

# AI-Driven Optimization of Wavelength-Switched Optical Interconnects for Low-Latency Data Center Networks

Kyano Erhout

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

## Abstract

The rapid expansion of cloud services, artificial intelligence workloads, and latency-sensitive applications has intensified pressure on data center networks to deliver extreme bandwidth with predictable low latency. Electrical packet switching, although highly mature and flexible, faces growing constraints in radix, energy per bit, and tail latency as port speeds and fabric scales increase. Wavelength-switched optical interconnects offer a promising architectural alternative by exploiting wavelength division multiplexing, reconfigurable optical switching, and circuit-like paths that can bypass contention and queuing in electronic fabrics. However, optical switching introduces new challenges in control complexity, synchronization, traffic prediction, fault management, and resource allocation. This paper examines how artificial intelligence can optimize wavelength-switched optical interconnects for low-latency data center networks from a system-level perspective. It analyzes architectural trade-offs between electrical and optical fabrics, discusses the role of machine learning in traffic forecasting, wavelength assignment, switch reconfiguration, and telemetry-driven control, and considers deployment governance, sustainability, robustness, fairness, and policy implications. Special attention is given to clock synchronization and hollow-core fiber as emerging enablers of low-latency wavelength-switched intra-data center interconnects. The paper argues that AI-driven optimization must be understood not merely as a control-layer innovation but as a socio-technical capability that shapes infrastructure economics, operational accountability, tenant fairness, and long-term sustainability.

## Keywords

wavelength switching, optical interconnects, data center networks, artificial intelligence, low latency, hollow core fiber, clock synchronization, network governance, sustainability, optical switching.

## 1. Introduction

Data center networks have become critical infrastructure for modern economies, supporting distributed computation, storage, machine learning, real-time communication, and a growing array of latency-sensitive services. Traditional data center architectures, including hierarchical and folded Clos designs, were built around electrical packet switching because it offers high statistical multiplexing, mature interoperability, and simple operational models [1], [2]. Yet the scale and traffic profile of contemporary data centers have exposed fundamental limits. Measurements of production data centers show that traffic is often bursty, asymmetric, and dominated by east-west communication between servers rather than north-south traffic with external clients [11], [12]. This internal traffic pattern places enormous pressure on the switching fabric, especially when applications require predictable tail latency rather than only

high average throughput [15]. Artificial intelligence workloads intensify these demands because training and inference systems frequently exchange large tensors, synchronize parameters, and coordinate many parallel workers. In such environments, even modest queuing delays can reduce accelerator utilization and increase the cost of computation.

Optical interconnects have long been proposed as a solution to the bandwidth and energy bottlenecks of electrical fabrics. Wavelength division multiplexing allows many independent channels to share a single fiber, while wavelength-selective switches and optical cross-connects can establish reconfigurable lightpaths without electronic packet processing at every hop [6], [7]. Wavelength-switched optical interconnects are particularly attractive for low-latency data center networks because they can create direct optical paths between endpoints, reducing the number of electronic buffering and forwarding stages. However, the benefits of optical switching are not automatic. Circuit-like behavior can waste capacity when traffic is unpredictable, and reconfiguration latency can undermine the very low-latency goals that motivate the architecture. Artificial intelligence is therefore increasingly viewed as an essential control technology for predicting traffic, scheduling wavelength assignments, managing contention, and coordinating hybrid electrical and optical resources [16], [20], [21]. This paper analyzes that intersection through a system-level lens, emphasizing architecture, governance, sustainability, robustness, fairness, and policy rather than isolated device performance.

## **2. Architectural Foundations and the Limits of Electrical Switching**

The modern data center fabric evolved from the need to connect thousands or tens of thousands of servers with high bisection bandwidth and low operational complexity. Early scalable designs, such as commodity Clos architectures, replaced expensive proprietary switches with regular topologies and multipath routing [1]. Subsequent systems such as VL2 and PortLand improved address virtualization, traffic engineering, and layer-two scalability, enabling large cloud deployments to treat the network as a flexible substrate for virtual machines and distributed applications [2], [3]. These architectures remain highly effective because electrical packet switching provides fine-grained statistical multiplexing, rapid adaptation to changing traffic, and well-understood congestion control. Hybrid electrical and optical designs such as Helios and c-Through demonstrated that optical circuits can be added to electrical fabrics to offload elephant flows or provide reconfigurable bandwidth between pods [4], [5]. Their results also revealed the central trade-off: optical circuits excel for large, stable flows, while electrical switching remains superior for small, bursty, or unpredictable communication.

As port speeds rise, electrical switching becomes increasingly difficult to scale economically and energetically. High-radix switches require complex silicon, advanced packaging, and substantial power delivery and cooling. The energy consumed by serialization, deserialization, buffering, and forwarding grows with aggregate bandwidth, and the thermal density of switch chips constrains rack and data hall design [22]. More fundamentally, electrical fabrics impose queuing at multiple stages, and queuing is the primary source of tail latency in congested data center networks. Congestion control mechanisms such as RTT-based schemes can mitigate incast and buffer buildup, but they cannot eliminate the physical and architectural sources of delay [15]. Wavelength-switched optical interconnects address this problem by offering paths that can be provisioned as lightpaths rather than packet-by-packet routes. In principle, an optical circuit can carry traffic with minimal electronic processing, reducing both latency and energy per bit. The challenge is to make such circuits responsive enough for dynamic data

center workloads without sacrificing the flexibility that makes electrical packet switching so successful.

### **3. Wavelength-Switched Optical Interconnects as a System-Level Response**

Wavelength-switched optical interconnects use tunable transceivers, arrayed waveguide gratings, wavelength-selective switches, and optical cross-connects to route signals across a fiber plant according to wavelength [6], [9]. A wavelength can be assigned to a source-destination pair, creating a dedicated optical channel that avoids electronic contention at intermediate switches. In a data center context, this capability can be used for rack-to-rack, pod-to-pod, or accelerator-to-accelerator communication, especially when traffic matrices exhibit stable heavy hitters or predictable collective communication patterns. Systems such as WaveCube and Projector explored scalable optical topologies and agile reconfigurable interconnects, showing that optical fabrics can provide high bandwidth and fault tolerance when combined with intelligent control [13], [14]. More recent designs such as Sirius demonstrated nanosecond-scale optical switching and flat network topologies, suggesting that the boundary between packet switching and circuit switching may become increasingly blurred [19]. These systems highlight that wavelength switching is not a single technology but a design space spanning switching speed, port count, reconfiguration granularity, and control-plane sophistication.

The system-level appeal of wavelength switching lies in its ability to bypass unnecessary electronic stages. In an electrical leaf-spine fabric, a packet may traverse multiple switch chips, each adding serialization delay, forwarding latency, buffer queuing, and energy cost. Optical lightpaths can reduce the number of electronic conversions, especially for long-lived flows or synchronized communication phases. However, optical switching also introduces new constraints. Wavelength resources are finite, and a wavelength assigned to one path cannot simultaneously serve another. Reconfiguration may require guard times, tuning delays, and synchronization across transceivers and switches. If the control plane cannot predict traffic accurately, the optical fabric may oscillate between underutilization and contention. Hybrid architectures that combine electrical packet switching for fine-grained traffic with optical wavelength switching for bulk or predictable traffic are therefore likely to dominate in the near term [4], [5], [10]. The optimization problem becomes one of deciding what should be switched optically, when, and under what governance constraints.

### **4. AI-Driven Optimization Across Layers**

Artificial intelligence offers a powerful toolkit for managing the complexity of wavelength-switched optical interconnects. Machine learning models can forecast traffic matrices from historical telemetry, application-level signals, and orchestration metadata, allowing the network to pre-provision wavelengths before demand materializes [20], [21]. Reinforcement learning can learn reconfiguration policies that balance latency, throughput, energy, and fairness across tenants. Graph neural networks can model the topology and state of a hybrid electrical and optical fabric, while time-series models can predict burstiness and incast patterns. Knowledge-defined networking extends this idea by separating a knowledge plane from the control and data planes, enabling closed-loop optimization based on analytics rather than static rules [16]. In such a framework, AI does not replace conventional routing or congestion control; it orchestrates them, adjusting optical paths, electrical paths, and traffic steering policies in response to changing conditions.

The application of AI to optical interconnects must nevertheless confront several structural challenges. First, training data may be scarce or unrepresentative, especially for rare failure events, reconfiguration transients, and adversarial traffic. Second, optimization objectives may conflict: minimizing tail latency may require overprovisioning, while minimizing energy may favor consolidation and lower utilization. Third, the control loop must operate at multiple timescales. Some decisions, such as wavelength assignment, may be made over milliseconds or seconds, while packet-level scheduling occurs in microseconds or nanoseconds. AI models that are too slow or too centralized cannot close these loops effectively. A practical architecture may therefore combine fast local heuristics with slower global learning, using telemetry and digital twins to validate policies before deployment [14], [16]. The goal is not full autonomy but bounded, auditable intelligence that improves optical fabric utilization without undermining reliability or operator trust.

## **5. Latency, Synchronization, and the Hollow-Core Fiber Frontier**

Low-latency data center networks require not only high bandwidth but also deterministic delay. Wavelength-switched optical interconnects can reduce queuing and forwarding delays, but they introduce synchronization requirements. Tunable lasers, wavelength-selective switches, and receivers must coordinate their state changes so that data is not lost during reconfiguration. Clock synchronization across endpoints and switches becomes critical, especially when optical paths are established and torn down rapidly. Recent work on low-latency wavelength-switched clock-synchronized intra-data center interconnects enabled by hollow core fiber points to an important direction: the physical medium itself can reduce propagation delay while supporting precise synchronization [17]. Hollow core fiber guides light primarily in air rather than glass, offering lower latency and potentially lower nonlinearity than conventional silica fiber. In a data center, even modest reductions in propagation delay can matter for distributed consensus, accelerator synchronization, and real-time inference.

The synchronization challenge is both technical and architectural. If optical paths are reconfigured frequently, the control plane must maintain a consistent view of time across transceivers, switches, and servers. Synchronization errors can cause guard-time inefficiency, packet loss, or incorrect wavelength assignment. AI can assist by predicting synchronization drift, scheduling reconfiguration during low-risk windows, and identifying anomalies that precede clock instability. Yet synchronization also has governance implications. Time is a shared resource, and different tenants may have different latency and accuracy requirements. A multi-tenant optical fabric must therefore allocate synchronization quality fairly and transparently, not merely maximize aggregate performance. Hollow core fiber and advanced clock distribution can expand the design space, but they also raise questions about cost, standardization, deployment disruption, and the embodied impact of new fiber infrastructure [18]. The long-term promise of low-latency optical interconnects depends on aligning device physics, control intelligence, and operational policy.

## **6. Deployment, Governance, and Operational Control**

Deploying wavelength-switched optical interconnects in production data centers is not simply a matter of installing new hardware. It requires changes in control planes, monitoring, fault management, and organizational workflows. Software-defined networking provides a foundation for programmable optical control, allowing operators to express policies that map application requirements to optical and electrical paths [9], [10]. However, optical devices often expose vendor-specific interfaces, and multi-vendor interoperability remains limited. An

AI-driven control plane must therefore abstract heterogeneous hardware while preserving enough low-level visibility to diagnose physical-layer problems. Telemetry is essential: optical signal-to-noise ratio, power levels, wavelength drift, switch state, and queue occupancy must be streamed into analytics pipelines. Without reliable telemetry, machine learning models may optimize against stale or incomplete views of the network [16]. Governance must also define who can reconfigure optical paths, under what conditions, and with what rollback mechanisms.

Operational control becomes more complex in multi-tenant environments. A cloud provider may offer latency-sensitive virtual networks, best-effort storage traffic, and internal management traffic on the same physical infrastructure. Wavelength switching can provide isolation and predictable performance, but it can also create new forms of contention if wavelengths are allocated unfairly or if reconfiguration disrupts other tenants. Policy engines must encode service-level objectives, priority classes, regulatory constraints, and fairness criteria. AI can help enforce these policies by detecting violations, predicting SLA risk, and recommending corrective actions, but it should not be the sole authority for consequential decisions. Human operators need explainable models, audit trails, and fail-safe defaults. The most credible deployments will likely use AI as an advisory and optimization layer within a governed control architecture, not as an opaque replacement for accountable network management [20], [21].

## **7. Sustainability, Economics, and Resource Allocation**

Sustainability is a first-order concern for data center networks because energy consumption, cooling, and embodied carbon scale with infrastructure expansion. Optical interconnects can improve energy efficiency by reducing electronic switching and conversion, but they also require lasers, amplifiers, cooling, and specialized materials. The net benefit depends on utilization, lifetime, manufacturing impact, and the energy mix of the hosting region [22]. Wavelength switching can improve sustainability by consolidating traffic onto fewer optical paths and powering down idle electrical resources, but aggressive consolidation may increase tail latency or reduce resilience. AI-driven optimization can search this trade-off by forecasting demand and selectively activating optical capacity only when needed. Yet sustainability metrics must include embodied carbon and material scarcity, not only operational power. Otherwise, an architecture that appears efficient at runtime may shift environmental burdens upstream.

Economic considerations are similarly multidimensional. Optical transceivers and switches have high capital costs, and reconfigurable fabrics may require new fiber plants, patch panels, and cooling designs. Operators must weigh these costs against savings in energy, latency-sensitive revenue, and accelerator utilization. AI can improve return on investment by increasing utilization and reducing overprovisioning, but it also introduces costs for data collection, model training, inference hardware, and skilled personnel. In multi-tenant data centers, the economic benefits of optical switching may accrue unevenly. Large tenants with predictable bulk traffic may benefit most from dedicated wavelengths, while small tenants may rely on electrical switching and shared best-effort paths. Fair resource allocation therefore requires pricing, quotas, and scheduling policies that prevent optical capacity from becoming a premium enclave available only to the largest customers. Sustainability and fairness are not secondary concerns; they shape whether wavelength-switched optical interconnects are socially and economically durable.

## **8. Robustness, Fairness, and Policy Implications**

Robustness is essential because data center networks are critical infrastructure. Optical fabrics can be vulnerable to fiber cuts, component failures, wavelength drift, and control-plane errors. AI models may improve fault prediction, but they can also introduce new failure modes through distribution shift, adversarial inputs, or correlated decision-making. A single global model that recommends reconfiguration across many pods could propagate errors rapidly if it is wrong or compromised. Robust design therefore favors diversity, local autonomy, graceful degradation, and fallback to electrical switching when optical paths are unavailable [2], [3]. Failure management should be tested under stress, including partial failures, asymmetric traffic, and degraded synchronization. The network must remain predictable enough for congestion control and application-level timeouts, even when optical resources are being reconfigured [15]. Resilience is not merely a hardware property; it is a property of the combined control, policy, and operational system.

Fairness and policy implications extend beyond technical performance. If AI-driven optical interconnects prioritize some tenants, applications, or regions over others, they may create new forms of digital inequality. A financial trading firm may pay for ultra-low-latency wavelengths, while a public-interest application may receive best-effort service. Within a single cloud provider, internal services may compete with external customers for optical capacity. Regulators and standards bodies may eventually need to address transparency, nondiscrimination, and critical-service priority in advanced data center networks. Privacy is also relevant because traffic prediction may rely on application metadata, tenant identifiers, or workload traces. Governance frameworks should specify what data can be collected, how long it is retained, and how models are audited. AI-driven optimization should be accompanied by human oversight, appeal mechanisms for tenants, and clear accountability when automated decisions cause outages or SLA violations [20], [21]. Policy is not an external constraint on technical design; it is part of the design space.

## **9. Case Illustrations and Cross-Domain Comparisons**

High-performance computing offers an instructive comparison. Supercomputing interconnects have long used optical links and specialized topologies to support collective communication and tightly coupled simulation. Photonic networks-on-chip and optical interconnection networks for high-performance computing demonstrated the potential of wavelength routing for predictable, high-bandwidth communication [6], [8]. Data centers, however, host more heterogeneous and dynamic workloads than many HPC systems. They must support multi-tenant isolation, elastic scaling, and diverse application requirements. This difference explains why hybrid electrical and optical architectures have been more common than fully optical designs. The HPC experience suggests that optical interconnects work best when communication patterns are known and synchronization is central. Data center AI workloads increasingly exhibit similar characteristics during training, but inference and microservices remain more bursty and unpredictable. AI-driven control can bridge this gap by classifying workloads and assigning them to the most appropriate switching regime.

Cross-domain comparisons with telecommunications and wide-area networks also illuminate governance challenges. Telecom operators have decades of experience with optical transport, wavelength management, and service-level agreements. Data center operators, by contrast, often prioritize rapid iteration and commodity economics. Bringing wavelength switching into the data center therefore requires translating telecom practices into a cloud-native operational culture. Systems such as Sirius and Projector illustrate the value of co-designing hardware, topology, and control [14], [19]. They also show that successful deployment depends on

workload-aware scheduling, not only faster switching. Similarly, financial exchanges and real-time control systems illustrate the premium placed on deterministic latency. If optical interconnects can provide predictable low latency at scale, they may enable new classes of applications, but they may also entrench advantages for organizations that can afford premium connectivity. The policy question is whether such capabilities should be treated as ordinary commercial services or as shared infrastructure with public-interest obligations [22].

## **10. Future Research Directions**

Future research should pursue several interconnected directions. First, AI models must become more sample-efficient and explainable so that operators can trust them in critical optical control loops. Digital twins of hybrid electrical and optical fabrics could allow safe exploration of reconfiguration policies before production deployment [16]. Second, hardware and control should be co-designed. Switching speed, tuning range, guard time, and synchronization accuracy determine what AI can achieve. Conversely, AI requirements should influence device specifications, such as telemetry granularity and programmable interfaces. Third, sustainability metrics should be integrated into optimization objectives from the beginning, including embodied carbon, water use, and material criticality [18], [22]. Fourth, fairness and policy-aware optimization should be studied as technical problems. This includes multi-objective scheduling, tenant-facing transparency, and mechanisms for auditing automated decisions. Fifth, standardization and interoperability are essential for broad adoption. Without open interfaces, AI-driven optical control may remain locked into single-vendor silos.

Another promising direction is the integration of wavelength switching with emerging accelerator-centric architectures. As machine learning training clusters grow, communication between accelerators becomes a dominant cost. Optical interconnects can provide reconfigurable all-to-all or torus-like connectivity, while AI schedulers predict collective communication phases and pre-establish lightpaths [14], [17], [19]. This requires close coordination between application frameworks, resource managers, and network control planes. It also raises questions about failure containment and fair sharing when multiple training jobs compete for optical bandwidth. Finally, researchers should examine the long-term governance of AI-driven infrastructure. Who is accountable when an automated wavelength assignment causes a cascading failure? How are tenants notified and compensated? What transparency is owed to regulators and the public? These questions are not peripheral to systems research; they determine whether advanced optical interconnects can be deployed responsibly and at scale.

## **11. Conclusion**

Wavelength-switched optical interconnects represent a significant architectural opportunity for low-latency data center networks. By reducing electronic switching stages and enabling reconfigurable lightpaths, they can improve latency, energy efficiency, and bandwidth scalability for demanding workloads. Artificial intelligence can make these fabrics more responsive by predicting traffic, assigning wavelengths, scheduling reconfiguration, and managing hybrid electrical and optical resources. Yet the success of AI-driven optimization depends on more than algorithmic performance. It depends on synchronization, telemetry, interoperability, governance, sustainability, fairness, and robust operational practice. Hollow core fiber and clock-synchronized wavelength switching point toward a future in which physical-layer advances and intelligent control combine to reduce latency further, but such advances must be embedded in accountable socio-technical systems. The most credible path

forward is not full automation or purely optical replacement, but a governed hybrid architecture in which AI optimizes within clear boundaries, human operators retain oversight, and policy ensures that the benefits of low-latency optical infrastructure are distributed fairly and sustainably.

## References

1. Al-Fares, M., Loukissas, A., & Vahdat, A. (2008). A scalable, commodity data center network architecture. In *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication* (pp. 63–74). Association for Computing Machinery.
2. Greenberg, A., Hamilton, J. R., Jain, N., Kandula, S., Kim, C., Lahiri, P., Maltz, D. A., Patel, P., & Sengupta, S. (2009). VL2: A scalable and flexible data center network. In *Proceedings of the ACM SIGCOMM 2009 Conference on Data Communication* (pp. 51–62). Association for Computing Machinery.
3. Mysore, R. N., Pamboris, A., Farrington, N., Huang, N., Miri, P., Radhakrishnan, S., Subramanya, V., & Vahdat, A. (2009). PortLand: A scalable fault-tolerant layer 2 data center network fabric. In *Proceedings of the ACM SIGCOMM 2009 Conference on Data Communication* (pp. 39–50). Association for Computing Machinery.
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. In *Proceedings of the ACM SIGCOMM 2010 Conference* (pp. 339–350). Association for Computing Machinery.
5. Wang, G., Andersen, D. G., Kaminsky, M., Papagiannaki, K., Ng, T. S. E., Kozuch, M., & Ryan, M. (2010). c-Through: Part-time optics in data centers. In *Proceedings of the ACM SIGCOMM 2010 Conference* (pp. 327–338). Association for Computing Machinery.
6. Kachris, C., & Tomkos, I. (2012). A survey on optical interconnects for data centers. *IEEE Communications Surveys & Tutorials*, 14(4), 1021–1036.
7. Biberman, A., & Bergman, K. (2012). Optical interconnection networks for high-performance computing systems. *Reports on Progress in Physics*, 75(4), 046402.
8. Shacham, A., Bergman, K., & Carloni, L. P. (2008). Photonic networks-on-chip for future generations of chip multiprocessors. *IEEE Transactions on Computers*, 57(9), 1246–1260.
9. 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.
10. Porter, G., Strong, R., Farrington, N., Forencich, A., Chen-Sun, P., Rosing, T., Fainman, Y., Papen, G., & Vahdat, A. (2013). Integrating microsecond circuit switching into the data center. In *Proceedings of the ACM SIGCOMM 2013 Conference* (pp. 447–458). Association for Computing Machinery.
11. 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). Association for Computing Machinery.

12. 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). Association for Computing Machinery.
13. Chen, K., Wen, X., Ma, X., Chen, Y., Xia, Y., Ni, X., & Zheng, Y. (2015). WaveCube: A scalable, fault-tolerant, high-performance optical data center architecture. In *Proceedings of IEEE INFOCOM 2015* (pp. 1903–1911). IEEE.
14. Ghobadi, M., Mahajan, R., Phanishayee, A., Devanur, N., Kulkarni, J., Ranade, G., Blanche, P.-A., Rastegarfar, H., Glick, M., & Kilper, D. (2016). Projector: Agile reconfigurable data center interconnect. In *Proceedings of the ACM SIGCOMM 2016 Conference* (pp. 216–229). Association for Computing Machinery.
15. Mittal, R., Lam, V. T., Dukkupati, N., Blem, E., Wassel, H., Ghobadi, M., Vahdat, A., Wang, Y., Wetherall, D., & Zats, D. (2015). TIMELY: RTT-based congestion control for the datacenter. In *Proceedings of the ACM SIGCOMM 2015 Conference* (pp. 537–550). Association for Computing Machinery.
16. Mestres, A., Rodriguez-Natal, A., Carner, J., Barlet-Ros, P., Alarcón, E., Solé, M., Muntés-Mulero, V., Meyer, D., Barkai, S., Hibbett, M. J., Estrada, G., Ma'ruf, K., Coras, F., Ermagan, V., Latapie, H., Cassar, C., Evans, D., Maino, F., Walrand, J., & Cabellos, A. (2017). Knowledge-defined networking. *ACM SIGCOMM Computer Communication Review*, 47(3), 2–10.
17. Zhou, Zichuan, et al. "Low-latency wavelength-switched clock-synchronized intra-data center interconnects enabled by hollow core fiber." *Optics Express* 31.15 (2023): 24739–24748.
18. Cheng, Q., Bahadori, M., Glick, M., Rumley, S., & Bergman, K. (2018). Recent advances in optical technologies for data centers: A review. *Optica*, 5(11), 1354–1370.
19. Ballani, H., Costa, P., Behrendt, R., Cletheroe, D., Haller, I., Jozwik, K., Karinou, F., Lange, S., Li, K., Liu, Y., Luk, W., Portnoi, M., Schedl, D., She, D., Shepherd, B., Thomsen, B., & Williams, H. (2020). Sirius: A flat datacenter network with nanosecond optical switching. In *Proceedings of the ACM SIGCOMM 2020 Conference* (pp. 349–363). Association for Computing Machinery.
20. 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.
21. Fadlullah, Z. M., Tang, F., Mao, B., Medhi, N., Akhtar, Z., Guerrero, I. F., & Guizani, M. (2017). State-of-the-art deep learning in computer networks: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 19(4), 2493–2528.
22. 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.