

Intelligent Resource Scheduling for Hybrid Optical-Electrical Interconnects in Cloud Computing Infrastructure

Michal Wodka

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
michal.wodka269@buffalo.edu

Amruta Ravel

Department of Electrical Engineering and Computer Science, University of Missouri,
Columbia, MO, USA.
amrutawork@missouri.edu

Giuliana Agnesi

Department of Computer Science, Binghamton University, Binghamton, NY, USA.
giuliana.agnesi@binghamton.edu

Larissa Sussebier

Department of Computer Science, University of North Texas, Denton, TX, USA.
larissa877@unt.edu

Abstract

Cloud computing infrastructure increasingly depends on interconnect fabrics that can sustain high bandwidth, low latency, and predictable energy behavior across heterogeneous workloads. Hybrid optical-electrical interconnects offer a promising architectural path because they combine the maturity, granularity, and buffering advantages of electrical switching with the bandwidth density, distance reach, and energy efficiency of optical communication. However, the coexistence of electrical and optical resources creates complex scheduling problems that cannot be reduced to conventional packet routing or virtual machine placement. This paper examines intelligent resource scheduling for hybrid optical-electrical interconnects from a system-level and socio-technical perspective. It argues that effective schedulers must coordinate workload characterization, topology awareness, optical circuit setup, electrical packet buffering, failure recovery, fairness constraints, and sustainability objectives. The analysis discusses structural trade-offs among reconfiguration latency, port count, spectrum assignment, and utilization, and it considers how machine learning and optimization-based control can operate within governance boundaries that protect tenant autonomy and network stability. The paper also explores deployment pathways, infrastructure economics, robustness under correlated failures, and policy implications for multi-tenant cloud environments. The central conclusion is that hybrid optical-electrical scheduling should be understood as an infrastructure governance problem as much as an algorithmic one. Sustainable and equitable cloud systems will require schedulers that can reason across time scales, ownership boundaries, energy sources, and service-level commitments while remaining explainable to operators, regulators, and tenants.

Keywords

hybrid optical-electrical interconnects; cloud computing; resource scheduling; data center networks; optical switching; sustainability; governance; robustness.

1. Introduction

Cloud computing has evolved from a utility-like abstraction for computation and storage into a foundational layer of digital economies, scientific research, and public services. Early cloud research emphasized elasticity, virtualization, and pay-as-you-go service models, but the physical substrate of cloud computing has become equally important because bandwidth, latency, and energy consumption increasingly determine what applications can be offered at scale [1]. Modern data centers host diverse workloads, including latency-sensitive microservices, distributed machine learning, high-performance analytics, storage replication, and interactive media. These workloads do not share a single traffic pattern, and their requirements often conflict. A scheduler that optimizes average throughput may starve latency-critical services, while a scheduler that reserves capacity for peak demand may waste energy during off-peak periods. The interconnect is therefore not merely a transport mechanism but a control surface through which cloud operators shape performance, cost, and sustainability [2].

Electrical packet switching has long dominated data center networks because it provides fine-grained multiplexing, mature congestion control, and flexible routing. Yet the growth of east-west traffic, accelerated computing, and disaggregated storage has exposed limits in bandwidth density, reach, and energy per bit [3]. Optical interconnects can carry high-capacity signals over longer distances with lower loss and potentially lower energy per bit, and they can support wavelength-division multiplexing for parallel channels [4]. However, optical switching often lacks the same fine-grained buffering and statistical multiplexing capabilities as electrical switching. Hybrid optical-electrical architectures attempt to combine the strengths of both domains by using electrical switches for bursty, short-lived, and latency-sensitive flows while using optical circuits for large, predictable, or long-lived transfers [5]. This combination creates a rich scheduling space, but also introduces new sources of complexity, including circuit setup delay, wavelength contention, topology reconfiguration, and dependence on traffic prediction [6].

Intelligent resource scheduling for hybrid optical-electrical interconnects must therefore operate across layers and time scales. At the millisecond scale, electrical switches handle packets and enforce fairness. At the second or minute scale, optical circuits may be established, torn down, or rerouted. At longer horizons, operators plan capacity, procure equipment, and negotiate energy contracts. A scheduler that treats these scales independently will miss opportunities for cross-layer optimization, but a scheduler that centralizes all decisions may become fragile, opaque, or politically unacceptable in multi-tenant environments. The tension between optimization and governance is central to this paper. The argument is that hybrid optical-electrical scheduling is not only a technical problem of matching flows to resources; it is also a problem of allocating shared infrastructure in ways that are robust, fair, sustainable, and accountable [7]. Recent work on low-latency wavelength-switched and clock-synchronized intra-data center interconnects enabled by hollow-core fiber illustrates how physical-layer advances can reduce latency and reconfiguration penalties, but such advances also raise expectations for tighter coordination between optical control and higher-level scheduling [8].

2. Architectural Foundations and Structural Trade-offs

Hybrid optical-electrical interconnects can be organized in several architectural patterns. In one pattern, a electrical leaf-spine fabric carries most traffic, while optical links provide offload paths for elephant flows or inter-pod bulk transfers. In another pattern, optical circuit switches form a reconfigurable core that connects pods or racks, and electrical switches handle local fan-out and burst absorption. A third pattern uses optical switching for rack-to-rack communication and electrical switching for intra-rack communication. Each pattern embodies a different trade-off between reconfiguration speed, port count, cabling complexity, energy use, and failure domain size [9]. The scheduler must understand these physical constraints because they determine which flows can be moved to optical paths and how quickly the fabric can respond to changing demand. For example, a hybrid architecture with slow micro-electromechanical optical switches may be excellent for long-lived backups but poor for interactive services. A architecture with fast wavelength-switched lasers may support more dynamic offloading but may require tighter clock synchronization and thermal control [10].

The most basic structural trade-off is between granularity and efficiency. Electrical packet switching offers fine granularity because every packet can be routed independently. Optical circuit switching offers high efficiency for large transfers because a dedicated lightpath avoids per-packet processing and buffering, but it may waste capacity if the circuit is underutilized. The scheduler must estimate flow size, duration, and predictability before deciding whether to reserve an optical circuit. Misclassification is costly. Sending a short flow over an optical circuit may incur setup delay and reduce effective throughput, while keeping a large flow on the electrical fabric may create congestion and increase energy per bit. This problem is compounded by the fact that cloud workloads are often bursty and correlated. A distributed training job may generate synchronized communication phases, while a web service may generate unpredictable request-response traffic. The scheduler needs both historical models and online adaptation to avoid brittle decisions [11].

A second trade-off concerns topology flexibility. Electrical fabrics can be over-subscribed and reconfigured logically through routing tables. Optical fabrics are often more physically constrained because wavelengths, ports, and switch states determine connectivity. If the optical topology is fixed or slowly reconfigurable, the scheduler must work within a set of available lightpaths. If the topology is highly reconfigurable, the scheduler must account for transient disruptions during switching and for the possibility that multiple tenants compete for the same optical resources [12]. In multi-tenant clouds, this competition raises fairness questions. A tenant with large, predictable flows may be able to reserve optical capacity, while many small tenants may remain on the electrical fabric. If optical capacity is faster or cheaper per bit, this separation can create performance inequality unless scheduling policies explicitly compensate for it. The design of hybrid interconnects therefore has distributive consequences that extend beyond engineering efficiency [13].

A third trade-off involves resilience. Electrical fabrics are generally robust because traffic can be rerouted around failed links or switches. Optical fabrics can also be protected, but protection may require reserved wavelengths or redundant paths that reduce utilization. In hybrid systems, failures can cross domains. An optical switch failure may push traffic onto electrical links, causing congestion and power spikes. An electrical switch failure may prevent local traffic from reaching optical gateways. The scheduler must therefore treat failures as normal operating conditions rather than rare exceptions. It should maintain graceful degradation, prioritize critical control traffic, and avoid oscillations in which recovery actions

in one domain trigger new failures in another [14]. This requirement argues for schedulers that are conservative in their assumptions and transparent in their fallback behavior.

3. Workload Characterization and Scheduling Objectives

Intelligent scheduling begins with workload characterization. Not all cloud traffic is equal, and not all traffic benefits from optical offloading. Latency-sensitive flows, such as remote procedure calls, database queries, and control messages, often require low jitter and immediate delivery. They may be poor candidates for optical circuits if setup time is significant. Throughput-oriented flows, such as virtual machine migration, backup, data replication, and model checkpointing, may tolerate some delay and can benefit from dedicated high-capacity paths. The scheduler must classify flows using features such as size, duration, source-destination pair, application class, tenant service-level agreement, and temporal pattern. Classification can be based on historical telemetry, packet sampling, or application-provided hints [15]. However, application hints can be inaccurate or strategic, especially in multi-tenant environments where tenants may seek priority. The scheduler therefore needs verification mechanisms and incentives that discourage misreporting.

Scheduling objectives are also multi-dimensional. Throughput, latency, energy, cost, fairness, and reliability cannot be optimized simultaneously without trade-offs. A scheduler that minimizes energy may consolidate traffic onto optical links and power down electrical ports, but this may reduce resilience and increase reconfiguration delays. A scheduler that minimizes latency may keep electrical capacity active even when utilization is low, increasing energy consumption. A scheduler that maximizes fairness may sacrifice some efficiency by reserving resources for small tenants. These trade-offs are not merely technical; they reflect organizational priorities and regulatory expectations. In regulated markets, for example, cloud providers may face requirements for service continuity, data localization, or nondiscriminatory access. The scheduler becomes a site where policy goals are translated into operational rules [16].

The presence of hybrid interconnects also changes the meaning of resource placement. In a purely electrical data center, placing two virtual machines in the same rack can reduce latency and network load. In a hybrid data center, placement may also determine whether traffic can use an optical path or must traverse electrical switches. If the optical fabric connects pods rather than racks, then placing related workloads in different pods may enable optical offloading but increase inter-pod latency. If the optical fabric connects racks within a pod, then placement can exploit high-capacity local lightpaths. The scheduler must therefore co-optimize placement and routing. This co-optimization is difficult because placement decisions persist for minutes or hours, while traffic patterns may change in seconds. A practical approach is to use hierarchical scheduling: a slower placement layer reserves capacity for workload groups, while a faster routing layer adapts to instantaneous demand within those reservations [17].

4. Intelligent Scheduling across Time, Space, and Spectrum

Intelligent scheduling in hybrid optical-electrical networks can be understood as reasoning across three dimensions: time, space, and spectrum. Time concerns when to establish, maintain, or tear down optical circuits and how long to keep electrical resources active. Space concerns which endpoints, racks, pods, or switches should be connected. Spectrum concerns which wavelengths or frequency channels are assigned to which circuits. These dimensions interact. A wavelength assignment that is efficient at one moment may block future

reconfigurations. A spatial topology that minimizes current congestion may create long detours for future traffic. A temporal policy that reacts too quickly may cause control-plane churn, while a policy that reacts too slowly may miss opportunities for offloading. The scheduler must balance stability and adaptivity [18].

Machine learning can help address this complexity by predicting demand, detecting anomalous traffic, and learning policies from operational data. Reinforcement learning, for example, can learn when to switch flows to optical paths by trial and error in simulation or in a controlled deployment. However, learning-based scheduling raises concerns about safety, explainability, and generalization. A policy trained on one data center may fail in another with different hardware, workload mix, or energy prices. A policy that is optimal on average may make rare but catastrophic decisions during failures. For these reasons, learning should be constrained by rule-based safety envelopes and validated through digital twins or canary deployments. The scheduler should be able to explain why a flow was moved to an optical path, why a wavelength was reserved, and why a tenant was denied priority [19].

Spectral efficiency is another key consideration. Wavelength-division multiplexing allows multiple optical circuits to share a fiber, but the number of wavelengths is finite and contention can occur. The scheduler must assign wavelengths to circuits in a way that minimizes fragmentation and supports future requests. This is similar to spectrum assignment in wireless networks, but with different constraints because optical channels are highly stable and can carry very high capacity. If the scheduler reserves wavelengths for premium tenants, it may reduce spectrum availability for best-effort traffic. If it allows dynamic sharing, it may require fast tunable lasers and careful synchronization. Recent advances in hollow-core fiber and clock-synchronized wavelength switching suggest that some physical-layer penalties can be reduced, but they do not eliminate the need for scheduling policies that manage contention fairly [8]. The scheduler must also consider the energy cost of tuning, cooling, and amplification, which may offset the savings from optical offloading if circuits are too short-lived [20].

Spatial scheduling involves deciding which endpoints should be connected by optical paths and which should remain on the electrical fabric. In a pod-based architecture, the scheduler may connect pods that exchange large amounts of data, while keeping intra-pod traffic on electrical switches. In a rack-scale architecture, the scheduler may connect specific top-of-rack switches for distributed storage or machine learning. The choice depends on traffic locality, cable topology, and switch port availability. If the scheduler has global visibility, it can avoid creating circuits that cross congested areas. If visibility is limited, it may rely on local agents that report summaries to a central controller. This distributed design can improve scalability but may lead to inconsistent decisions. The governance challenge is to define authority boundaries: what can a local agent decide autonomously, and what requires central approval? In multi-tenant clouds, these boundaries also determine whether tenants can influence optical resource allocation directly or only through service-level agreements.

5. Robustness, Fairness, and Governance

Robustness in hybrid optical-electrical interconnects requires more than redundant hardware. It requires scheduling policies that anticipate correlated failures, cascading effects, and control-plane partitions. For example, a fiber cut may cause multiple optical circuits to fail simultaneously. If the scheduler reroutes all affected traffic to the electrical fabric, the electrical switches may become overloaded. If it instead throttles low-priority traffic, it may preserve stability but violate service expectations. A robust scheduler should have

precomputed fallback plans, clear priority classes, and mechanisms for graceful degradation. It should also avoid single points of failure in the control plane. If a central optical controller fails, local electrical switches should continue forwarding traffic with safe defaults. If a machine learning model becomes unavailable, the system should fall back to static rules rather than freeze or make arbitrary decisions [2].

Fairness is especially challenging because hybrid interconnects create new dimensions of differentiation. Tenants may differ in their ability to generate predictable large flows, their willingness to pay for premium optical paths, and their tolerance for latency. A scheduler that allocates optical capacity to the largest flows may achieve high utilization but disadvantage small tenants. A scheduler that enforces equal access to optical paths may waste capacity on flows that do not benefit from them. Fairness must therefore be defined in context. One approach is to guarantee minimum electrical resources for all tenants while offering optical resources as a best-effort or auctioned tier. Another approach is to use weighted fair sharing across both electrical and optical domains, with weights based on service-level agreements and social or regulatory priorities. The choice has implications for market structure and innovation. If optical resources are concentrated among a few large tenants, smaller tenants may face higher costs or lower performance, reducing competition in cloud markets [3].

Governance enters because scheduling decisions are not purely technical. They determine who gets access to scarce capacity, how energy is consumed, and how failures are communicated. In multi-tenant environments, tenants may demand transparency about prioritization, isolation, and performance. Regulators may demand auditability of critical infrastructure. Operators may demand operational simplicity and predictable costs. A governance framework for hybrid optical-electrical scheduling should therefore include clear policies for resource reservation, priority, preemption, and dispute resolution. It should also include technical mechanisms for logging decisions, verifying compliance, and enabling independent audits. These mechanisms must be designed carefully because excessive transparency can reveal sensitive tenant information, while insufficient transparency can undermine trust [18]. The scheduler is thus a socio-technical artifact that encodes organizational values and power relationships.

6. Deployment, Sustainability, and Economic Infrastructure

Deployment of hybrid optical-electrical interconnects is incremental in most cloud environments. Operators rarely replace an entire network at once. Instead, they introduce optical links for specific use cases, such as inter-pod bulk transfer, storage replication, or machine learning training. Early deployments may use static or manually configured optical paths, with scheduling limited to routing and admission control. As operational confidence grows, operators may introduce dynamic optical switching and automated scheduling. This incremental path allows learning and risk management, but it also creates heterogeneity. Some parts of the data center may support fast optical reconfiguration, while others rely on slow or static optics. The scheduler must operate across this heterogeneity without assuming uniform capabilities [4]. It may need to discover available optical resources dynamically and negotiate with legacy electrical domains.

Sustainability is a major driver for hybrid interconnects. Data centers consume significant energy for computation, cooling, and networking. Optical links can reduce energy per bit over distance, but they also require lasers, amplifiers, and cooling. The net benefit depends on utilization, distance, and the energy mix of the local grid. A scheduler that considers energy price and carbon intensity can shift large transfers to times or locations with cleaner energy. It

can also consolidate traffic onto optical links to power down electrical ports. However, such consolidation may reduce redundancy and increase risk. It may also create rebound effects if lower network costs encourage more data movement. Sustainable scheduling must therefore consider total lifecycle impacts, including manufacturing and disposal of optical components. It should also align with broader infrastructure planning, such as renewable energy procurement and cooling optimization [12]. The scheduler becomes one component of a larger energy management system.

Economic considerations are equally important. Hybrid optical-electrical infrastructure requires capital investment in optical switches, transceivers, fibers, and control software. It also requires operational expenditure for energy, maintenance, and skilled personnel. The business case depends on the value of reduced latency, increased bandwidth, and energy savings. In some cases, optical offloading may delay the need for expensive electrical switch upgrades. In other cases, the complexity of hybrid management may outweigh the benefits. The scheduler can influence the business case by improving utilization and reducing manual configuration. It can also enable new services, such as guaranteed low-latency paths for financial trading or high-throughput paths for scientific data transfer. These services may generate revenue but may also raise fairness concerns if they prioritize wealthy tenants over public-interest workloads [15]. Infrastructure economics and social policy are therefore intertwined.

7. Policy, Standards, and Cross-Domain Implications

Policy and standards will shape how hybrid optical-electrical scheduling evolves. Standards for optical networking, data center management, and cloud interoperability can reduce integration costs and improve multi-vendor compatibility. They can also define safety, security, and privacy requirements. For example, standards for telemetry and control interfaces can enable independent auditing of scheduling decisions. Standards for energy reporting can support carbon-aware scheduling. Standards for failure notification can improve resilience across operators. However, standards must balance flexibility and interoperability. Overly rigid standards may stifle innovation, while overly loose standards may lead to fragmentation and lock-in. Policymakers should encourage open interfaces and transparent governance without mandating specific algorithms [16].

Cross-domain implications extend beyond cloud providers. Telecommunications operators, energy utilities, research networks, and government agencies all depend on cloud infrastructure. Hybrid optical interconnects can support edge computing, content delivery, and scientific collaboration. They can also concentrate power in the hands of a few large operators if optical capacity is expensive and difficult to deploy. Public policy may therefore need to promote competition, universal access, and resilience. For example, governments may fund shared optical infrastructure for research and public services. They may require cloud providers to disclose performance and energy metrics. They may also establish dispute resolution mechanisms for tenants that are unfairly deprioritized. These policies are not merely external constraints; they can guide the design of scheduling systems by clarifying societal priorities [19].

Security is another cross-domain concern. Optical circuits can be difficult to tap without detection, but they can also be disrupted by physical attacks on fibers or switches. The control plane for optical scheduling may be a target for denial-of-service attacks or manipulation. If an attacker can influence traffic classification or demand prediction, they may redirect optical resources or cause congestion. The scheduler must therefore include authentication,

authorization, and anomaly detection. It should also limit the impact of compromised components by isolating control domains and enforcing least privilege. In multi-tenant environments, security policies must prevent one tenant from inferring another tenant's traffic patterns through side channels in optical resource allocation. These requirements add complexity but are essential for trustworthy infrastructure [20].

8. Conclusion

Intelligent resource scheduling for hybrid optical-electrical interconnects is a defining challenge for next-generation cloud computing infrastructure. The combination of electrical and optical resources offers substantial benefits in bandwidth, latency, and energy efficiency, but it also creates complex trade-offs that cannot be resolved by isolated optimization. Schedulers must coordinate workload classification, topology management, wavelength assignment, circuit setup, failure recovery, fairness, and sustainability. They must operate across multiple time scales and organizational boundaries. They must be adaptive enough to exploit changing conditions and stable enough to avoid control-plane churn and cascading failures. They must be efficient enough to justify the cost of hybrid infrastructure and transparent enough to earn the trust of tenants, operators, and regulators.

The central argument of this paper is that hybrid optical-electrical scheduling should be treated as an infrastructure governance problem as much as an algorithmic one. Technical decisions about when to switch a flow to an optical path, which tenant receives priority, and how to respond to failures are also decisions about access, equity, and risk. Sustainable and equitable cloud systems will require schedulers that can reason across technical and social dimensions while remaining explainable and accountable. Future research should develop safe learning-based control, standardized telemetry, carbon-aware scheduling, and participatory governance models. It should also examine how hybrid interconnects interact with edge computing, disaggregated hardware, and emerging applications such as distributed artificial intelligence. By integrating system-level design with policy and sustainability analysis, the field can move toward cloud infrastructure that is not only faster and more efficient but also more robust, fair, and socially legitimate.

References

1. Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., & Zaharia, M. (2010). A view of cloud computing. *Communications of the ACM*, 53(4), 50–58. <https://doi.org/10.1145/1721654.1721672>
2. 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. <https://doi.org/10.2200/S00516ED2V01Y201306CAC024>
3. 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>
4. 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>
5. Kachris, C., & Tomkos, I. (2012). A survey on optical interconnects for data centers. *IEEE Communications Surveys & Tutorials*, 14(4), 1021–1036. <https://doi.org/10.1109/SURV.2011.122111.00069>

6. 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. <https://doi.org/10.1145/1851182.1851223>
7. Chen, K., Singla, A., Singh, A., Ramachandran, K., Xu, L., Zhang, Y., Wen, X., & Chen, Y. (2013). OSA: An optical switching architecture for data center networks with unprecedented flexibility. *IEEE/ACM Transactions on Networking*, 22(2), 498–511. <https://doi.org/10.1109/TNET.2013.2253120>
8. 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.
9. Bazzaz, H. H., Tewari, M., Wang, G., Porter, G., Ng, T. S., Andersen, D. G., Kaminsky, M., Kozuch, M. A., & Vahdat, A. (2013). Switching the optical divide: Fundamental challenges for hybrid electrical/optical datacenter networks. *Proceedings of the ACM Symposium on Cloud Computing*, 1–13. <https://doi.org/10.1145/2523616.2523618>
10. 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. <https://doi.org/10.1109/MCOM.2013.6588648>
11. Shalf, J. (2020). The future of computing beyond Moore's Law. *Philosophical Transactions of the Royal Society A*, 378(2166), 20190061. <https://doi.org/10.1098/rsta.2019.0061>
12. Baliga, J., Ayre, R. W. A., Hinton, K., & Tucker, R. S. (2011). Green cloud computing: Balancing energy in processing, storage, and transport. *Proceedings of the IEEE*, 99(1), 149–167. <https://doi.org/10.1109/JPROC.2010.2060451>
13. Kliazovich, D., Bouvry, P., & Khan, S. U. (2012). GreenCloud: A packet-level simulator of energy-aware cloud computing data centers. *The Journal of Supercomputing*, 62(3), 1263–1283. <https://doi.org/10.1007/s11227-010-0504-1>
14. Heller, B., Seetharaman, S., Mahadevan, P., Yiakoumis, Y., Sharma, P., Banerjee, S., & McKeown, N. (2010). ElasticTree: Saving energy in data center networks. *Proceedings of the 7th USENIX Symposium on Networked Systems Design and Implementation*, 249–264.
15. Buyya, R., Yeo, C. S., Venugopal, S., Broberg, J., & Brandic, I. (2009). Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility. *Future Generation Computer Systems*, 25(6), 599–616. <https://doi.org/10.1016/j.future.2008.12.001>
16. Zhang, Q., Cheng, L., & Boutaba, R. (2010). Cloud computing: State-of-the-art and research challenges. *Journal of Internet Services and Applications*, 1(1), 7–18. <https://doi.org/10.1007/s13174-010-0007-6>
17. Georgakilas, K., Tzanakaki, A., Anastasopoulos, M., & Simeonidou, D. (2016). Converged optical network and data center virtual infrastructure planning. *Journal of Optical Communications and Networking*, 8(7), 459–470. <https://doi.org/10.1364/JOCN.8.000459>

18. Bari, M. F., Boutaba, R., Esteves, R., Granville, L. Z., Podlesny, M., Rabbani, M. G., Zhang, Q., & Zhani, M. F. (2013). Data center network virtualization: A survey. *IEEE Communications Surveys & Tutorials*, 15(2), 909–928.
<https://doi.org/10.1109/SURV.2012.090512.00043>
19. Gai, K., Qiu, M., Zhao, H., & Tao, L. (2016). Dynamic energy-aware cloudlet-based mobile cloud computing model for green computing. *Journal of Network and Computer Applications*, 59, 46–54. <https://doi.org/10.1016/j.jnca.2015.05.016>
20. Cisco. (2023). Cisco annual internet report (2018–2023) white paper. Cisco.
<https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html>