

Time-Sensitive Networking over Low-Latency Optical Data Center Fabrics

Ermenegildo Cadorna

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

Vendela Mohamed

Department of Computer Science, University of Central Florida, Orlando, FL, USA.
vendelamohamed97@ucf.edu

Abstract

The convergence of time-sensitive networking and low-latency optical data center fabrics represents a significant architectural shift in how deterministic communication, massive parallelism, and energy-aware infrastructure are coordinated. Time-sensitive networking offers mechanisms for bounded latency, scheduled traffic, clock synchronization, and resource reservation, but its original deployment context was largely industrial and automotive Ethernet. Data centers, by contrast, are characterized by high fan-in, bursty workloads, multi-tenant isolation, and rapid reconfiguration. Optical fabrics introduce wavelength switching, spatial parallelism, hollow-core fiber, and low-loss propagation, which can reduce latency and energy per bit while also complicating synchronization, scheduling, and fault management. This paper examines the system-level implications of running time-sensitive networking over low-latency optical data center fabrics. It analyzes architectural trade-offs between packet-switched and circuit-switched domains, the role of centralized and distributed control, and the governance challenges that arise when deterministic guarantees must coexist with fair multi-tenant access. The discussion also considers robustness under link and switch failures, sustainability pressures, and policy questions related to transparency, accountability, and infrastructure ownership. The paper argues that optical data center fabrics can extend time-sensitive networking beyond niche industrial use cases, but only if synchronization, scheduling, control-plane semantics, and operational governance are co-designed rather than treated as separate optimization problems. The conclusion outlines research directions for reconfigurable optical determinism, cross-domain standards, and equitable resource allocation in future data center ecosystems.

Keywords

time-sensitive networking; optical data center fabrics; low-latency interconnects; clock synchronization; hollow-core fiber; deterministic networking; governance; sustainability; fairness; robustness.

1. Introduction

The modern data center is no longer merely a warehouse of commodity servers connected by best-effort Ethernet. It is a socio-technical infrastructure that supports financial trading, industrial control, artificial intelligence training, real-time analytics, cloud telephony, and distributed storage. These workloads increasingly require not only high throughput but also predictable latency, low jitter, and bounded tail behavior. Time-sensitive networking has emerged as a family of standards and mechanisms designed to provide such guarantees within

Ethernet networks [1]. Its core promise is that certain classes of traffic can be scheduled, shaped, and protected so that they arrive within specified time windows even under contention. Early deployments concentrated on industrial automation, automotive systems, and professional audio-video bridging, where physical processes impose hard deadlines [2]. More recently, researchers have explored ultra-low-latency networking, deterministic networking, and 5G integration, expanding the conceptual reach of time-sensitive networking into broader networked systems [3]. At the same time, industrial communication research has clarified that deterministic behavior depends on tight coupling among scheduling, synchronization, topology, and configuration [4]. This paper asks what happens when those requirements are projected onto low-latency optical data center fabrics. Optical interconnects can reduce propagation delay, energy per bit, and switching overhead, but they also introduce wavelength contention, circuit setup latency, and new failure modes. The central argument is that time-sensitive networking over optical data center fabrics is not simply a matter of replacing copper links with fiber. It requires rethinking the relationship between determinism and flexibility, between centralized governance and distributed autonomy, and between performance optimization and sustainable operation. The remainder of the paper examines foundations, optical architectures, synchronization, integration trade-offs, governance, robustness, fairness, sustainability, deployment pathways, and future research directions.

2. Foundations of Time-Sensitive Networking

Time-sensitive networking is best understood as a set of extensions to Ethernet that transform a best-effort medium into a scheduled and synchronized communication substrate. The IEEE 802.1Qbv amendment introduced time-aware shaper concepts that allow egress ports to open and close transmission gates according to a global schedule [1]. This enables time-critical frames to traverse a network with bounded latency and low jitter, provided that all bridges share a common time base and that schedules are feasible. The scheduling problem itself is computationally difficult because it must satisfy deadlines, queue constraints, and topology dependencies simultaneously [5]. Research on satisfiability modulo theories has shown that schedule synthesis for time-triggered multi-hop networks can be formulated as a constraint-solving problem, but scalability remains a concern as network size and traffic diversity increase [6]. Timing analysis for rate-constrained traffic further complicates the picture because not all traffic is strictly periodic; some flows require latency bounds while still allowing burstiness and multiplexing [7]. In parallel, data center transport research has developed congestion control mechanisms such as Data Center TCP, which uses explicit congestion notification to maintain low latency and high throughput in bursty environments [8]. Programmable control planes, including OpenFlow and P4, have made it possible to reconfigure forwarding behavior dynamically [9][10]. These developments suggest that time-sensitive networking in data centers cannot be reduced to static schedules alone. It must interact with congestion control, traffic engineering, and programmable data planes. The challenge is that time-sensitive networking assumes a degree of stability and administrative coherence that multi-tenant data centers do not always provide. When tenants, workloads, and virtual networks change rapidly, deterministic schedules must be recomputed or adapted without violating guarantees for existing flows. This tension between determinism and dynamism is the central design problem for any attempt to run time-sensitive networking over optical data center fabrics.

3. Optical Data Center Fabric Architectures

Optical data center fabrics have been proposed as a response to the bandwidth, latency, and energy limits of electrical switching. Surveys of optical interconnects for data centers show that optics can provide high-bandwidth links, low attenuation, and potentially lower energy per bit over distance [11]. Subsequent work has examined optical interconnection networks in data centers and identified challenges in switching granularity, buffering, and cost-effective integration [12]. Hybrid electrical and optical architectures, such as Helios, demonstrated that circuit-switched optical paths can complement packet-switched electrical networks for elephant flows and bulk transfers [13]. More recent reviews of optical technologies for data centers highlight advances in silicon photonics, wavelength-division multiplexing, and integrated transceivers that make optical fabrics more manufacturable and deployable [14]. Fiber-optic transmission research has also shown that the next decades of networking will depend on tighter integration between photonic devices, modulation formats, and network control [15]. In an optical data center fabric, the physical layer may include wavelength-switched paths, spatial-division multiplexing, and low-latency fiber such as hollow-core fiber. These technologies can reduce propagation delay and support reconfigurable topologies, but they also change the assumptions of time-sensitive networking. Packet-level scheduling assumes that frames can be queued and released at precise moments. Optical circuit switching, by contrast, may require path establishment and wavelength assignment before data can flow. If the circuit setup time is not negligible, then the end-to-end latency bound must include not only transmission and propagation but also reconfiguration and synchronization overhead. The architecture therefore must decide which traffic classes remain in the electrical packet domain and which are assigned to optical circuits. This decision has implications for determinism, utilization, fault recovery, and fairness among tenants.

4. Synchronization, Latency, and Clock Distribution

Clock synchronization is the foundation of any time-sensitive network because scheduled gates and time-triggered transmissions depend on a shared notion of time. In optical data center fabrics, synchronization becomes more complex because signals may traverse different wavelengths, paths, and media with different propagation characteristics. A recent demonstration of low-latency wavelength-switched clock-synchronized intra-data center interconnects enabled by hollow core fiber shows that optical paths can be designed to support both switching and synchronization simultaneously [16]. This is important because it suggests that optical determinism is not inherently incompatible with time-sensitive networking. However, the same work also highlights that synchronization must be engineered at the physical and network layers together. Spatial-division multiplexing offers another dimension for increasing capacity and parallelism, but it introduces challenges in skew, crosstalk, and per-core or per-mode synchronization [17]. Hollow-core photonic bandgap fibers promise lower latency and reduced nonlinearity compared with conventional silica fibers, but their deployment requires careful attention to loss, splicing, and long-term reliability [18]. In a time-sensitive optical data center, synchronization domains may need to span electrical switches, optical circuit switches, wavelength converters, and endpoint network interface cards. Each domain may have different jitter characteristics and failure modes. A single grandmaster clock may be insufficient if the optical fabric is partitioned into administrative or physical islands. Hierarchical synchronization, holdover clocks, and redundant time distribution may be necessary. The latency budget must also account for the time required to align schedules across domains. If an optical path is reconfigured, the endpoints may need to renegotiate time slots or adjust guard bands. These requirements increase the complexity of the control plane and make operational governance more important.

Without rigorous synchronization, deterministic guarantees degrade into statistical expectations, which are inadequate for applications such as industrial control, financial settlement, or coordinated artificial intelligence training.

5. Architectural Integration and Structural Trade-offs

Integrating time-sensitive networking with optical data center fabrics requires explicit trade-offs among determinism, utilization, scalability, and cost. Packet-switched electrical networks excel at statistical multiplexing and fine-grained resource sharing, but they struggle to provide hard latency bounds under congestion. Optical circuit switching can provide dedicated paths with predictable latency, but it may waste capacity when traffic is bursty or when circuits remain idle. A hybrid architecture can assign latency-critical, high-priority flows to scheduled electrical paths while assigning bulk or long-lived flows to optical circuits. This division reduces contention for time-sensitive traffic and allows optical paths to be optimized for throughput and energy. However, the boundary between domains must be managed carefully. If traffic crosses between electrical and optical domains, the latency bound must include conversion, queuing, and circuit setup. The control plane must also coordinate schedules across domains. Software-defined networking and programmable data planes can help by providing a unified view of topology, traffic, and policy [9][10]. Yet centralized control introduces a single point of failure and may not scale to thousands of endpoints with microsecond-level requirements. Distributed scheduling may improve resilience but can lead to inconsistent decisions or conflicting reservations. Congestion control mechanisms such as Data Center TCP operate on different timescales and assumptions than time-aware scheduling [8]. Integrating them requires careful prioritization so that congestion control does not undermine deterministic guarantees. Another trade-off concerns buffering. Time-sensitive networking often relies on queues and gate schedules to absorb bursts and enforce deadlines. Optical switches may have limited or no buffering, which shifts the burden to edge devices and can increase end-to-end latency. The architecture must therefore decide where buffering, shaping, and scheduling occur. These choices affect not only performance but also cost, power consumption, and fault tolerance.

6. Control, Governance, and Operational Policy

Governance is a critical but often underappreciated dimension of deterministic networking. In a single administrative domain, time-sensitive networking can be configured through engineering rules, offline schedules, and centralized management. In a multi-tenant data center, however, the infrastructure is shared among cloud providers, enterprises, research consortiums, and latency-sensitive applications. Each tenant may require different guarantees, yet the physical fabric is common. Governance must determine who can reserve time slots, how resources are accounted, and how conflicts are resolved. Standards such as IEEE 802.1Qbv provide technical mechanisms, but they do not prescribe economic or policy models [1]. Industrial perspectives on time-sensitive networking emphasize that deployment success depends on lifecycle management, configuration validation, and interoperability [4]. Ultra-low-latency research similarly notes that deterministic networking must interoperate with broader network management and service-level agreements [3]. In an optical data center, governance becomes more complex because wavelength assignment and circuit provisioning may be controlled by different operational teams. A tenant may have authority over virtual machines and application scheduling but not over optical path establishment. This separation can create accountability gaps when latency violations occur. Policy must therefore define clear interfaces between tenant-level requests and infrastructure-level provisioning. It must

also address transparency. If a tenant cannot observe optical path changes, synchronization events, or maintenance activities, it cannot diagnose performance anomalies. Regulatory and contractual frameworks may need to specify minimum notice periods, audit trails, and remediation procedures. Finally, governance must consider fairness. If premium tenants can reserve optical circuits and scheduled time slots, lower-priority tenants may experience starvation or unpredictable performance. A fair governance model must balance revenue, efficiency, and equitable access.

7. Robustness, Fairness, and Sustainability

Robustness in a time-sensitive optical data center fabric depends on redundancy, fast failure detection, and graceful degradation. Optical links can fail due to fiber cuts, connector contamination, transceiver faults, or wavelength drift. Electrical switches can fail due to software bugs, power loss, or configuration errors. If time-sensitive traffic is scheduled along a single path, a failure can violate deadlines immediately. Protection paths, fast rerouting, and schedule reconfiguration are therefore essential. However, redundancy consumes resources and may reduce the capacity available for revenue-generating traffic. The system must decide which flows require hard protection and which can tolerate temporary degradation. Fairness is closely related to robustness because failure recovery may prioritize some tenants over others. A fair recovery policy should be transparent and consistent with service-level agreements. Sustainability adds another dimension. Data centers consume significant electricity, and global estimates of information and communication technology energy use continue to shape policy debates [20]. Recalibrating data center energy-use estimates has shown that efficiency improvements can offset some growth, but demand for artificial intelligence and cloud services continues to rise [19]. Optical fabrics can reduce energy per bit by lowering switching and transmission losses, but they also require cooling, lasers, and specialized components. Time-sensitive networking can improve energy efficiency by allowing over-provisioned resources to sleep or by consolidating latency-critical traffic onto fewer active paths. Yet strict determinism may require idle guard bands and redundant capacity, which can increase energy use. The sustainability case for optical time-sensitive networking therefore depends on utilization. If optical circuits are underutilized, the energy savings may not materialize. If schedules are too rigid, hardware may remain powered but idle. A holistic approach must integrate energy-aware scheduling, dynamic topology control, and workload placement. Fairness and sustainability also intersect with policy, because energy pricing, carbon accounting, and infrastructure siting affect who bears the environmental cost of low-latency services.

8. Deployment Pathways and Cross-Domain Comparisons

Deployment of time-sensitive networking over optical data center fabrics will likely proceed through incremental pathways rather than a single architectural replacement. One pathway begins in specialized clusters, such as financial trading, industrial edge clouds, or research testbeds, where deterministic latency has clear economic or scientific value. In these environments, operators may accept custom hardware, dedicated fibers, and centralized scheduling. A second pathway begins with optical circuit switching for bulk traffic, gradually adding time-sensitive scheduling for selected flows. A third pathway integrates time-sensitive networking with programmable data planes and software-defined networking, allowing policies to be expressed at higher levels and translated into optical and electrical configurations [9][10]. Cross-domain comparisons are instructive. Industrial automation has long used time-triggered architectures with static schedules and rigorous certification [4].

Telecom networks have developed deterministic networking and segment routing to support service-level guarantees [3]. High-performance computing has explored low-latency interconnects and optical fabrics for tightly coupled workloads [13]. Each domain has different assumptions about ownership, failure tolerance, and traffic patterns. Data centers combine the scale of telecom, the multi-tenancy of cloud computing, and the latency sensitivity of industrial control. This combination makes governance and fairness especially challenging. A deployment pathway that works for a single enterprise may not scale to a public cloud. Interoperability standards will be necessary, but standards alone cannot resolve economic conflicts. Deployment also depends on physical infrastructure. Hollow-core fiber and advanced optical components may not be available in all locations [18]. Retrofitting existing data centers may be more difficult than building new ones. The pace of adoption will therefore vary by region, market segment, and regulatory environment. Cross-domain learning suggests that successful deployments will require close collaboration among network architects, application developers, facility operators, and policymakers.

9. Future Research Directions

Future research should address several open problems. First, synchronization and scheduling must be co-designed for reconfigurable optical fabrics. The required reference demonstrates that wavelength-switched clock-synchronized interconnects are feasible, but more work is needed on scalability, fault tolerance, and multi-domain synchronization [16]. Second, control-plane architectures must balance centralization and distribution. Machine learning and optimization may help generate schedules, but they must be explainable and auditable for governance. Third, fairness metrics for deterministic networking need further development. Latency bounds alone do not capture whether access is equitable or whether small tenants can compete with large ones. Fourth, sustainability models should integrate optical fabric energy consumption, cooling, and embodied carbon. Data center energy estimates provide a starting point, but more granular models are needed [19][20]. Fifth, security and privacy must be considered. Deterministic schedules may leak information about tenant workloads, and optical circuit switching may create new side channels. Sixth, interoperability across standards and vendors remains a practical barrier. Research should test whether time-sensitive networking mechanisms can operate over spatial-division multiplexed fibers and heterogeneous optical switches [17]. Seventh, deployment research should examine organizational factors. Technical performance is necessary but not sufficient; adoption depends on trust, training, and regulatory acceptance. Finally, cross-domain testbeds that combine industrial control, cloud services, and optical fabrics would provide valuable evidence. Such testbeds could reveal whether deterministic guarantees hold under realistic failures, maintenance, and multi-tenant contention.

10. Conclusion

Time-sensitive networking over low-latency optical data center fabrics is a promising but demanding research and engineering direction. Optical interconnects can reduce latency, increase capacity, and improve energy efficiency, while time-sensitive networking can provide bounded latency, low jitter, and predictable service. Bringing them together requires more than incremental optimization. It requires architectural decisions about circuit and packet domains, synchronization across heterogeneous media, control-plane coordination, and governance of shared deterministic resources. The analysis in this paper suggests that the most viable path is a hybrid architecture in which optical circuits serve high-volume and latency-sensitive flows, while electrical packet switching handles fine-grained multiplexing

and dynamic traffic. Such an architecture must be supported by robust synchronization, transparent policy, fair resource allocation, and sustainability-aware scheduling. Failure modes, multi-tenancy, and regulatory pressures mean that technical mechanisms alone will not determine success. Deployment will depend on standards, economic incentives, and organizational trust. Future research should prioritize reconfigurable optical determinism, cross-domain interoperability, fairness metrics, and energy-aware control. If these challenges are addressed, time-sensitive optical data center fabrics could become a foundational infrastructure for applications that require both scale and determinism.

References

1. Institute of Electrical and Electronics Engineers. (2016). IEEE Standard for Local and Metropolitan Area Networks—Bridges and Bridged Networks—Amendment 25: Enhancements for Scheduled Traffic (IEEE Std 802.1Qbv-2015). IEEE.
2. Finn, N. (2018). Introduction to time-sensitive networking. *IEEE Communications Standards Magazine*, 2(2), 22–28.
3. Nasrallah, A., Thyagaturu, A. S., Alharbi, Z., Wang, C., Shao, X., Reisslein, M., & ElBakoury, H. (2019). Ultra-low latency (ULL) networks: The IEEE TSN and IETF DetNet standards and related 5G ULL research. *IEEE Communications Surveys & Tutorials*, 21(1), 88–145.
4. Bello, L. L., & Steiner, W. (2019). A perspective on IEEE time-sensitive networking for industrial communication and automation systems. *Proceedings of the IEEE*, 107(6), 1094–1120.
5. Craciunas, S. S., Oliver, R. S., Chmelik, M., & Steiner, W. (2016). Scheduling real-time communication in IEEE 802.1Qbv time sensitive networks. In *Proceedings of the 24th International Conference on Real-Time Networks and Systems* (pp. 183–192). Association for Computing Machinery.
6. Steiner, W. (2010). An evaluation of SMT-based schedule synthesis for time-triggered multi-hop networks. In *2010 IEEE Real-Time Systems Symposium* (pp. 375–384). IEEE.
7. Zhao, L., Pop, P., & Steiner, W. (2018). Timing analysis of rate-constrained traffic in TSN networks. In *2018 IEEE Real-Time Systems Symposium* (pp. 160–171). IEEE.
8. Alizadeh, M., Greenberg, A., Maltz, D. A., Padhye, J., Patel, P., Prabhakar, B., Sengupta, S., & Sridharan, M. (2010). Data center TCP (DCTCP). *ACM SIGCOMM Computer Communication Review*, 40(4), 63–74.
9. 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.
10. Bosshart, P., Daly, D., Gibb, G., Izzard, M., McKeown, N., Rexford, J., Schlesinger, C., Talayco, D., Vahdat, A., Varghese, G., & Walker, D. (2014). P4: Programming protocol-independent packet processors. *ACM SIGCOMM Computer Communication Review*, 44(3), 87–95.
11. Kachris, C., & Tomkos, I. (2012). A survey on optical interconnects for data centers. *IEEE Communications Surveys & Tutorials*, 14(4), 1021–1036.

12. 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.
13. 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.
14. 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.
15. Winzer, P. J., Neilson, D. T., & Chraplyvy, A. R. (2018). Fiber-optic transmission and networking: The previous 20 and next 20 years. *Optics Express*, 26(18), 24190–24239.
16. Zhou, Z., Dzieciol, H., Clark, K., Luo, Y., Richardson, D., Poletti, F., ... & Liu, Z. (2023). Low-latency wavelength-switched clock-synchronized intra-data center interconnects enabled by hollow core fiber. *Optics Express*, 31(15), 24739–24748.
17. Puttnam, B. J., Rademacher, G., & Luís, R. S. (2021). Space-division multiplexing for optical fiber communications. *Optica*, 8(9), 1186–1203.
18. Richardson, D. J., Wheeler, N. V., Petrovich, M. N., & Poletti, F. (2013). Hollow-core photonic bandgap fibres. *Nature Photonics*, 7(5), 354–362.
19. Masanet, E., Shehabi, A., Lei, N., Smith, S., & Koomey, J. (2020). Recalibrating global data center energy-use estimates. *Science*, 367(6481), 984–986.
20. Andrae, A. S. G., & Edler, T. (2015). On global electricity usage of communication technology: Trends to 2030. *Challenges*, 6(1), 117–157.