

Digital Twin-Based Optimization of Data Center Optical Interconnection Networks

Naomh Farrelly

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

Jacinta Lasa

Department of Computer Science, University of New Hampshire, Durham, NH, USA.
jacintalasa747@unh.edu

Wiktor Kapral

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.
wiktork@uab.edu

Kajetan Irla

Department of Computer Science, University of Central Florida, Orlando, FL, USA.
kajetanirla74@ucf.edu

Abstract

Data center optical interconnection networks are becoming central to the performance, resilience, and energy profile of large-scale digital infrastructure. As workloads diversify across artificial intelligence training, latency-sensitive inference, storage replication, and multi-tenant cloud services, static provisioning and coarse-grained traffic engineering increasingly fail to capture the dynamic coupling between physical optics, packet-level scheduling, workload placement, and facility constraints. This paper examines digital twin-based optimization of data center optical interconnection networks from a system-level perspective. It argues that digital twins can improve optical network operation not merely by simulating components, but by creating a continuously synchronized decision layer that links telemetry, physical models, workload forecasts, and policy constraints. The analysis considers architectural trade-offs among electrical switching, optical circuit switching, wavelength switching, and hybrid fabrics, with attention to observability, control latency, fault isolation, and deployment cost. It then discusses optimization loops that span physical-layer impairment estimation, topology reconfiguration, workload-aware routing, and sustainability-aware scheduling. The paper also explores governance, robustness, fairness, and policy implications, including data sovereignty, accountability for automated decisions, and equitable access to shared infrastructure. Case illustrations from hyperscale, colocation, and edge contexts show that digital twin value depends less on model fidelity alone than on institutional arrangements that determine what may be sensed, changed, and trusted. The conclusion identifies research directions for verifiable twins, multi-domain federation, and human oversight in critical optical infrastructure.

Keywords

digital twin; data center networks; optical interconnects; network optimization; sustainability; robustness; governance; fairness.

1. Introduction

Large-scale data centers have evolved from collections of servers into tightly integrated computational infrastructures whose behavior is shaped by power delivery, cooling, virtualization, software orchestration, and network topology [1]. Within this stack, the optical interconnection network has become a decisive subsystem because it determines how efficiently compute, storage, and accelerator resources can be pooled and reconfigured. Early data center network designs emphasized scalable electrical fabrics and commodity switching, often using folded Clos or leaf-spine topologies to provide predictable bisection bandwidth [2]. Yet the growth of east-west traffic, distributed storage, and accelerator-centric workloads has strained purely electrical approaches, motivating hybrid and all-optical architectures that exploit wavelength division, optical circuit switching, and low-latency optical paths [3]. Measurement studies have long shown that data center traffic is bursty, heterogeneous, and often poorly captured by simple aggregate models, which complicates static capacity planning [4]. Alternative server-centric and modular topologies have sought to improve fault tolerance and incremental growth, but they also introduce complex routing and management requirements [5]. The economic pressure to reduce cost per delivered service further raises the stakes for network optimization because overprovisioning wastes capital and energy while underprovisioning degrades service quality [6].

Digital twins offer a conceptual and operational response to this complexity. A digital twin is not simply a simulation; it is a virtual representation that maintains a meaningful synchronization with a physical asset or system and supports monitoring, analysis, prediction, and control [7]. In industrial settings, digital twins have been used to reduce unpredictable emergent behavior, optimize maintenance, and support lifecycle decisions [8]. Literature reviews distinguish between digital models, digital shadows, and digital twins according to the direction and automation of data exchange, highlighting that the strongest forms of twinning involve bidirectional interaction rather than one-way visualization [9]. For data center optical networks, this distinction matters because a useful twin must ingest physical-layer telemetry, topology state, workload demand, and policy constraints, then recommend or execute reconfigurations without destabilizing production traffic. Advances in reinforcement learning and sequential decision-making provide algorithmic tools for such control, but they also introduce questions of safety, explainability, and operational trust [10], [11]. The low-latency wavelength-switched interconnect research illustrates how physical-layer innovations can reduce synchronization uncertainty and latency, yet such gains must be coordinated with higher-layer control to be realized in production [12]. This paper therefore analyzes digital twin-based optimization as a socio-technical architecture rather than as a purely algorithmic problem.

2. Foundations of Digital Twins in Data Center Optical Interconnection Networks

A digital twin for an optical interconnection network requires a model of the physical layer, the logical topology, and the workload environment. At the physical layer, the twin must represent transceiver behavior, wavelength availability, amplifier characteristics, fiber propagation, connector losses, and impairment accumulation [13]. Optical interconnects for data centers have advanced through parallel optics, silicon photonics, coherent detection, and wavelength switching, each introducing different monitoring and calibration requirements [14]. At the logical layer, the twin must represent switch configurations, routing tables, optical circuit assignments, protection paths, and quality-of-service policies. At the workload layer, it must represent virtual machines, containers, accelerator jobs, storage replication

flows, and service-level objectives. These layers are not independent. A change in wavelength assignment can alter latency and bit error rate, which can affect transport protocols, which can in turn shift traffic patterns and congestion. A digital twin that treats these layers separately may produce locally optimal decisions that are globally harmful.

The fidelity of a digital twin is often discussed in terms of model accuracy, but operational value depends on synchronization latency and actionability. A high-fidelity model updated hourly cannot support rapid protection switching or congestion-aware wavelength reassignment. Conversely, a lightweight model updated in milliseconds may sacrifice physical accuracy but provide sufficient signal for coarse routing decisions. This trade-off resembles the distinction between fast control loops and slow planning loops in autonomic systems. The twin should therefore be decomposed into multiple time scales, with fast shadow models for telemetry validation and fault detection, medium-term models for traffic engineering, and slow models for capacity planning and sustainability optimization. This multi-rate architecture reduces computational burden and makes the twin more governable because different decisions can be audited at different levels of scrutiny.

Data acquisition is a foundational challenge. Optical networks generate telemetry from transceivers, optical channel monitors, switch counters, packet probes, and workload orchestrators. These sources vary in sampling rate, timestamp accuracy, and semantic consistency. Without careful data governance, the twin may learn spurious correlations or perpetuate measurement biases. The problem is not only technical but institutional: different teams may own the optical layer, the packet layer, the compute scheduler, and the facility management system. A digital twin that crosses these boundaries must negotiate data sharing agreements, access controls, and retention policies. In this sense, the twin is also an organizational artifact that makes dependencies visible and negotiable [19]. If governance is weak, the twin may become a dashboard that no operational team trusts or a control system that no team is willing to authorize.

3. Architectural Trade-offs in Optical Interconnection and Digital Twin Instrumentation

The architecture of the optical interconnection network determines what a digital twin can observe and control. Electrical packet switching offers fine-grained statistical multiplexing and mature traffic management, but it consumes significant power and may struggle with very high radix and bandwidth-density requirements [2]. Optical circuit switching can provide high-bandwidth, low-latency paths for elephant flows, but it requires reconfiguration time and careful traffic prediction to avoid underutilization [3]. Wavelength switching adds another degree of freedom by allowing multiple channels to be routed without full circuit reestablishment, but it increases complexity in wavelength assignment and conflict resolution. Hybrid architectures seek to combine the flexibility of electrical switching with the efficiency of optical bypass, yet they create new control problems because the twin must decide which flows should remain in the electrical domain and which should be handed to the optical domain. These decisions depend on flow size, duration, deadline, and tolerance for reconfiguration delay.

Digital twin instrumentation must be co-designed with this architecture. If the twin is expected to optimize wavelength assignment, the network must expose wavelength occupancy, optical signal-to-noise ratio, and transceiver tuning latency. If the twin is expected to optimize workload placement, it must receive job characteristics and service-level objectives from the orchestrator. If the twin is expected to support sustainability goals, it must

access power usage effectiveness, renewable availability, and carbon intensity signals. Instrumentation adds cost, power, and failure modes, so it must be justified by decision value. A common mistake is to collect exhaustive telemetry without a clear control objective. A more disciplined approach is to identify the decisions the twin will support, determine the minimum observability needed for those decisions, and then design the telemetry accordingly. This reduces data volume and makes the twin easier to validate.

Architectural trade-offs also affect robustness. Centralized digital twins can maintain a global view and coordinate optimization across domains, but they introduce a single point of failure and a potential bottleneck. Decentralized or hierarchical twins can improve resilience and reduce latency, but they may produce conflicting local decisions. Federated twins can share learned models without raw telemetry, but they require trust mechanisms and may struggle with heterogeneous data schemas. The choice among these patterns is not purely technical; it reflects governance preferences regarding control, privacy, and accountability. A hyperscale operator may prefer a centralized twin because it owns the entire stack, while a colocation facility with multiple tenants may require a federated approach to preserve tenant boundaries. Edge data centers may use lightweight twins that coordinate with regional twins for capacity planning but make local decisions for latency-sensitive services.

4. Optimization and Control Loops across Physical, Network, and Workload Layers

Digital twin-based optimization can be organized into nested control loops. The fastest loop monitors physical-layer health, detects anomalies, and triggers protection switching or power adjustments. This loop requires low-latency telemetry and deterministic action. A slightly slower loop performs traffic engineering by adjusting routing weights, wavelength assignments, or optical circuit schedules based on predicted demand. A slower planning loop evaluates topology upgrades, transceiver refresh cycles, and capacity expansion. At the slowest scale, the twin supports sustainability and policy planning by comparing energy sourcing, workload shifting, and carbon-aware scheduling. These loops are coupled: a protection switch changes topology, which affects traffic engineering; a traffic engineering decision changes link utilization, which affects failure risk; a capacity plan changes the available degrees of freedom for all faster loops.

Reinforcement learning is attractive for these problems because optical network control involves sequential decisions under uncertainty [11]. A twin can serve as a safe training environment where policies are tested against simulated failures, traffic surges, and configuration changes before deployment [10]. However, simulation-to-reality transfer remains difficult. The physical layer may behave differently under aging, temperature variation, or manufacturing variance. Workload behavior may shift due to external events. A policy that performs well in simulation may fail in production if the twin does not represent rare but consequential conditions. This argues for conservative deployment, human-in-the-loop approval, and fallback mechanisms. The twin should recommend actions with confidence estimates and explanations, not merely output opaque decisions. It should also support counterfactual analysis so operators can ask what would have happened under alternative configurations.

Optimization objectives are inherently multi-objective. Operators may seek to minimize latency, maximize throughput, reduce energy, lower cost, improve fairness among tenants, and maintain resilience. These objectives can conflict. For example, routing all traffic along the shortest optical path may reduce latency but concentrate risk on a few links. Shifting workloads to regions with low carbon intensity may increase network congestion or violate

data sovereignty requirements. A digital twin can help make these trade-offs explicit by allowing operators to adjust weights, constraints, and policies and observe predicted consequences. Yet the twin cannot resolve political or ethical disagreements about which trade-offs are acceptable. Those decisions require governance. The technical system can provide transparency, but it cannot replace legitimate decision-making processes.

5. Deployment, Governance, and Operational Sustainability

Deployment of digital twins in production optical networks is as much an organizational change as a technical upgrade. Operations teams must trust the twin's telemetry, models, and recommendations. This trust is built through validation, incident review, and gradual autonomy. A useful pattern is to begin with a read-only twin that monitors and predicts, then move to a recommend-only twin that suggests actions, and finally to a closed-loop twin that executes low-risk actions under strict constraints. Each stage should have clear metrics for accuracy, safety, and operational benefit. Without such staging, operators may reject the twin after a single incorrect prediction, or they may over-trust it and fail to notice model drift.

Governance frameworks must address data ownership, access control, auditability, and liability. In multi-tenant data centers, the optical network may carry traffic for competing customers. A digital twin that aggregates telemetry across tenants could leak sensitive information about traffic patterns, business cycles, or security incidents. Privacy-preserving techniques such as aggregation, differential privacy, and secure multiparty computation can help, but they add complexity and may reduce model accuracy. Policy must also address cross-border data flows because digital twins may replicate telemetry to centralized cloud regions. Data sovereignty rules may prohibit certain telemetry from leaving a jurisdiction, which affects twin placement and federation [19]. The code that governs these systems is itself a form of regulation, shaping what is possible and what is permissible [20].

Sustainability is a major driver for digital twin adoption. Data centers consume substantial energy, and global estimates have been debated due to differences in scope and methodology [16], [17]. Optical interconnections can improve energy efficiency by reducing the number of electrical switching stages and enabling longer-reach, lower-loss links, but their net benefit depends on transceiver power, cooling, and utilization [18]. A digital twin can support sustainability by correlating workload placement with renewable availability, optimizing optical bypass to reduce switching energy, and identifying stranded capacity. It can also support circular economy goals by modeling component lifetimes and maintenance schedules. However, sustainability optimization must avoid rebound effects. If efficiency gains lower the cost of computation, demand may increase, offsetting energy savings. The twin should therefore be integrated with broader capacity and business planning rather than treated as a purely technical efficiency tool.

6. Robustness, Fairness, and Policy Implications

Robustness in optical interconnection networks involves resilience to fiber cuts, transceiver failures, software bugs, and adversarial attacks. A digital twin can improve robustness by maintaining a real-time model of failure propagation and recovery options. It can simulate cascading failures and identify hidden dependencies. Yet the twin itself can become a target. If an attacker compromises telemetry or model inputs, the twin may issue harmful reconfigurations. Security must therefore be designed into the twin from the beginning, including authenticated telemetry, signed model updates, and anomaly detection. The twin

should also be able to operate in a degraded mode if connectivity to centralized analytics is lost. Local fallback policies and safe default configurations are essential.

Fairness is an emerging concern in shared digital infrastructure. Optical networks often serve multiple tenants with different service-level agreements, budgets, and strategic importance. Optimization that maximizes aggregate throughput may starve small tenants or delay latency-sensitive applications. A digital twin can encode fairness constraints, such as min-max delay, weighted fair sharing, or reserved capacity for critical services. However, fairness definitions are contested. Is fairness about equal outcomes, equal opportunity, or protection of the weakest? The twin can make the consequences of different definitions visible, but it cannot choose among them without legitimate governance. In addition, fairness must be maintained across time. A tenant that is deprioritized during congestion may suffer long-term competitive disadvantage. Policy should therefore include transparency, appeal mechanisms, and periodic review of automated decisions.

Policy implications extend to accountability. If a digital twin autonomously reconfigures an optical network and causes a service outage, who is responsible? The operator, the model developer, the data provider, or the governance body that authorized autonomy? Clear documentation of models, datasets, and decision logs can support accountability [21], [22]. Model cards and datasheets provide structured ways to record intended use, performance characteristics, and limitations. For optical networks, similar documentation should include topology assumptions, failure scenarios tested, telemetry sources, and known blind spots. Regulatory frameworks may eventually require such documentation for critical infrastructure. The challenge is to avoid excessive compliance burden while ensuring that automated systems remain inspectable and contestable.

7. Case Illustrations and Cross-Domain Comparisons

Hyperscale operators provide the most mature context for digital twin-based optical optimization. They own the compute, network, and facility layers, giving them broad observability and authority to reconfigure. Their twins can integrate workload schedulers, optical switches, and power management. The main challenge is scale and complexity. Even a well-designed twin may struggle to model millions of components. Hyperscale operators often use hierarchical twins with regional and local instances. They also invest heavily in validation because a small error can affect many customers. In this context, digital twins are less about radical autonomy and more about incremental optimization, anomaly detection, and capacity planning.

Colocation data centers present a different picture. The operator may control the building and cross-connect infrastructure but not the tenants' servers or workloads. The twin may only observe coarse traffic volumes and power draw. Optimization is constrained by tenant contracts and privacy. Federated learning and confidential computing can help, but they require cooperation. The value of the twin may lie in cross-connect provisioning, fault localization, and energy reporting rather than fine-grained traffic engineering. Edge data centers introduce another set of constraints. They are numerous, lightly staffed, and often connected by heterogeneous links. A centralized twin may be impractical due to latency and connectivity, so lightweight local twins with periodic synchronization are more appropriate. These twins can optimize cooling, local caching, and failover to nearby sites.

Cross-domain comparisons can illuminate design choices. Manufacturing digital twins emphasize asset lifecycle, maintenance, and quality control [8], [9]. Data center optical twins

emphasize real-time traffic and failure response. Healthcare digital twins emphasize privacy and regulation. Energy grid digital twins emphasize stability and market coordination. Each domain shapes the twin's temporal resolution, governance, and risk tolerance. Data center optical networks can learn from grid twins about resilience and from healthcare twins about privacy. They can also contribute to cross-domain methods for multi-rate synchronization, safe reinforcement learning, and policy-aware optimization. The most transferable lesson is that digital twins are not neutral technical objects. They embed assumptions about who can see what, who can decide what, and what counts as success.

8. Future Directions and Research Agenda

Several research directions emerge. First, verifiable digital twins are needed. Operators should be able to prove that a twin's recommendations satisfy safety and policy constraints. This requires formal methods, runtime verification, and interpretable models. Second, multi-domain federation must mature. Optical networks interact with compute, storage, power, and cooling. Federated twins could optimize across domains without centralizing sensitive data. Third, human oversight must be designed rather than added later. Interfaces should show uncertainty, counterfactuals, and affected tenants. Operators need training to recognize model drift and automation bias. Fourth, sustainability metrics must be standardized and integrated with carbon accounting. Fifth, security must address adversarial telemetry and model poisoning. Finally, fairness must be operationalized in ways that are transparent and contestable.

A promising architectural direction is the policy-aware digital twin. Instead of optimizing a single objective, the twin maintains a policy layer that encodes legal, contractual, and ethical constraints. The optimization engine proposes actions, and the policy layer verifies or rejects them. This separation allows policies to be updated without retraining models and makes accountability clearer. Another direction is the use of digital twins for resilience exercises. Operators can simulate fiber cuts, ransomware attacks, and coordinated failures. These exercises can reveal hidden dependencies and improve incident response. They can also build trust in the twin by demonstrating its behavior under stress. Over time, digital twins may become part of the control plane itself, but only if they are robust, governable, and aligned with human values.

9. Conclusion

Digital twin-based optimization of data center optical interconnection networks offers a powerful way to manage complexity across physical optics, packet routing, workload placement, and sustainability constraints. The most important insight is that a digital twin is not merely a model. It is a socio-technical system that depends on telemetry, governance, organizational trust, and policy constraints. Its value comes from linking observation to action in a controlled, auditable manner. Optical networks are particularly challenging because they combine high-speed physical dynamics with multi-tenant policy requirements and significant capital cost. A successful twin must therefore be multi-rate, multi-domain, and policy-aware. It must support human oversight, degrade gracefully, and document its limitations. It must also address fairness, security, and sustainability as first-class objectives rather than afterthoughts. The future of data center optical interconnection will likely involve a mix of electrical and optical switching, hierarchical control, and federated analytics. Digital twins can help orchestrate this mix, but only if they are embedded in legitimate governance structures and subjected to continuous validation. The research agenda should therefore

advance not only model fidelity and optimization algorithms but also accountability, fairness, and resilience in the institutions that deploy them.

References

1. Barroso, L. A., Hölzle, U., & Ranganathan, P. (2019). *The datacenter as a computer: Designing warehouse-scale machines* (3rd ed.). Springer.
2. Al-Fares, M., Loukissas, A., & Vahdat, A. (2008). A scalable, commodity data center network architecture. *ACM SIGCOMM Computer Communication Review*, 38(4), 63–74.
3. Farrington, N., Porter, G., Radhakrishnan, S., Bazzaz, H. H., Subramanya, V., Fainman, Y., Papen, G., & Ford, J. (2010). Helios: A hybrid electrical/optical switch architecture for modular data centers. *ACM SIGCOMM Computer Communication Review*, 40(4), 339–350.
4. 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). ACM.
5. Guo, C., Lu, G., Li, D., Wu, H., Zhang, X., Shi, Y., Tian, C., Zhang, Y., & Lu, S. (2009). BCube: A high performance, server-centric network architecture for modular data centers. *ACM SIGCOMM Computer Communication Review*, 39(4), 63–74.
6. 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.
7. Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In F.-J. Kahlen, S. Flumerfelt, & A. Alves (Eds.), *Transdisciplinary perspectives on complex systems* (pp. 85–113). Springer.
8. Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415.
9. Kritzinger, W., Karner, M., Traar, G., Henjes, J., & Sihn, W. (2018). Digital twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine*, 51(11), 1016–1022.
10. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533.
11. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.
12. 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.
13. 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.

14. Kachris, C., & Tomkos, I. (2012). A survey on optical interconnects for data centers. *IEEE Communications Surveys & Tutorials*, 14(4), 1021–1036.
15. Shalf, J. (2020). The future of computing beyond Moore's Law. *Philosophical Transactions of the Royal Society A*, 378(2166), 20190061.
16. Koomey, J., & Masanet, E. (2020). Does not compute: Avoiding pitfalls in assessing the internet's energy and carbon impacts. *Joule*, 4(1), 21–26.
17. Masanet, E., Shehabi, A., Lei, N., Smith, S., & Koomey, J. (2020). Recalibrating global data center energy-use estimates. *Science*, 367(6481), 984–986.
18. Shuja, J., Gani, A., Shamshirband, S., Ahmad, R. W., & Bilal, K. (2016). Sustainable cloud data centers: A survey of enabling techniques and technologies. *Renewable and Sustainable Energy Reviews*, 62, 195–214.
19. DeNardis, L. (2014). *The global war for internet governance*. Yale University Press.
20. Lessig, L. (2006). *Code: Version 2.0*. Basic Books.
21. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 220–229). ACM.
22. Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92.