, O., Jinno, M., Lord, A., & Yoo, S. J. B. (2012). Elastic optical networking: A new dawn for the optical layer? *IEEE Communications Magazine*, 50(2), S12–S20. <https://doi.org/10.1109/MCOM.2012.6146481>
5. Jinno, M., Takara, H., Kozicki, B., Tsukishima, Y., Sone, Y., & Matsuoka, S. (2009). Spectrum-efficient and scalable elastic optical path network: Architecture, benefits, and enabling technologies. *IEEE Communications Magazine*, 47(11), 66–73. <https://doi.org/10.1109/MCOM.2009.5307468>
6. Tomkos, I., Azodolmolky, S., Sole-Pareta, J., Careglio, D., & Palkopoulou, E. (2014). A tutorial on the flexible optical networking paradigm: State of the art, trends, and research challenges. *Proceedings of the IEEE*, 102(9), 1317–1337. <https://doi.org/10.1109/JPROC.2014.2334693>
7. Zhou, Z., Wei, J., Luo, Y., Clark, K. A., Sillescu, E., Deakin, C., ... & Liu, Z. (2023). Communications with guaranteed bandwidth and low latency using frequency-referenced multiplexing. *Nature Electronics*, 6(9), 694–702.
8. Napoli, A., Bohn, M., Rafique, D., Stavdas, A., Sambo, N., Poti, L., Nölle, M., Fischer, J. K., Riccardi, E., Pagano, A., Castoldi, P., & Tomkos, I. (2015). Next generation elastic optical networks: The vision of the European research project IDEALIST. *IEEE Communications Magazine*, 53(2), 152–162. <https://doi.org/10.1109/MCOM.2015.7045403>
9. Hong, C.-Y., Kandula, S., Mahajan, R., Zhang, M., Gill, V., Nanduri, M., & Wattenhofer, R. (2013). Achieving high utilization with software-driven WAN. *ACM SIGCOMM Computer Communication Review*, 43(4), 15–26. <https://doi.org/10.1145/2534169.2486012>
10. Jain, S., Kumar, A., Mandal, S., Ong, J., Poutievski, L., Singh, A., Venkata, S., Wanderer, J., Zhou, J., Zhu, M., Zolla, J., Hözl, U., Stuart, S., & Vahdat, A. (2013). B4: Experience with a globally-deployed software defined WAN. *ACM SIGCOMM Computer Communication Review*, 43(4), 3–14. <https://doi.org/10.1145/2534169.2486019>
11. Al-Fares, M., Loukissas, A., & Vahdat, A. (2008). A scalable, commodity data center network architecture. *ACM SIGCOMM Computer Communication Review*, 38(4), 63–74. <https://doi.org/10.1145/1402958.1402967>

12. 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. <https://doi.org/10.1145/1496091.1496103>
13. Benson, T., Akella, A., & Maltz, D. A. (2010). Network traffic characteristics of data centers in the wild. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement* (pp. 267–280). ACM. <https://doi.org/10.1145/1879141.1879175>
14. Kandula, S., Sengupta, S., Greenberg, A., Patel, P., & Chaiken, R. (2009). The nature of data center traffic: Measurements and analyses. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement* (pp. 202–208). ACM. <https://doi.org/10.1145/1644893.1644918>
15. Chiaraviglio, L., Mellia, M., & Neri, F. (2012). Minimizing ISP network energy cost: Formulation and solutions. *IEEE/ACM Transactions on Networking*, 20(2), 463–476. <https://doi.org/10.1109/TNET.2011.2161487>
16. Sterbenz, J. P. G., Hutchison, D., Çetinkaya, E. K., Jabbar, A., Rohrer, J. P., Schöller, M., & Smith, P. (2010). Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. *Computer Networks*, 54(8), 1245–1265. <https://doi.org/10.1016/j.comnet.2010.03.005>
17. Clark, D. D., Wroclawski, J., Sollins, K. R., & Braden, R. (2005). Tussle in cyberspace: Defining tomorrow's Internet. *IEEE/ACM Transactions on Networking*, 13(3), 462–475. <https://doi.org/10.1109/TNET.2005.850224>
18. Frischmann, B. M. (2012). *Infrastructure: The social value of shared resources*. Oxford University Press.
19. Open Networking Foundation. (2014). *SDN architecture, issue 1.0*. Open Networking Foundation.
20. Ceccarelli, D., & Lee, Y. (2018). Framework for abstraction and control of traffic engineered networks (ACTN) (RFC 8453). Internet Engineering Task Force. <https://doi.org/10.17487/RFC8453>